Parallel Belief Revision

James Delgrande and **Yi Jin**
School of Computing Science
Simon Fraser University
Burnaby, B.C. V5A 1S6
Canada
{jim,yij}@cs.sfu.ca

Abstract

A recalcitrant problem in approaches to iterated belief revision is that, after first revising by a formula and then by a formula that is inconsistent with the first formula, all information in the original formula is lost. As noted by various researchers, this phenomenon is made explicit in the second postulate (C2) of the well-known Darwiche-Pearl framework, and so this postulate has been a point of criticism of this and related approaches. In contrast, we argue that the true culprit of this problem arises from a basic assumption of the AGM framework, that new information is represented by a single formula. We propose a more general framework for belief revision (called *parallel belief revision*) in which individual items of new information are represented by a set of formulas. In this framework, if one revises by a set of formulas, and then by the negation of some members of this set, then other members of the set are still believed after the revision. Hence the aforecited problem is discharged. We present first a basic approach to parallel belief revision, and next an approach that combines the basic approach with that of Jin and Thielscher. Postulates and semantic conditions characterizing these approaches are given, and representation results provided.

Introduction

Belief revision is the area of knowledge representation concerned with how an agent may incorporate new information about a domain. It is generally accepted that there is no single "best" revision operator, and different agents may have different revision functions. However, revision functions are not arbitrary, but are usually considered as being guided, or characterized, by various *rationality criteria*. The original and best-known set of postulates is called the *AGM postulates* (Alchourrón, Gärdenfors, & Makinson 1985). Subsequently, there has been a great deal of attention paid to *iterated belief revision*, which addresses logical relations among a sequence of (possibly conflicting) observations.

While there has been much progress in the area of iterated belief revision, all such work suffers from the following problem: if one revises by a formula and then by a formula that is inconsistent with this formula, the agent's beliefs are exactly the same as if only the second revision had taken place. For example, consider where an agent receives a report that a particular object is a red bird. If K is the agent's original set of beliefs and $*$ is a revision operator, we can represent the result of this revision by $K*(r \wedge b)$. Then if we next discover that the object is not a bird ($\neg b$), we also lose the information that the object is red; i.e. the beliefs resulting from $K*(r \wedge b)*\neg b$ are exactly the same as those obtained by $K*\neg b$. While this *may* be a desired result in some cases, it certainly shouldn't be a *necessary* outcome. This example can be exaggerated to emphasize the point: Consider where an agent with no initial non-tautological beliefs is presented with a huge number of facts, only to have one of these items subsequently negated. Then all other information is lost, except for the newly-negated item. Clearly, this is too strong a condition to impose on every revision function in all circumstances. We will refer to this as the "drowning problem" of iterated revision, noting that it is quite distinct from a similarly-named problem that arises in some approaches to nonmonotonic reasoning.

Our thesis is that this isn't a problem with these approaches per se, but rather that a more nuanced representation of the item(s) for revision is required. To this end, we develop an account of what we call *parallel belief revision*, in which the second argument to a revision function is a finite set of formulas. Thus, we distinguish $K*\{\alpha \wedge \beta\}$ and $K*\{\alpha, \beta\}$. In the former, revision is by a single formula, and if a subsequent revision contradicts this formula, then belief in this item of information may well be lost. On the other hand, it generally makes sense that α be believed in $K*\{\alpha,\beta\}*\{\neg\beta\}$, since if one element of the input set is contradicted, this need not affect belief in other elements. In this paper then, we develop approaches to parallel belief revision, and show how the aforecited problem is resolved.

In the next section we review the area of belief revision and further motivate our approach. It proves to be the case that parallel belief revision is, in a sense, largely independent of other accounts of iterated belief revision. Hence, we first give an account of the most basic approach to parallel revision. We then show how this approach can be combined with the approach to iterated revision of Jin and Thielscher (2007). (The combination of parallel revision with other approaches is deferred to the full paper.) In each case, postulates characterizing the revision function are given, along with a semantic account, in terms of total orders on sets of worlds, and representation results. We conclude with a comparison to related work.

Background

Formal Preliminaries

We assume a propositional language $\mathcal{L}$ generated from a finite set $\mathcal{P}$ of atomic propositions. The language is that of classical propositional logic, i.e., with the classical consequence relation $\vdash$. $Cn(A)$ is the set of logical consequences of A, that is $Cn(A) = \{\alpha \in \mathcal{L} \mid A \vdash \alpha\}$. $\top$ stands for some arbitrary tautology. For a (finite) set of formulas S, $\wedge S$ is the conjunction of members of S; and $\overline{S} = \{\neg\alpha \mid \alpha \in S\}$. Given two sets of formulas A and B, $A + B$ denotes the *expansion* of A by B, that is $A + B = Cn(A \cup B)$. Expansion of a set A by a formula β is defined analogously. Two sentences α and β are *logically equivalent*, written as $\alpha \equiv \beta$, iff $\alpha \vdash \beta$ and $\beta \vdash \alpha$. A propositional *interpretation* (also referred to as a *possible world*) is a mapping from $\mathcal{P}$ to {true, false}. The set of all interpretations is denoted by $\Theta_{\mathcal{L}}$. A *model* of a sentence α is an interpretation w that makes α true according to the usual definition of truth, and is denoted by $w \models \alpha$. For $\mathcal{W} \subseteq \Theta_{\mathcal{L}}$, we also write $\mathcal{W} \models \alpha$ if $w \models \alpha$ for every $w \in \mathcal{W}$. For a set of sentences A, $Mod(A)$ is the set of all models of A. For simplicity, $Mod(\{\alpha\})$ is also written as $Mod(\alpha)$. Conversely, given a set of possible worlds $\mathcal{W} \subseteq \Theta_{\mathcal{L}}$, we denote by $\mathcal{T}(\mathcal{W})$ the set of sentences which are true in all elements of $\mathcal{W}$, that is $\mathcal{T}(\mathcal{W}) = \{\alpha \in \mathcal{L} \mid w \models \alpha \text{ for all } w \in \mathcal{W}\}$.

A total preorder $\preceq$ (possibly indexed) is a reflexive, transitive binary relation, s.t. either $\alpha \preceq \beta$ or $\beta \preceq \alpha$ for every α, β. As well, $\alpha \prec \beta$ iff $\alpha \preceq \beta$ and $\beta \not\preceq \alpha$. As usual, $\alpha = \beta$ abbreviates $\alpha \preceq \beta$ and $\beta \preceq \alpha$. Given a set S and total preorder $\preceq$ defined on members of S, we denote by $\min(S, \preceq)$ the set of minimal elements of S in $\preceq$.

Belief Revision

In the AGM theory, beliefs of an agent are modelled by a *belief set* K, i.e. a set K such that $K = Cn(K)$. Belief revision is modeled as a function from belief sets and formulas to belief sets. However, various researchers have argued that, in order to address iterated belief revision, it is more appropriate to consider *belief states* (also called *epistemic states*) as objects of revision. A belief state $\mathcal{K}$ effectively encodes preferential information regarding how the revision function itself changes under a revision.[1] The belief set corresponding to belief state $\mathcal{K}$ is denoted $Bel(\mathcal{K})$. Formally, a revision operator $*$ maps a belief state $\mathcal{K}$ and new information α to a revised belief state $\mathcal{K} * \alpha$. Then, in the spirit of (Darwiche & Pearl 1997), the AGM postulates for revision can be reformulated as follows:

($\mathcal{K} * 1$) $Bel(\mathcal{K} * \alpha) = Cn(Bel(\mathcal{K} * \alpha))$

($\mathcal{K} * 2$) $\alpha \in Bel(\mathcal{K} * \alpha)$

($\mathcal{K} * 3$) $Bel(\mathcal{K} * \alpha) \subseteq Bel(\mathcal{K}) + \alpha$

($\mathcal{K} * 4$) If $\neg\alpha \notin Bel(\mathcal{K})$ then $Bel(\mathcal{K}) + \alpha \subseteq Bel(\mathcal{K} * \alpha)$

($\mathcal{K} * 5$) $Bel(\mathcal{K} * \alpha)$ is inconsistent, only if $\nvdash \neg\alpha$

($\mathcal{K} * 6$) If $\alpha \equiv \beta$ then $Bel(\mathcal{K} * \alpha) \equiv Bel(\mathcal{K} * \beta)$

($\mathcal{K} * 7$) $Bel(\mathcal{K} * (\alpha \wedge \beta)) \subseteq Bel(\mathcal{K} * \alpha) + \beta$

($\mathcal{K} * 8$) If $\neg\beta \notin Bel(\mathcal{K} * \alpha)$ then
$Bel(\mathcal{K} * \alpha) + \beta \subseteq Bel(\mathcal{K} * (\alpha \wedge \beta))$

See (Gärdenfors 1988) for motivation and interpretation of these postulates.

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 any belief state $\mathcal{K}$ can induce, as its preferential information, a total preorder on the set of possible worlds. Formally, given a belief state $\mathcal{K}$, a *faithful ranking* on $\mathcal{K}$ is a total preorder $\preceq_{\mathcal{K}}$ on the possible worlds $\Theta_{\mathcal{L}}$, s.t., 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 $w_1 \prec_{\mathcal{K}} w_2$

Intuitively, $w_1 \preceq_{\mathcal{K}} w_2$ if w_1 is at least as plausible as w_2.

It follows directly from the results of (Katsuno & Mendelzon 1991) that a revision operator $*$ satisfies ($\mathcal{K} * 1$)–($\mathcal{K} * 8$) iff there exists a faithful ranking $\preceq_{\mathcal{K}}$ for an arbitrary belief state $\mathcal{K}$, such that for any sentence α:

$$Bel(\mathcal{K} * \alpha) = \begin{cases} \mathcal{L} & \text{if } \vdash \neg\alpha \\ \mathcal{T}(\min(Mod(\alpha), \preceq_{\mathcal{K}})) & \text{otherwise} \end{cases}$$

Iterated Belief Revision

The AGM postulates do not address properties of iterated belief revision. This has led to the development of additional postulates for iterated revision; the best-known approach is that of Darwiche and Pearl (1997) (DP). They propose the following postulates, adapted according to our notation:

C1 If $\beta \vdash \alpha$, then $Bel((\mathcal{K} * \alpha) * \beta) = Bel(\mathcal{K} * \beta)$.

C2 If $\beta \vdash \neg\alpha$, then $Bel((\mathcal{K} * \alpha) * \beta) = Bel(\mathcal{K} * \beta)$.

C3 If $\alpha \in Bel(\mathcal{K} * \beta)$, then $\alpha \in Bel((\mathcal{K} * \alpha) * \beta)$.

C4 If $\neg\alpha \notin Bel(\mathcal{K} * \beta)$, then $\neg\alpha \notin Bel((\mathcal{K} * \alpha) * \beta)$.

Darwiche and Pearl show that an AGM revision operator satisfies Postulates (C1)–(C4) iff the way it revises faithful rankings satisfies the conditions:

CR1 If $w_1, w_2 \models \alpha$, then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K}*\alpha} w_2$.

CR2 If $w_1, w_2 \not\models \alpha$, then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K}*\alpha} w_2$.

CR3 If $w_1 \models \alpha$ and $w_2 \not\models \alpha$, then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K}*\alpha} w_2$.

CR4 If $w_1 \models \alpha$ and $w_2 \not\models \alpha$, then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K}*\alpha} w_2$.

The DP postulates have been criticized in two aspects. On one hand, the DP postulates are too permissive, in that they support revision operators which allow arbitrary dependencies among the items of information which an agent acquires along its way. Consequently, Jin and Thielscher (2007) have proposed the so-called postulate of independence:

Ind If $\neg\alpha \notin Bel(\mathcal{K}*\beta)$ then $\alpha \in Bel((\mathcal{K}*\alpha)*\beta)$

[1] This glosses over a number of issues on the nature of a revision function, which need not concern us here. See (Rott 2001; Nayak, Pagnucco, & Peppas 2003) for more on this issue.

Postulate (Ind) strengthens both (C3) and (C4). Thus, their suggested set of postulates consists of (C1), (C2), and (Ind). They also give necessary and sufficient condition for an AGM revision operator to satisfy (Ind):

IndR If $w_1 \models \alpha$ and $w_2 \models \neg\alpha$, then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K}*\alpha} w_2$.

On the other hand, the DP postulates are too strong. In particular, Postulate (C2) has been accused by many researchers of being responsible for the "drowning" problem (Lehmann 1995; Konieczny & Pino Pérez 2000). As another example, consider a scenario proposed by Konieczny and Pino Pérez (2000):

Example 1. Suppose an electric circuit contains an adder and a multiplier. The atomic propositions a and m denote respectively that the adder and the multiplier are working. Initially we have no information about this circuit, and we then learn that the adder and the multiplier are working ($\alpha = a \wedge m$). Thereafter, someone tells us that the adder is actually not working ($\beta = \neg a$).

As argued in (Konieczny & Pino Pérez 2000), there is obviously no reason to "forget" that the multiplier is working; however by (C2) we must have $(K * \alpha) * \beta = K * \beta$, since $\beta \vdash \neg\alpha$. Hence, in this case (C2) appears to be too strong.

Intuitively, such examples are compelling. However, the case against (C2) isn't clear cut. First, as a technical defense of (C2), it can be observed that many researchers who are against (C2) appear to be in favour of Postulate (C1). However, the semantic characterization of Postulate (C2) (viz. (CR2)) seems as reasonable as that of (C1) (viz. (CR1)): If being informed about α does not change the relative plausibility of α-worlds, why should the relative ordering of $\neg\alpha$-worlds be changed? This idea is also articulated in (Spohn 1988), which argues that it is only reasonable to change the relative ordering between α-worlds and $\neg\alpha$-worlds.

As an informal defense of (C2), it can be observed that in Example 1 it is implicitly assumed that a and m are separate items of information. However, in the AGM approach, the simultaneous revision by a and m is represented by a conjunction. What gets lost is the relation, if any, between these items of information. It could be that the new information should be treated as an undecomposable unit; in this case, the behaviour imposed by (C2) in Example 1 would be perfectly reasonable. Thus, if we are told by someone that both the adder and multiplier are working, and then determine ourselves that the adder is not working, it would make sense to give up *in toto* all information provided by that person.

The above discussion shows that there are at least two situations where the agent can learn several pieces of new information simultaneously: either these pieces of information are to be treated as separate items, or else they together make up an undecomposable item of information. Clearly, accounts of iterated belief revision are not sufficient to deal with both situations. Thus, Example 1 doesn't provide a counterexample to (C2), so much as it highlights the limitations of the expressibility of revision functions. More precisely, it suggests the necessity of generalizing AGM revision functions so that both above-mentioned situations can be handled. This topic is developed in the next section.

Parallel Revision

We have argued that $\mathcal{K} * \{\alpha \wedge \beta\}$ should be treated differently from $\mathcal{K} * \{\alpha, \beta\}$ with respect to iterated revision. However, we would want to relate these two instances, in the simplest case, as follows[2]

For (finite) set of formulas S, $Bel(\mathcal{K} * S) = Bel(\mathcal{K} * \{\wedge S\})$.

Consider the binary case, assuming $\alpha \wedge \beta \not\vdash \bot$:

$$Bel(\mathcal{K} * \{\alpha, \beta\}) = Bel(\mathcal{K} * \{\alpha \wedge \beta\}).$$

On the right hand side of the equality, we revise by a single item of information, $\alpha \wedge \beta$; if this item is shown false (e.g. in later revising by $\neg\beta$) then this information has been contradicted and it is reasonable that all original information (including α) may be lost. Hence, possibly $\alpha \notin Bel(\mathcal{K} * \{\alpha \wedge \beta\} * \{\neg\beta\})$. This argument doesn't apply to $Bel(\mathcal{K} * \{\alpha, \beta\})$, where we revise by a set consisting of two items of information. If one of these elements is subsequently believed to be false then one would nonetheless want to retain the other element where "reasonable".[3]

Semantically, this has the following ramifications. In the faithful ordering resulting from a revision $\mathcal{K} * \{\alpha, \beta\}$, we have that the least $\alpha \wedge \beta$ worlds are ranked lower than the least $\neg(\alpha \wedge \beta)$ worlds in $\preceq_{\mathcal{K}*\{\alpha,\beta\}}$ (since, of course, the least $\alpha \wedge \beta$ worlds are minimal in $\preceq_{\mathcal{K}*\{\alpha,\beta\}}$). The key intuition in parallel revision is that these considerations extend to subsets of the set of formulas for revision. Hence following the revision $\mathcal{K} * \{\alpha, \beta\}$, we will also require that the least $\alpha \wedge \neg\beta$ worlds be ranked below the least $\neg\alpha \wedge \neg\beta$ worlds, and similarly for the least $\neg\alpha \wedge \beta$, $\neg\alpha \wedge \neg\beta$ worlds.

Essentially then, for a revision $\mathcal{K} * S$, changes to the underlying ranking on worlds will depend not just on the set S, but also on subsets of S. In the next subsection we formalize this intuition. The approach is largely independent of previous approaches to iterated revision, and so in the next section we combine the basic approach with that of (Jin & Thielscher 2007) to yield what we suggest is the appropriate general model for iterated belief revision.

The Basic Approach

In this section, we develop the basic approach to parallel revision, wherein new information is represented by a set of formulas. The intuition is that each formula of the set represents an undecomposable (with respect to revision) piece of information. To distinguish this from standard belief revision, we denote a parallel revision operator by $\otimes$. Formally, $\otimes$ maps a belief state $\mathcal{K}$ and set of formulas S to a revised belief state $\mathcal{K} \otimes S$. We assume henceforth that the second argument to $\otimes$ is a *finite* set of formulas.

To begin, we adapt the AGM postulates for parallel revision; the following are analogous to postulates given in (Zhang *et al.* 1997), adapted for belief states.

$(\mathcal{K} \otimes 1)$ $Cn(Bel(\mathcal{K} \otimes \mathrm{S})) = Bel(\mathcal{K} \otimes \mathrm{S})$

[2] In the final section we briefly consider where $S \vdash \bot$, in which case we may want to violate this constraint; however for the present we hew as closely as possible to the standard (AGM) approach.

[3] A case where this wouldn't apply is $\mathcal{K} * \{\alpha, \alpha \wedge \beta\}$ where clearly one requires that $\alpha \wedge \beta \notin Bel(\mathcal{K} * \{\alpha, \alpha \wedge \beta\} * \{\neg\alpha\})$.

($\mathcal{K} \otimes 2$) $S \subseteq Bel(\mathcal{K} \otimes \mathrm{S})$

($\mathcal{K} \otimes 3$) $Bel(\mathcal{K} \otimes \mathrm{S}) \subseteq Bel(\mathcal{K}) + \mathrm{S}$

($\mathcal{K} \otimes 4$) If $Bel(\mathcal{K}) \cup \mathrm{S}$ is consistent, then $Bel(\mathcal{K}) + \mathrm{S} \subseteq Bel(\mathcal{K} \otimes \mathrm{S})$

($\mathcal{K} \otimes 5$) If S is consistent and $S \neq \emptyset$, then $Bel(\mathcal{K} \otimes \mathrm{S})$ is consistent

($\mathcal{K} \otimes 6$) If $S_1 \equiv S_2$, then $Bel(\mathcal{K} \otimes \mathrm{S}_1) = Bel(\mathcal{K} \otimes \mathrm{S}_2)$

($\mathcal{K} \otimes 7$) $Bel(\mathcal{K} \otimes (\mathrm{S}_1 \cup \mathrm{S}_2)) \subseteq Bel(\mathcal{K} \otimes \mathrm{S}_1) + \mathrm{S}_2$

($\mathcal{K} \otimes 8$) If $Bel((\mathcal{K} \otimes \mathrm{S}_1)) \cup \mathrm{S}_2$ is consistent, then
$$Bel((\mathcal{K} \otimes \mathrm{S}_1)) + \mathrm{S}_2 \subseteq Bel(\mathcal{K} \otimes (\mathrm{S}_1 \cup \mathrm{S}_2))$$

Note that ($\mathcal{K} \otimes 6$) yields $Bel(\mathcal{K} \otimes S) = Bel(\mathcal{K} \otimes (\wedge S))$.

The basic approach to parallel revision involves three new postulates, along with three corresponding semantic conditions. The first deals with a limiting case, that of revising by the empty set. For the other two postulates, in the AGM approach one has that, semantically, the set of worlds corresponding to $Bel(\mathcal{K} \otimes S)$ is made up of a set of *least* S-worlds in $\preceq_{\mathcal{K}}$, and that this set is *maximal* in size. The second and third postulates extend these notions to subsets of S.

First, as noted, the set S for revision could be the empty set. In this case, in which there is no input information, the belief state should remain unchanged. We adopt the slightly weaker condition that $\mathcal{K}$ and $\mathcal{K} \otimes \emptyset$ behave identically with respect to further revisions:

($\mathcal{K} \otimes \emptyset$) $Bel((\mathcal{K} \otimes \emptyset) \otimes \mathrm{S}) = Bel(\mathcal{K} \otimes \mathrm{S})$

Note that this means $\mathcal{K} \otimes \emptyset$ and $\mathcal{K} \otimes \{\top\}$ behave differently, specifically when $Bel(\mathcal{K})$ is inconsistent; this is reflected in the additional prerequisite condition to ($\mathcal{K} \otimes 5$). Obviously, ($\mathcal{K} \otimes \emptyset$) corresponds to the following semantical condition:

(PE) $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K} \otimes \emptyset} w_2$

The next postulate characterises the basic approach to parallel revision.

($\mathcal{K} \otimes P$) For $S_1, S_2 \subseteq \mathcal{L}$ where $S_1 \cup S_2 \not\vdash \bot$, and
$$S_1 \cup \overline{S_2} \not\vdash \bot, \text{ we have } S_1 \subseteq Bel((\mathcal{K} \otimes (S_1 \cup S_2) \otimes \overline{S_2}).$$

Hence for a revision of $\mathcal{K}$ by S (here $= S_1 \cup S_2$), with $S_1 \subseteq S$, we have that S_1 is preserved in revising by the negations of $S_2 = S \setminus S_1$. The intuition is that, in revising by S, all elements of S are believed; if some members of S are subsequently disbelieved then, insofar as possible, the remaining members of S are still believed. From a semantic point of view, consider the following condition on a faithful ordering $\preceq_{\mathcal{K}}$, where $w_1, w_2 \in \Theta_{\mathcal{L}}$:

(PR) If $X \subset Y \subseteq S$ and $w_1 \in \min(\overline{X}, \preceq_{\mathcal{K} \otimes S})$ and $w_2 \in \min(\overline{Y}, \preceq_{\mathcal{K} \otimes S})$ then $w_1 \prec_{\mathcal{K} \otimes S} w_2$.

That is, in the AGM approach we have that in $\mathcal{K} \otimes S$, the least S worlds are ranked below the least $\neg(\wedge S)$ worlds in $\preceq_{\mathcal{K} \otimes S}$; the condition PR extends this minimal ranking, insofar as possible, to subsets of S. We obtain the representation result:

Theorem 1. *Let $\otimes$ be a revision operator satisfying Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). Then $\otimes$ satisfies ($\mathcal{K} \otimes \emptyset$) and ($\mathcal{K} \otimes P$) iff it revises faithful rankings according to (PE) and (PR).*

Some examples will make the ramifications of this approach clear. Consider first $\mathcal{K} * S$ where $S = \{a, b, c\}$, and a, b, c are atoms. We obtain that in a faithful ranking resulting from the revision $\preceq_{\mathcal{K} \otimes S}$, the least $\{a, b, c\}$ worlds are strictly less than the least $\{a, b, \neg c\}$ worlds, which in turn are strictly less than the least $\{a, \neg b, \neg c\}$ worlds.

We get:

$$\begin{array}{rcl} a \wedge c & \in & Bel(\mathcal{K} * \{a, b, c\} * \{\neg b\}) \\ a & \in & Bel(\mathcal{K} * \{a, b, c\} * \{\neg b, \neg c\}) \\ a & \in & Bel(\mathcal{K} * \{a, b, c\} * \{\neg b\} * \{\neg c\}) \\ a & \in & Bel(\mathcal{K} * \{a, b, c\} * \{\neg b \vee \neg c\}) \end{array}$$

The final result follows from the factoring result in AGM revision; in fact it is straightforward to show that, assuming the antecedent conditions in ($\mathcal{K} \otimes P$), that

$$S_1 \subseteq Bel((\mathcal{K} \otimes (S_1 \cup S_2) \otimes \{\neg(\wedge S_2)\}).$$

Clearly, for atoms a, b, we don't generally obtain that $b \in Bel(\mathcal{K} * \{a, a \vee b\})$. However, we do obtain:

$$b \in Bel(\mathcal{K} * \{a, a \vee b\} * \neg a).$$

Thus in this case we establish a preference between a and b, to the effect of "accept a, but if it is subsequently found to be false, accept b."

Essentially then, in revising by a set S, we pay attention to not just the minimal S elements in the faithful ordering, but also to the minimal element of subsets of S. With this in mind, we adopt the last postulate to characterize the basic approach. The idea is that in revising by set S, for $S_1 \subset S$, the minimal $S_1 \cup (\overline{S \setminus S_1})$ worlds should be the same before and after revising by S. This is given in the next postulate:

($\mathcal{K} \otimes S$) For $S_1 \subseteq S$, $S \not\vdash \bot$, $\quad Bel(\mathcal{K} \otimes (S_1 \cup (\overline{S \setminus S_1})))$
$$= Bel(\mathcal{K} \otimes S \otimes (S_1 \cup (\overline{S \setminus S_1}))).$$

We have the corresponding condition on a faithful ordering:

(PS) For $S_1 \subset S$ and $S \not\vdash \bot$,
$$\min(S_1 \cup (\overline{S \setminus S_1}), \preceq_{\mathcal{K}}) = \min(S_1 \cup (\overline{S \setminus S_1}), \preceq_{\mathcal{K} \otimes S}).$$

We obtain the representation result.

Theorem 2. *Let $\otimes$ be a revision operator satisfying Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). Then $\otimes$ satisfies ($\mathcal{K} \otimes \emptyset$) and ($\mathcal{K} \otimes S$) iff it revises faithful rankings according to (PE) and (PS).*

Parallel Revision and Iterated Revision

The basic approach only deals with limited situations where we first revise by a set of formulas then by the negations of some of these formulas. In this section, we extend the basic approach to deal with more general cases. We first show that the straightforward generalization of the well-known iterated revision postulates are problematic and insufficient. Then, we will present a postulate of *evidence retainment*, which offers a better solution to the drowning problem.

We start with the following generalization of the DP postulates as suggested by (Zhang 2004).

(C1$^{\otimes}$) If $S_2 \vdash S_1$, then $Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2) = Bel(\mathcal{K} \otimes \mathrm{S}_2)$.

(C2$^{\otimes}$) If $S_1 \cup S_2$ is inconsistent, then $Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2) = Bel(\mathcal{K} \otimes \mathrm{S}_2)$

(C3$^\otimes$) If $S_1 \subseteq Bel(\mathcal{K} \otimes \mathrm{S}_2)$, then $S_1 \subseteq Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$.

(C4$^\otimes$) If $S_1 \cup Bel(\mathcal{K} \otimes \mathrm{S}_2)$ is consistent, then $S_1 \cup Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$ is also consistent.

We remark that, while (C1$^\otimes$), (C3$^\otimes$) and (C4$^\otimes$) still seem as reasonable as their counterparts, (C2$^\otimes$) is not desirable. First, previous criticisms of (C2) apply equally well to (C2$^\otimes$). Second, (C2$^\otimes$) is clearly inconsistent with ($\mathcal{K} \otimes P$) in the presence of the (adapted to sets) AGM postulates. Hence, we do not consider (C2$^\otimes$) further as a general postulate for parallel revision.

For reference, the semantical conditions for the DP postulates can be generalized as follows:

(C1$^\otimes$R) If $w_1, w_2 \models S$, then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K} \otimes S} w_2$

(C2$^\otimes$R) If $w_1, w_2 \not\models S$, then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K} \otimes S} w_2$

(C3$^\otimes$R) If $w_1 \models S$ and $w_2 \not\models S$, then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K} \otimes S} w_2$

(C4$^\otimes$R) If $w_1 \models S$ and $w_2 \not\models S$, then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K} \otimes S} w_2$

To show (C2$^\otimes$) is undesirable from another perspective, one may argue that (C2$^\otimes$R) is overly strong: in the case where w_2 satisfies significantly more sentences of S than w_1, it is perfectly reasonable that $w_2 \prec_{\mathcal{K} \otimes S} w_1$ even if $w_1 \preceq_{\mathcal{K}} w_2$.

Similarly, we can also generalize the postulate of independence and its corresponding semantical condition:

(Ind$^\otimes$) If $S_1 \cup Bel(\mathcal{K} \otimes \mathrm{S}_2)$ is consistent, then $S_1 \subseteq Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$

(Ind$^\otimes$R) If $w_1 \models S$ and $w_2 \not\models S$, then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K} \otimes S} w_2$

Note that, among the above-mentioned postulates, (C2$^\otimes$) is the only one which deals with the case where S_1 and S_2 are jointly inconsistent. This suggest that we need some new postulates in order to address the "drowning" problem. As already argued, it is too radical to give up all formulas of S_1 (as imposed by (C$^\otimes$2)) just because $S_1 \cup S_2$ is inconsistent. The problem is, what formulas in S_1 should be retained? Intuitively, a formula $\alpha \in S_1$ should be kept if there is no evidence (in S_1) against α after learning S_2. To formalize this idea, we need the following definition:

Definition 1. Let S_1, S_2 be two sets of sentences. We denote by $S_1||S_2$ the set of all subsets of S_1 that are consistent with S_2, that is $S \in S_1||S_2$ iff:

1. $S \subseteq S_1$
2. $S \cup S_2$ is consistent

Formally, the fact that there exists evidence in S_1 against α after learning S_2 (given the original belief state $\mathcal{K}$) can be expressed as: $\exists S \in S_1||S_2 \ \neg\alpha \in Bel(\mathcal{K} \otimes (\mathrm{S} \cup \mathrm{S}_2))$.

Based on above discussion, we obtain the so-called postulate of *evidence retainment*:[4]

(Ret$^\otimes$) If $\alpha \in S_1$ and $\alpha \notin Bel((\mathrm{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$, then $\exists S \in S_1||S_2$ such that $\neg\alpha \in Bel(\mathrm{K} \otimes (\mathrm{S} \cup \mathrm{S}_2))$

[4]This postulate is inspired by the postulate of *core retainment* (Hansson 1997), which says a formula α is removed from a belief set K by a contraction with β only if there is some evidence in K that shows that α contributes to the implication of β.

Equivalently, the postulate of evidence retainment can be rephrased as follows:

(Ret$^\otimes$) If $\alpha \in S_1$ and $\forall S \in S_1||S_2 \ \neg\alpha \notin Bel(\mathrm{K} \otimes (\mathrm{S} \cup \mathrm{S}_2))$, then $\alpha \in Bel((\mathrm{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$.

Recall Example 1 with $S_1 = \{a, m\}$ and $S_2 = \{\neg a\}$. Since $S_1||S_2 = \{\{m\}\}$, Postulate (Ret$^\otimes$) implies that $(\mathcal{K} \otimes S_1) \otimes S_2 \vdash m$, which gives us the desired result. Note that, in case a and m make up a single piece of information (i.e. $S_1 = \{a \wedge m\}$) Postulate (Ret$^\otimes$) does not apply.

To give a formal justification for (Ret$^\otimes$), we will show a representation theorem.

Definition 2. Let S be a set of sentences and w a possible world. Then $S|w$ denotes the set of all element of S which are true in w, i.e., $S|w = \{\alpha \in S \mid w \models \alpha\}$.

The following theorem gives a necessary and sufficient semantical condition for (Ret$^\otimes$):

Theorem 3. *Suppose $\otimes$ is a parallel revision operator satisfies Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). Then $\otimes$ satisfies (Ret$^\otimes$) iff it revises faithful rankings in the following manner:*

(Ret$^\otimes$R) *If $S|w_2 \subset S|w_1$, then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K} \otimes S} w_2$*

Arguably, (Ret$^\otimes$R) is very natural and intuitive. It essentially says: if w_1 confirms more new information (in S) than w_2, and w_1 is at least as plausible as w_2, then w_1 becomes more plausible than w_2 after revising by S. It is not difficult to see that (Ret$^\otimes$) implies (Ind$^\otimes$) and ($\mathcal{K} \otimes P$).

Proposition 4. *Suppose $\otimes$ is a parallel revision operator satisfying Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). Then $\otimes$ satisfies (Ret$^\otimes$) only if it satisfies (Ind$^\otimes$), ($\mathcal{K} \otimes P$).*

Based on similar considerations, we present two additional postulates which also seem quite intuitive, and which naturally extend (C3$^\otimes$) and (C4$^\otimes$).

(PC3$^\otimes$) If $\forall S \in S_1||S_2 \ \alpha \in Bel(\mathcal{K} \otimes (\mathrm{S} \cup \mathrm{S}_2))$, then $\alpha \in Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$

(PC4$^\otimes$) If $\forall S \in S_1||S_2 \ \neg\alpha \notin Bel(\mathcal{K} \otimes (\mathrm{S} \cup \mathrm{S}_2))$, then $\neg\alpha \notin Bel((\mathcal{K} \otimes \mathrm{S}_1) \otimes \mathrm{S}_2)$

Essentially, (PC3$^\otimes$) says if all evidences in S_1 supports α after learning S_2, then α must be believed; (PC4$^\otimes$) says if no evidence in S_1 is against α, then there is no reason to believe $\neg\alpha$. We present a representation theorem for (PC3$^\otimes$) and (PC4$^\otimes$) as the formal justification.

Theorem 5. *Suppose $\otimes$ is a parallel revision operator satisfying Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). Then $\otimes$ satisfies (PC3$^\otimes$) and (PC4$^\otimes$) iff it revises faithful rankings in the following manner:*

(PC3$^\otimes$R) *If $S|w_2 \subseteq S|w_1$, then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K} \otimes S} w_2$*

(PC4$^\otimes$R) *If $S|w_2 \subseteq S|w_1$, then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K} \otimes S} w_2$*

It can be observed that (PC3$^\otimes$) and (PC4$^\otimes$) extend (C3$^\otimes$) and (C4$^\otimes$), respectively.

Proposition 6. *Suppose $\otimes$ is a parallel revision operator satisfying Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). Then $\otimes$ satisfies (PC3$^\otimes$) and (PC4$^\otimes$) only if it satisfies (C3$^\otimes$) and (C4$^\otimes$), respectively.*

Moreover, the semantical conditions of (PC3$^\otimes$) and (PC4$^\otimes$) requires that the relative ordering of two possible worlds remain unchanged, provided they satisfy the same subset of the new information.

Proposition 7. *(PC3$^\otimes R$) and (PC3$^\otimes R$) imply the following semantical condition:*

If $S|w_2 = S|w_1$, then $w_1 \preceq_K w_2$ iff $w_1 \preceq_{K \otimes S} w_2$

Clearly (PC3$^\otimes$) and (PC4$^\otimes$) imply postulate ($\mathcal{K} \otimes \emptyset$). As well, it is not difficult to see that (PC3$^\otimes$) and (PC4$^\otimes$) together imply (C1$^\otimes$) and ($\mathcal{K} \otimes S$).

Proposition 8. *Suppose $\otimes$ is a parallel revision operator satisfying Postulates ($\mathcal{K} \otimes 1$)–($\mathcal{K} \otimes 8$). If $\otimes$ satisfies (PC3$^\otimes$) and (PC4$^\otimes$) then it also satisfies (C1$^\otimes$) and ($\mathcal{K} \otimes S$).*

Based on the above development, we suggest a general parallel revision operator should satisfy the AGM postulates (extended to sets), (Ret$^\otimes$), (PC3$^\otimes$), and (PC4$^\otimes$).

Conclusions and Related Work

We have developed an account of *parallel belief revision*, in which the second argument to a revision function is a finite set of formulas. Each formula of the set represents an individual item of new information, whereas multiple parts of a single item are represented by a conjunction. We then showed how the drowning problem of iterated revision is addressed, by presenting several postulates for parallel belief revision. In our account we present both a basic approach, consisting of three new postulates, and a "preferred" account of iterated parallel revision, also consisting of three new postulates, and wherein these postulates imply the three basic postulates. In all cases, corresponding semantic conditions are given and representation results derived.

The idea of the changing of an agent's beliefs in light of a set of formulas isn't very new. (Fuhrmann & Hansson 1994) proposes *package contraction*, which is concerned with removing a set of formulas from a belief set. On the other hand, the general contraction introduced by (Zhang *et al.* 1997) studies how to contract a belief set so that it is consistent with a set of formulas. However, neither of these papers are concerned with iterated belief change per se. As well, (Nayak 1994) anticipates some of the properties of parallel revision, in an approach where both the belief state and input are represented by epistemic entrenchment relations.

Syntactically, our approach resembles *set revision* (Zhang & Foo 2001) or *multiple revision* (Peppas 2004). There are two main differences. First, our focus is on iterated revision, and in particular constraints that need to be imposed on an agent's underlying epistemic state in order to effect plausible revisions. Second, (Zhang & Foo 2001) and (Peppas 2004) mainly study infinite sets. Therefore, set revision or multiple revision might be useful for investigating infinite non-monotonic reasoning, our approach is more suitable for modelling the evolution of an agent's belief state.

Our account of parallel revision is intended to extend the AGM approach. In particular, revising by an inconsistent set of formulas yields an inconsistent belief set. For future work, an obvious, interesting, and (we believe) straightforward extension is to combine our approach with that of a *merging* operator. Given some "reasonable" merging operator Δ, one could relate merging and revision as follows:

$$Bel(\mathcal{K} \otimes S) = Bel(\mathcal{K} \otimes (\Delta S)).$$

Hence in this case, one would expect that $Bel(\mathcal{K} \otimes \{\alpha, \neg\alpha, \beta\})$ would be consistent and entail β, while $Bel(\mathcal{K} \otimes \{\alpha \wedge \neg\alpha \wedge \beta\})$ would of course be inconsistent. In this way, merging operators would be employed in the service of revision, in order to obtain consistent revisions in some cases where the input is inconsistent.

Last, as indicated in an example in the basic approach, parallel revision may be used to encode preferences over the revision formulas; a second, intriguing direction for future research is to further explore this phenomenon.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change. *JSL* 50(2):510–530.

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

Fuhrmann, A., and Hansson, S. 1994. A survey of multiple contractions. *JLLI* 3:39–76.

Gärdenfors, P. 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press.

Hansson, S. 1997. *A Textbook of Belief Dynamics*. Kluwer Academic Publishers.

Jin, Y., and Thielscher, M. 2007. Iterated belief revision, revised. *AIJ* 171(1):1–18.

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

Konieczny, S., and Pino Pérez, R. 2000. A framework for iterated revision. *JANCL* 10(3-4).

Lehmann, D. 1995. Belief revision, revised. In *Proc. IJCAI*, 1534–1540. Morgan Kaufmann.

Nayak, A. C.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *AIJ* 146(2):193–228.

Nayak, A. 1994. Iterated belief change based on epistemic entrenchment. *Erkenntnis* 41:353–390.

Peppas, P. 2004. The limit assumption and multiple revision. *JLC* 14(3):355–371.

Rott, H. 2001. *Change, Choice and Inference*. Oxford: Clarendon Press.

Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume II. Kluwer Academic Publishers. 105–134.

Zhang, D., and Foo, N. 2001. Infinitary belief revision. *JPL* 30(6):525–570.

Zhang, D.; Chen, S.; Zhu, W.; and Chen, Z. 1997. Representation theorems for multiple belief changes. In *Proc. IJCAI*, 89–94. Morgan Kaufmann.

Zhang, D. 2004. Properties of iterated multiple belief revision. *Proc. LPNMR*, 314–325. Springer.