Permissive Belief Revision

Maria R. Cravo, João P. Cachopo, Ana C. Cachopo, and João P. Martins

Instituto Superior Técnico
Dep. Eng. Informática, GIA
Av. Rovisco Pais
1049-001 Lisboa — Portugal
{mrcravo,jcachopo,acardoso,jpm}@gia.ist.utl.pt

Abstract. We propose a new operation of belief revision, called permissive belief revision. The underlying idea of permissive belief revision is to replace the beliefs that are abandoned by traditional theories with weaker ones, entailed by them, that still keep the resulting belief set consistent. This framework allows us to keep more beliefs than what is usual using existent belief base-based revision theories.

1 Introduction

In this paper we define a new kind of belief revision. We call it permissive belief revision, and its main advantage over traditional belief revision operations is that more beliefs are kept after revising a set of beliefs. To achieve this result, permissive revision takes the beliefs abandoned by some traditional belief revision operation, and weakens them, adding their weakened versions to the result of the traditional operation. In this way, "some parts" of the abandoned beliefs are still kept.

Throughout the article we use the following notation: lower case greek letters (α, β, ...) represent meta-variables that range over single formulas; lower case roman letters (a, b, ...) represent single atomic formulas; upper case roman letters (A, B, ...) represent sets of formulas; $\mathcal{L}$ represents the language of classical logic (either propositional or first-order logic).

In Section 2 we briefly describe the work in belief revision that is relevant for the understanding of this article. In Section 3 we give some motivations, and an example, that will provide a better understanding of what is gained with permissive revision. After this, in Sections 4 and 5, we formally define this operation, and present some examples. In Section 6 we prove some properties about permissive revision, and show that it satisifies suitable counterparts for the AGM postulates. Finally, in Section 7 we discuss some relevant issues about our theory, in Section 8 we make a comparison with other approaches and in Section 9 we point out some directions in which the present work may evolve.

2 Belief revision

One of the main sources of inspiration in belief revision, the AGM theory, follows the work of [1]. This theory deals with deductively closed sets of sentences, called sets of beliefs. According to the AGM theory, there are three operations on sets of beliefs: expansions, contractions, and revisions.

The AGM theory presents a drawback from a computational point of view, since it deals with infinite sets of beliefs. Both [8, 9] and [6] modified AGM by working with a finite set of propositions, called a *belief base*, B, and using the set of consequences of B, defined as $Cn(B) = \{\phi : B \vdash \phi\}$.[1]

We also define permissive revision on finite sets of beliefs. The traditional revision of a *consistent* belief base B with a formula ϕ, represented by $(B * \phi)$, consists in changing B in such a way that it contains ϕ and is consistent (if ϕ is consistent). The case of interest is when $B \cup \{\phi\}$ is inconsistent, because, otherwise, ϕ can just be added to B.

To perform the revision $(B * \phi)$ when $B \cup \{\phi\}$ is inconsistent, we have to remove some belief(s) from B, before we can add ϕ. In other words, in a revision $(B * \phi)$ some belief(s) *must* be discarded from B.

3 Motivations

The idea of permissive revision is to transform the beliefs that were discarded in a traditional revision into weaker versions and to add them to the result of the revision. Permissive revision, thus, corresponds to a "smaller" change in beliefs than traditional revision, while keeping the goal of having a consistent result.

Conjunctions are the most obvious candidates to be weakened. This aspect was already recognized by [7], who discussed that revision theories sometimes require to give up too many beliefs, without providing a solution to the problem. While Lehmann only presents the problem regarding conjunctions, we argue that this problem is more general and that it can arise with other kinds of formulas.

To illustrate the main idea behind the weakening of conjunctions, suppose, for instance, that some traditional revision operation provides the result:

$$(\{a \wedge b, a \Rightarrow c\} * \neg c) = \{a \Rightarrow c, \neg c\}$$

Permissive revision, represented by $\circledast$, weakens the abandoned formula, $a \wedge b$ to b, and adds this to the result of traditional revision:

$$(\{a \wedge b, a \Rightarrow c\} \circledast \neg c) = \{b, a \Rightarrow c, \neg c\}$$

4 Formalization

By now, it should be clear that the main task in defining permissive revision is the definition of a function Wk, which weakens the formula that was removed during traditional revision. Actually, since there may be more than one such formula, we consider the conjunction of all the removed formulas, and weaken it into a new formula which will then be added to the result of traditional revision to obtain permissive revision.

[1] $\vdash$ represents the classical derivability operation.

The function Wk will have different definitions, depending on whether we are using classical logic, a non-monotonic logic or some other logic. In this article, we restrict ourselves to classical first order logic.

Weakening a formula depends, naturally, on the set of formulas into which we will be adding the result. Therefore, the function Wk will depend on the formula to weaken and a set of formulas:

$$\mathrm{Wk} : \mathcal{L} \times 2^{\mathcal{L}} \rightarrow \mathcal{L}$$

$\mathrm{Wk}(\phi, W)$ can be interpreted as "Weaken the formula ϕ, in such a way that after the weakened formula is added to W, the resulting set is not inconsistent".

Given such a function, we can formally define the permissive revision of a set of formulas W with a formula ϕ, $(W \circledast \phi)$. Let *Abandoned* be the conjunction of all the formulas which were abandoned during the traditional revision of W with ϕ, $\mathit{Abandoned} = \bigwedge(W - (W * \phi))$. Then, the permissive revision of W with ϕ is given by

$$(W \circledast \phi) = (W * \phi) \cup \{\mathrm{Wk}(\mathit{Abandoned}, (W * \phi))\}$$

Let us now see how a formula is weakened. Obviously, this depends on the type of formula in question. The example in the previous section conveys the main ideas behind weakening conjunctions. However, there are other logical symbols besides conjunctions. Considering the usual logical symbols, $\{\neg, \Rightarrow, \wedge, \vee, \exists, \forall\}$, we have the following definition for Wk.[2]

$$\mathrm{Wk}(\phi, W) = \begin{cases} \phi & \text{if } W \cup \{\phi\} \text{ is consistent} \\ \mathrm{WkN}(\phi, W) & \text{if } \phi \text{ is a negation} \\ \mathrm{WkI}(\phi, W) & \text{if } \phi \text{ is an implication} \\ \mathrm{WkD}(\phi, W) & \text{if } \phi \text{ is a disjunction} \\ \mathrm{WkC}(\phi, W) & \text{if } \phi \text{ is a conjunction} \\ \mathrm{WkE}(\phi, W) & \text{if } \phi \text{ is an existential rule} \\ \mathrm{WkU}(\phi, W) & \text{if } \phi \text{ is a universal rule} \\ \top & \text{otherwise} \end{cases}$$

Note that, although Wk will only be used, in the context of permissive revision, to weaken a formula ϕ known to be inconsistent with W, the weakening process is recursive (on the structure of formulas), and there may be sub-formulas which are consistent with W. That's the reason for the first case. As for the last case, which means that ϕ is an atomic formula inconsistent with W, there is no weaker formula we can give than a valid formula.

[2] This definition of Wk has some steps similar to the conversion to Conjunctive-Normal-Form, and it could be simpler if the knowledge base were required to be in a canonical form (CNF for instance). However, the syntactic differences between two logically equivalent formulas are important from the knowledge representation point of view.

Next, we define each of the weakening functions mentioned above. We should keep in mind that a good weakening function should allow us to keep as much information as possible. In order to do that for non-atomic formulas, we weaken each sub-formula and combine the results.

When $\phi = \neg\alpha$, for some atomic formula α, there is nothing we can retain of the weakening of ϕ. However, if α is a non-atomic formula, $a \vee b$, for instance, we can apply logical transformations to ϕ to bring to the surface a kind of formula we know how to handle. In this case $\neg(a \vee b)$ is logically equivalent to $(\neg a) \wedge (\neg b)$.

$$\mathrm{WkN}(\phi, W) = \begin{cases} \mathrm{Wk}(\neg\alpha \wedge \neg\beta, W) & \text{if } \phi = \neg(\alpha \vee \beta) \\ \mathrm{Wk}(\neg\alpha \vee \neg\beta, W) & \text{if } \phi = \neg(\alpha \wedge \beta) \\ \mathrm{Wk}(\alpha \wedge \neg\beta, W) & \text{if } \phi = \neg(\alpha \Rightarrow \beta) \\ \mathrm{Wk}(\alpha, W) & \text{if } \phi = \neg\neg\alpha \\ \mathrm{Wk}(\forall(x)\neg\alpha(x), W) & \text{if } \phi = \neg\exists(x)\alpha(x) \\ \mathrm{Wk}(\exists(x)\neg\alpha(x), W) & \text{if } \phi = \neg\forall(x)\alpha(x) \\ \top & \text{otherwise} \end{cases}$$

Weakening an implication is treated in a similar way, transforming the implication into the logically equivalent disjunction, and weakening the result instead.

$$\mathrm{WkI}(\alpha \Rightarrow \beta, W) = \mathrm{Wk}(\neg\alpha \vee \beta, W)$$

If $\phi = \alpha \vee \beta$, and it is inconsistent with W (otherwise WkD would not be used), then both α and β are inconsistent with W. So, to weaken ϕ we have to individually weaken both α and β, in W, and combine the results with the disjunction again.

$$\mathrm{WkD}(\alpha \vee \beta, W) = \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$$

Conjunction seems to be a more complex case. To help understand its definition we present some examples. First, consider the set $W = \{a \wedge b\}$ and its revision with $\neg a$. Using permissive revision, we use $\mathrm{Wk}(a \wedge b, \{\neg a\})$ and expect it to give b. We just have to abandon one of the elements of the conjunction and keep the other. However, if each element is itself a non-atomic formula, the contradiction may be deeper inside in either one or in both of the elements of the conjunction. For instance, given $W = \{(a \wedge b) \wedge (c \wedge d)\}$ and revising it with $\neg(b \wedge c)$ we would like to get $(a \wedge (c \wedge d)) \vee ((a \wedge b) \wedge d)$, i.e., if it's not possible to have both b and c, then we would like to have either a, b and d or a, c and d. This is the result of $\mathrm{WkC}((a \wedge b) \wedge (c \wedge d), \{\neg(b \wedge c)\})$, according to the following definition.

$$\begin{aligned} \mathrm{WkC}(\alpha \wedge \beta, W) = {} & (\mathrm{Wk}(\alpha, W) \wedge \mathrm{Wk}(\beta, W \cup \{\mathrm{Wk}(\alpha, W)\})) \vee \\ & (\mathrm{Wk}(\beta, W) \wedge \mathrm{Wk}(\alpha, W \cup \{\mathrm{Wk}(\beta, W)\})) \end{aligned}$$

Handling existentially quantified formulas will be done through skolemization, weakening the formula which results from the elimination of the existential quantifier.

$$\mathrm{WkE}(\exists(x)\alpha(x), W) = \mathrm{Wk}(\alpha(p), W), \quad \text{where } p \text{ is a Skolem constant}$$

Finally, the result of weakening universally quantified formulas is just $\top$. This means that, in what concerns this kind of formula, permissive revision brings nothing new. In Section 9, we discuss some alternatives to the weakening of universally quantified formulas.

$$\mathrm{WkU}((\forall x)\alpha(x), W) = \top$$

5 Examples

In this section we present some examples, to illustrate permissive revision. In all the examples we present, permissive revision keeps more beliefs than traditional revision. Of course, this is not always the case. Sometimes both revisions give the same result.

Example 1 (Weakening of conjunctions). In the first situation both conjuncts are inconsistent with the result of traditional revision; in the second situation only one of the conjuncts is inconsistent; and in the third situation none of the conjuncts by itself is inconsistent, only the conjunction of them causes the inconsistency.

1. **Both conjuncts are inconsistent**

$$W = \{a \wedge (b \wedge c), a \Rightarrow d, b \Rightarrow d\}$$

suppose

$$(W * \neg d) = \{a \Rightarrow d, b \Rightarrow d, \neg d\}$$

then

$$\begin{aligned}
&\mathrm{Wk}(a \wedge (b \wedge c), (W * \neg d)) = \\
&\quad = (\mathrm{Wk}(a, (W * \neg d)) \wedge \mathrm{Wk}(b \wedge c, (W * \neg d) \cup \{\mathrm{Wk}(a, (W * \neg d))\})) \vee \\
&\quad\quad (\mathrm{Wk}(b \wedge c, (W * \neg d)) \wedge \mathrm{Wk}(a, (W * \neg d) \cup \{\mathrm{Wk}(b \wedge c, (W * \neg d))\})) \\
&\quad = (\top \wedge \mathrm{Wk}(b \wedge c, (W * \neg d))) \vee (\mathrm{Wk}(b \wedge c, (W * \neg d)) \wedge \top) \\
&\quad = \mathrm{Wk}(b \wedge c, (W * \neg d)) \\
&\quad = (\mathrm{Wk}(b, (W * \neg d)) \wedge \mathrm{Wk}(c, (W * \neg d) \cup \{\mathrm{Wk}(b, (W * \neg d))\})) \vee \\
&\quad\quad (\mathrm{Wk}(c, (W * \neg d)) \wedge \mathrm{Wk}(b, (W * \neg d) \cup \{\mathrm{Wk}(c, (W * \neg d))\})) \\
&\quad = (\top \wedge c) \vee (c \wedge \top) \\
&\quad = c
\end{aligned}$$

and

$$(W \circledast \neg d) = \{a \Rightarrow d, b \Rightarrow d, \neg d, c\}$$

Note that in the traditional revision we can no longer derive c, but this is still a consequence of the permissive revision.

2. **Only one of the conjuncts is inconsistent**

$$W = \{a \wedge b, a \Rightarrow c\}$$

suppose

$$(W * \neg c) = \{a \Rightarrow c, \neg c\}$$

then

$$\begin{aligned}
&\mathrm{Wk}(a \wedge b, (W * \neg c)) = \\
&\quad = (\mathrm{Wk}(a, (W * \neg c)) \wedge \mathrm{Wk}(b, (W * \neg c) \cup \{\mathrm{Wk}(a, (W * \neg c))\})) \vee \\
&\quad\quad (\mathrm{Wk}(b, (W * \neg c)) \wedge \mathrm{Wk}(a, (W * \neg c) \cup \{\mathrm{Wk}(b, (W * \neg c))\})) \\
&\quad = (\top \wedge \mathrm{Wk}(b, (W * \neg c) \cup \{\top\})) \vee (b \wedge \mathrm{Wk}(a, (W * \neg c) \cup \{b\})) \\
&\quad = (\top \wedge b) \vee (b \wedge \top) \\
&\quad = b
\end{aligned}$$

and

$$(W \circledast \neg c) = \{a \Rightarrow c, \neg c, b\}$$

Like before, we keep more beliefs than traditional revision, namely b.

3. **None of the conjuncts by itself is inconsistent**

$$W = \{a \wedge b, (a \wedge b) \Rightarrow c\}$$

suppose

$$(W * \neg c) = \{(a \wedge b) \Rightarrow c, \neg c\}$$

then

$$\begin{aligned}
&\mathrm{Wk}(a \wedge b, (W * \neg c)) = \\
&\quad = (\mathrm{Wk}(a, (W * \neg c)) \wedge \mathrm{Wk}(b, (W * \neg c) \cup \{\mathrm{Wk}(a, (W * \neg c))\})) \vee \\
&\quad\quad (\mathrm{Wk}(b, (W * \neg c)) \wedge \mathrm{Wk}(a, (W * \neg c) \cup \{\mathrm{Wk}(b, (W * \neg c))\})) \\
&\quad = (a \wedge \mathrm{Wk}(b, (W * \neg c) \cup \{a\})) \vee (b \wedge \mathrm{Wk}(a, (W * \neg c) \cup \{b\})) \\
&\quad = (a \wedge \top) \vee (b \wedge \top) \\
&\quad = a \vee b
\end{aligned}$$

and

$$(W \circledast \neg c) = \{(a \wedge b) \Rightarrow c, \neg c, a \vee b\}$$

Example 2 (Weakening of disjunctions). We now present one example of weakening a disjunction. Obviously, both disjuncts are inconsistent with the result of traditional revision, otherwise the disjunction would not be inconsistent. Furthermore, the disjuncts are both non-atomic, otherwise the result would be $\top$.

$$W = \{(a \wedge b) \vee (c \wedge d), b \Rightarrow e, d \Rightarrow e, a \Rightarrow f, c \Rightarrow f\}$$

suppose

$$(W * \neg e) = \{b \Rightarrow e, d \Rightarrow e, a \Rightarrow f, c \Rightarrow f, \neg e\}$$

then

$$\begin{aligned} &\mathrm{Wk}((a \wedge b) \vee (c \wedge d), (W * \neg e)) = \\ &\quad = \mathrm{WkD}((a \wedge b) \vee (c \wedge d), (W * \neg e)) \\ &\quad = \mathrm{Wk}(a \wedge b, (W * \neg e)) \vee \mathrm{Wk}(c \wedge d, (W * \neg e)) \\ &\quad = \mathrm{WkC}(a \wedge b, (W * \neg e)) \vee \mathrm{WkC}(c \wedge d, (W * \neg e)) \\ &\quad = a \vee c \end{aligned}$$

and

$$(W \circledast \neg e) = \{a \vee c, b \Rightarrow e, d \Rightarrow e, a \Rightarrow f, c \Rightarrow f, \neg e\}$$

Note that in the traditional revision we can no longer derive, for instance f, but this is still a consequence of the permissive revision.

Example 3 (Weakening an existentially quantified formula).

$$W = \{\exists(x)a(x) \wedge b(x), \forall(x)a(x) \Rightarrow c(x)\}$$

suppose

$$(W * \forall(x)\neg c(x)) = \{\forall(x)a(x) \Rightarrow c(x), \forall(x)\neg c(x)\}$$

then

$$\begin{aligned} &\mathrm{Wk}(\exists(x)a(x) \wedge b(x), (W * \forall(x)\neg c(x))) = \\ &\quad = \mathrm{WkE}(\exists(x)a(x) \wedge b(x), (W * \forall(x)\neg c(x))) \\ &\quad = \mathrm{Wk}(a(p) \wedge b(p), (W * \forall(x)\neg c(x))) \\ &\quad = \mathrm{WkC}(a(p) \wedge b(p), (W * \forall(x)\neg c(x))) \\ &\quad = b(p) \end{aligned}$$

where p is a Skolem constant and

$$(W \circledast \forall(x)\neg c(x)) = \{\forall(x)a(x) \Rightarrow c(x), \forall(x)\neg c(x), b(p)\}$$

In words, permissive revision, unlike traditional revision, allows us to keep the belief $\exists(x)b(x)$.

6 Properties

We now prove two essential properties of the Wk function. By essential properties, we mean that it would be unacceptable for the Wk function not to satisfy them. The first property ensures that we don't produce an inconsistent set when we add the result of weakening a formula to the result of the traditional revision. The second property ensures that we are not able to derive new conclusions from the result of weakening a formula, that were not derivable from the formula itself.

The next theorem guarantees the first of these properties.

Theorem 1 *Let W be a consistent set of formulas, and ϕ any formula. Then $W \cup \{\mathrm{Wk}(\phi, W)\}$ is consistent.*

Proof. If ϕ is consistent with W, then $\mathrm{Wk}(\phi, W) = \phi$ and the result follows trivially. Otherwise, we will prove by induction on the structure of the formula ϕ that the weakening function produces a formula consistent with W.

If ϕ is a literal (an atomic formula or the negation of an atomic formula) or a universally quantified formula, then $\mathrm{Wk}(\phi, W) = \top$, and therefore $W \cup \{\top\}$ is consistent, provided that W is consistent.

The cases where ϕ is of the form $\neg\alpha$ or $\alpha \Rightarrow \beta$, reduce to one of the other cases, since the weakening of ϕ in these cases reduces to the weakening of a logically equivalent formula, with either a quantifier, a disjunction or a conjunction.

Assume that α, β and $\gamma(p)$, where p is some constant, are formulas that verify the theorem. Since $W \cup \{\mathrm{Wk}(\gamma(p), W)\}$ is consistent by hypothesis, then $W \cup \{\mathrm{Wk}(\exists(x)\gamma(x), W)\}$ is also consistent, by definition of WkE. Accordingly, given that $W \cup \{\mathrm{Wk}(\alpha, W)\}$ is consistent, and, therefore, $W \cup \{\mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)\}$ is consistent, we prove that $W \cup \{\mathrm{Wk}(\alpha \vee \beta, W)\}$ is also consistent. Finally, let $W' = W \cup \{\mathrm{Wk}(\alpha, W)\}$, which, as we have seen, is consistent. Since, by hypothesis, $W' \cup \{\mathrm{Wk}(\beta, W')\}$ is consistent, i.e., $W \cup \{\mathrm{Wk}(\alpha, W), \mathrm{Wk}(\beta, W')\}$ is consistent, we have that $W \cup \{\mathrm{Wk}(\alpha, W) \wedge \mathrm{Wk}(\beta, W')\}$ is consistent, from where it follows trivially that $W \cup \{\mathrm{Wk}(\alpha \wedge \beta, W)\}$ is consistent, which finishes our proof. □

Theorem 2 guarantees that the result of weakening a formula is not stronger than the original formula, i.e., we do not introduce new beliefs.

Theorem 2 *Let W be a set of formulas, and ϕ any formula. Then $\phi \vdash \mathrm{Wk}(\phi, W)$.*

Proof. If $\phi \vdash \bot$ then $\phi \vdash \psi$ for every formula ψ, and in particular for $\psi = \mathrm{Wk}(\phi, W)$. If ϕ is consistent with W, then $\mathrm{Wk}(\phi, W) = \phi$ and, obviously, $\phi \vdash \phi = \mathrm{Wk}(\phi, W)$. Otherwise, as above, we will prove by induction on the structure of the formula ϕ that the weakening function produces a formula not stronger than the original.

The structure of this proof is similar to the previous one: if ϕ is a literal or a universally quantified formula, then $\mathrm{Wk}(\phi, W) = \top$, and $\phi \vdash \top$; if ϕ is of the form $\neg\alpha$ or $\alpha \Rightarrow \beta$, the weakening of ϕ reduces to the weakening of a logical equivalent formula, with either a quantifier, a disjunction or a conjunction.

By eliminating the existential quantifier, we have that $\exists(x)\gamma(x) \vdash \gamma(p)$ for some Skolem constant p. By hypothesis, $\gamma(p) \vdash \mathrm{Wk}(\gamma(p), W) = \mathrm{Wk}(\exists(x)\gamma(x), W)$, and, therefore, $\exists(x)\gamma(x) \vdash \mathrm{Wk}(\exists(x)\gamma(x), W)$.

Assume that α and β are formulas that verify the theorem. Given that, by hypothesis, $\alpha \vdash \mathrm{Wk}(\alpha, W)$, then $\alpha \vdash \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$, and, likewise, since $\beta \vdash \mathrm{Wk}(\beta, W)$ then $\beta \vdash \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$. Joining the two, we have that $\alpha \vee \beta \vdash \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$, i.e., $\alpha \vee \beta \vdash \mathrm{Wk}(\alpha \vee \beta, W)$. To finish the proof, let's see that conjunction preserves the theorem: from $\alpha \vdash \mathrm{Wk}(\alpha, W)$ and $\beta \vdash \mathrm{Wk}(\beta, W \cup \{\mathrm{Wk}(\alpha, W)\})$, it follows trivially that $\alpha \wedge \beta \vdash \mathrm{Wk}(\alpha \wedge \beta, W)$. $\square$

Although we don't consider it essential, we now prove another theorem that will be needed when we prove the satisfaction of the AGM postulates for our theory. The theorem says that the results of weakening a formula, with respect to two logically equivalent sets, are the same.

Theorem 3 *Let W and W' be two sets of formulas, such that $Cn(W) = Cn(W')$, and ϕ any formula. Then* $\mathrm{Wk}(\phi, W) = \mathrm{Wk}(\phi, W')$.

Proof. The only use of W in the definition of Wk is to check whether $W \cup \{\phi\}$ is consistent. Since $Cn(W) = Cn(W')$, $W \cup \{\phi\}$ is consistent, iff $W' \cup \{\phi\}$ is consistent. $\square$

We now show that permissive revision satisfies the AGM postulates for revision, if the traditional revision satisfies these postulates. Since these postulates refer to a revision operation on belief sets (theories or closed sets), and permissive revision was defined on bases (finite sets), the first thing to do is to define a corresponding permissive revision on theories.

Let W be a base, T the theory generated by W, i.e., $T = Cn(W)$, and ϕ a formula. The permissive revision of the theory T with the formula ϕ, $(T \circledast_T \phi)$, is defined by

$$(T \circledast_T \phi) = Cn(W \circledast \phi)$$

Note that this definition implies that $(T \circledast_T \phi)$ depends not only on T and ϕ, but also on the base W from which T was generated.

We recall that permissive revision is defined in terms of a traditional revision, $*$, by:

$$(W \circledast \phi) = (W * \phi) \cup \{\mathrm{Wk}(\mathit{Abandoned}, (W * \phi))\}$$

Before we prove anything on permissive revision, let us assume that the traditional revision on bases, $*$, satisfies suitable counterparts to the AGM postulates.

$(*1)$	$(W * \phi)$ is a base
$(*2)$	$\phi \in (W * \phi)$
$(*3)$	$(W * \phi) \subseteq W \cup \{\phi\}$
$(*4)$	If $\neg\phi \notin Cn(W)$, then $W \cup \{\phi\} \subseteq (W * \phi)$
$(*5)$	$(W * \phi)$ is inconsistent, iff $\neg\phi \in Cn(\emptyset)$
$(*6)$	If $\phi \Leftrightarrow \psi \in Cn(\emptyset)$, then $(W * \phi) - \{\phi\} = (W * \psi) - \{\psi\}$

Of these postulates, only postulate $(*6)$ is not a straightforward counterpart to the corresponding AGM postulate. The straightforward counterpart would be

$$\text{If } \phi \Leftrightarrow \psi \in Cn(\emptyset), \text{ then } (W * \phi) = (W * \psi)$$

Since we are dealing with bases, and not closed sets, it is not reasonable to expect such a result (unless, of course, ϕ and ψ are not only equivalent, but also the same formula). What is reasonable to assume is that the revisions of the same base with two equivalent (but different) formulas, only differ between them in these formulas.

If we now define a traditional revision on theories as

$$(T *_T \phi) = Cn(W * \phi)$$

where T is the theory generated by the base W, $T = Cn(W)$, it is trivial to show that the AGM postulates are satisfied by $*_T$.

$(*_T1)$	$(T *_T \phi)$ is a theory, i.e., $(T *_T \phi) = Cn(T *_T \phi)$
$(*_T2)$	$\phi \in (T *_T \phi)$
$(*_T3)$	$(T *_T \phi) \subseteq Cn(T \cup \{\phi\})$
$(*_T4)$	If $\neg\phi \notin T$, then $Cn(T \cup \{\phi\}) \subseteq (T *_T \phi)$
$(*_T5)$	$(T *_T \phi)$ is inconsistent iff $\neg\phi \in Cn(\emptyset)$
$(*_T6)$	If $\phi \Leftrightarrow \psi \in Cn(\emptyset)$, then $(T *_T \phi) = (T *_T \psi)$

Now that we have established the postulates satisfied by the traditional revision on which permissive revision is based, we prove the following theorem.

Theorem 4 *Let $\circledast_T$ be a permissive revision on theories defined as before:*

$$(T \circledast_T \phi) = Cn(W \circledast \phi).$$

Then, for any theory T (generated from a base W), and any formulas ϕ and ψ, $\circledast_T$ satisfies the AGM postulates.

$(\circledast_T1)$	$(T \circledast_T \phi)$ *is a theory, i.e.,* $(T \circledast_T \phi) = Cn(T *_T \phi)$
$(\circledast_T2)$	$\phi \in (T \circledast_T \phi)$
$(\circledast_T3)$	$(T \circledast_T \phi) \subseteq Cn(T \cup \{\phi\})$
$(\circledast_T4)$	*If* $\neg\phi \notin T$, *then* $Cn(T \cup \{\phi\}) \subseteq (T \circledast_T \phi)$
$(\circledast_T5)$	$(T \circledast_T \phi)$ *is inconsistent, iff* $\neg\phi \in Cn(\emptyset)$
$(\circledast_T6)$	*If* $\phi \Leftrightarrow \psi \in Cn(\emptyset)$, *then* $(T \circledast_T \phi) = (T \circledast_T \psi)$

Proof.
($\circledast_T 1$) $(T \circledast_T \phi)$ is a theory, i.e., $(T \circledast_T \phi) = Cn(T *_T \phi)$.
By definition of $\circledast_T$.

($\circledast_T 2$) $\phi \in (T \circledast_T \phi)$.
By definition of $\circledast_T$ and $(*2)$.

($\circledast_T 3$) $(T \circledast_T \phi) \subseteq Cn(T \cup \{\phi\})$.
By definition of $\circledast_T$ and $\circledast$

$$(T \circledast_T \phi) = Cn(W \circledast \phi) = Cn((W * \phi) \cup \{\mathrm{Wk}(\mathit{Abandoned}, (W * \phi))\})$$

By $(*3)$ $((W * \phi) \subseteq W \cup \{\phi\})$, Theorem 2 ($\mathit{Abandoned} \vdash \mathrm{Wk}(\mathit{Abandoned}, W)$), and monotonicity

$$Cn((W * \phi) \cup \{\mathrm{Wk}(\mathit{Abandoned}, (W * \phi))\}) \subseteq Cn(W \cup \{\phi\} \cup \{\mathit{Abandoned}\})$$

By definition of *Abandoned*, $\mathit{Abandoned} = \bigwedge(W - (W * \phi))$, we have that $W \vdash \mathit{Abandoned}$, so

$$Cn(W \cup \{\phi\} \cup \{\mathit{Abandoned}\}) = Cn(W \cup \{\phi\})$$

Finally, since

$$Cn(W \cup \{\phi\}) = Cn(Cn(W) \cup \{\phi\}) = Cn(T \cup \{\phi\})$$

our proof of ($\circledast_T 3$) is complete.

($\circledast_T 4$) If $\neg\phi \notin T$, then $Cn(T \cup \{\phi\}) \subseteq (T \circledast_T \phi)$.

$$Cn(T \cup \{\phi\}) = Cn(Cn(W) \cup \{\phi\}) = Cn(W \cup \{\phi\})$$

By $(*4)$, since $\neg\phi \notin T$, we have that $W \cup \{\phi\} \subseteq (W * \phi)$. This, together with monotonicity, implies that

$$Cn(W \cup \{\phi\}) \subseteq Cn(W * \phi)$$

Now, if $\neg\phi \notin T$, then, by definition, $\mathit{Abandoned} = \bigwedge\{\} = \top$, and $Wk(\mathit{Abandoned}, (W * \phi)) = \top$. So, $(W \circledast \phi) = (W * \phi)$, and $(T \circledast_T \phi) = Cn(W * \phi)$.

($\circledast_T 5$) $(T \circledast_T \phi)$ is inconsistent, iff $\neg\phi \in Cn(\emptyset)$.
By definition,

$$(T \circledast_T \phi) = Cn(W \circledast \phi) = Cn((W * \phi) \cup \{\mathrm{Wk}(\mathit{Abandoned}, (W * \phi))\}).$$

If $\neg\phi \in Cn(\emptyset)$, then, by $(*5)$, $(W * \phi)$ is inconsistent, and so is $(T \circledast_T \phi)$.
If $\neg\phi \notin Cn(\emptyset)$, then, by $(*5)$, $(W * \phi)$ is consistent, and, by Theorem 1 so is $(W * \phi) \cup \{\mathrm{Wk}(\mathit{Abandoned}, (W * \phi))\}$.

($\circledast_T 6$) If $\phi \Leftrightarrow \psi \in Cn(\emptyset)$, then $(T \circledast_T \phi) = (T \circledast_T \psi)$.
By definition,

$$(T \circledast_T \phi) = Cn(W \circledast \phi) = Cn((W * \phi) \cup \{\mathrm{Wk}(Abandoned_\phi, (W * \phi))\}),$$

where $Abandoned_\phi = \bigwedge(W - (W * \phi))$, and

$$(T \circledast_T \psi) = Cn(W \circledast \psi) = Cn((W * \psi) \cup \{\mathrm{Wk}(Abandoned_\psi, (W * \psi))\}),$$

where $Abandoned_\psi = \bigwedge(W - (W * \psi))$.
Since, by ($*6$), $(W * \phi) - \{\phi\} = (W * \psi) - \{\psi\}$, we have that

$$Abandoned_\phi = Abandoned_\psi$$

Since, by ($*_T 6$), $Cn(W * \phi) = Cn(W * \psi)$, by Theorem 3, we have that

$$\mathrm{Wk}(Abandoned_\phi, (W * \phi)) = \mathrm{Wk}(Abandoned_\psi, (W * \psi))$$

which ends our proof. □

7 Discussion

Traditional belief revision theories [9, 6] may produce different results when revising logically equivalent theories with the same formula, i.e., they are syntax-dependent. For example, the fact that both a and b are true may be represented either by $\{a \wedge b\}$ or by $\{a, b\}$. These two representations will provide different results when revised with $\neg a$. This should be expected, since we are dealing with syntax-based approaches and finite belief sets. However, *in this example*, if permissive revision is applied to weaken the formulas removed by traditional revision, the result will be the same, $\{b\}$. This allows us to conclude that, in some cases, the syntax-dependency of traditional approaches is nullified by permissive revision. Furthermore, the weakening function by itself is syntax-dependent, as the following example shows: $\mathrm{Wk}(a \vee b \Rightarrow c, \{a, \neg c\}) = \top$, but $\mathrm{Wk}((a \Rightarrow c) \wedge (b \Rightarrow c), \{a, \neg c\}) = b \Rightarrow c$. Again, such a behaviour should be expected: the weakening function completely relies on the syntax of the formula to be weakened. If this dependency was considered a flaw, the use of a canonical form for the formula to be weakened would very easily eliminate it. However, since the weakening function is applied to the result of a traditional theory, which is syntax-dependent, it wouldn't make much sense. Theorem 3, on the other hand, shows that this function is *not* dependent on the syntax of the set in respect to which a formula is weakened.

A preliminary report of the work presented in this article appears in [4]. The present article contains, in addition to the preliminary report, Theorem 3 and the proof of the AGM postulates.

8 Comparison with other approaches

To the best of our knowledge, when we submitted this article no similar work had been done, except for our preliminary report [4]. The work presented in [2] however, also aims at minimizing the loss of information by weakening information involved in conflicts rather than completely removing it. We will first convey the main ideas behind their work, and then present some comments on the comparison of both approaches.

In [2] it is assumed that the available information is given as an ordered knowledge base (KB), i.e., a ranking of information as logical sentences: $KB = \{S_1, S_2, \ldots, S_n\}$. When revising a KB with a formula ϕ, they start with $i = 1$ and $KB = \{\phi\}$; for each S_i, if it is consistent with KB then $KB \leftarrow KB \cup S_i$. Otherwise, all possible disjunctions (of the formulas in conflict) of size 2 are computed. If they are consistent with KB then they are added to KB. Otherwise, all possible disjunctions of size 3 are computed, and so on.

One major difference between both approaches is that [2] is a "complete" revision operation, while ours can be applied to the result of any traditional revision operation.[3] So, our theory is also more permissive in the sense that it allows any traditional theory to choose the formulas to weaken.

The only example in [2] has its knowledge base in clausal form, although this does not seem to be a requirement. If we convert the examples in this paper to clausal form, both approaches produce exactly the same results in all the examples. Further work is needed to prove whether this is always so. However, if we use our original examples, the results are not same. Actually, since in all our examples only one formula is removed, the work of [2] simply discards that formula, while ours weakens it.

9 Future work

As we saw in Section 4, universal rules are weakened to $\top$, which is obviously too drastic a solution. This aspect can be improved in two directions. When considering a monotonic logic, a universal rule can be weakened following the general ideas presented in Section 4. For instance, if we have $\forall(x)a(x) \land b(x)$, and revise this with $\neg a(p)$, the universal rule must be abandoned, but it can be weakened to $\forall(x)b(x)$.

In another direction, i.e., when considering a non-monotonic logic, the most natural way of weakening a universal rule is to turn it into the "corresponding" default rule. Of course, defining the exact meaning of "corresponding" default rule will depend on the particular non-monotonic logic being considered, but we can state this informally as turning a universal like "All As are Bs" into the default "Typically, As are Bs". See [10] for an approach to this problem, using Default Logic.

We intend to implement permissive revision on top of SNePSwD [5]. The fact that this system is a belief revision system, with an underlying non-monotonic

[3] Our approach could even be applied to the result of theirs.

logic [3] will be particularly helpful. This system already has mechanisms for determining the consistency of belief sets, and keeping a record of inconsistent sets, which will be necessary for the implementation of the weakening function.

10 Acknowledgements

This work was partially supported by IDMEC—Institute of Mechanical Engineering, Intelligent Systems Centre.

References

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

[2] Salem Benferhat, Souhila Kaci, Daniel Le Berre, and Mary-Anne Williams. Weakening conflicting information for iterated revision and knowledge integration. In *Proceedings of IJCAI-2001*, Seattle, Washington, USA, 2001. Morgan Kaufmann Publishers, Inc.

[3] Maria R. Cravo. SWMC: A logic for default reasoning and belief revision (a new version). Technical Report GIA 93/02, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal, November 1993.

[4] Maria R. Cravo, João P. Cachopo, Ana C. Cachopo, and João P. Martins. Permissive belief revision (preliminary report). "Third Workshop on Nonmonotonic Reasoning, Action, and Change" (NRAC99), "Sixteenth International Joint Conference on Artificial Intelligence" (IJCAI99), August 1999.

[5] Maria R. Cravo and João P. Martins. SNePSwD, a newcomer to the SNePS family. *Journal of Experimental and Theoretical Artificial Intelligence*, 5:135–148, 1993.

[6] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20(2):175–203, 1991.

[7] Daniel Lehmann. Belief revision revised. In *Proceedings of IJCAI-95*, pages 1534–1540, Montreal, Canada, 1995. Morgan Kaufmann Publishers, Inc.

[8] Bernhard Nebel. A knowledge level analysis of belief revision. In *Proceedings of KR-89*, pages 301–311. Morgan Kaufmann Publishers, Inc., Toronto, Canada, 1989.

[9] Bernhard Nebel. *Reasoning and revision in hybrid representation systems*, volume 422 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Heidelberg, Germany, 1990.

[10] C.F. Pimentel and M. R. Cravo. Permissive belief revision using default logic. Technical Report GIA 01/01, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal, 2001.