A New Approach to Base Revision

Paolo Di Giusto[1] Guido Governatori[2]

[1] Regione Emilia-Romagna, v.le A.Moro 52, 40122 Bologna, Italy,
e-mail coc0788@iperbole.bologna.it
[2] CIT, Griffith University, Nathan QLD 4111, Australia
e-mail G.Governatori@cit.gu.edu.au

Abstract. We present three approaches to revision of belief bases, which are also examined in the case in which the sentences in the base are partitioned between those which can and those which cannot be changed; the approaches are shown to be semantically equivalent. A new approach is then presented, based on the modification of individual rules, instead of deletion. The resulting base is semantically equivalent to that generated by the other approaches, in the sense that it has the same models, but the rule part alone has fewer models, that is, is subjected to a smaller change.

Introduction

Belief revision faces the problem of maintaining the consistency of a system of beliefs when new pieces of information are added, and, at the same time, it should preserve as many beliefs compatible with the new data as possible. In other words, given a set of beliefs and a new belief, we want to find a new set of beliefs which includes the new belief and differs as little as possible from the old set. Classical logic has the *ex absurdo sequitur quodlibet* principle, i.e., from a contradiction anything follows; so, finding a contradiction in a system should render it completely useless. We know that things go differently. Contradictions are tackled trying to repair systems locally, and keeping modifications as small as possible. Systems (not only logical ones, but also normative systems) are very conservative – or, at least, this is how we behave about them. Belief revision theories are the logic tools for keeping change small. Different specifications of the problem result in different formal treatments.

The first watershed is between theories dealing with sets closed w.r.t. logical consequence (*belief sets*) and theories dealing with (finite) non-closed sets (*belief bases*). Belief sets are simpler; the principle of irrelevance of syntax, for instance, is trivially satisfied: it does not matter how I describe a set, because what is important is the set (closed w.r.t. logical consequence) – the principle of extensionality. Belief bases, on the contrary, are much more realistic. Our knowledge is described by finite sets of sentences, and so are the norms regulating our daily behavior and all sorts of databases. Any change affects only finite sets of sentences. This makes more difficult to satisfy (and even to state clearly) the intuitively appealing principle of irrelevance of syntax.[1]

There are then two distinctions, originating from the same source: the first is made between purely logical approaches and approaches assuming the existence of extra-logical properties of the sentences; the second among approaches looking for a unique resulting revision and approaches admitting multiple results. The source of the problem is that there may be different candidates for a revision (of a set or base). While this is not a theoretical problem, it is a hard practical difficulty, when trying to put the theory of belief revision at work (e.g., in a deductive database system). We are confronted with two alternatives: a non-deterministic choice among the candidates, or building a unique result as a combination of the candidates; relying on extra-logical properties of the sentences may, sometimes, find a unique revision, but, more often, it simply elects a subset of candidates, and the problem of finding a unique result is still there.

A special case of extra-logical properties, pertaining to finite bases, is the distinction of sentences between those that can be affected by the revision and those that cannot be affected. The distinction is rather common: it suffices to consider the defeasible and non-defeasible rules of non-monotonic logics, or, from another point of view, the distinction between observed facts and theoretical constructs of (some approaches to) epistemology. As a mnemonic, we shall speak of rules (sentences that can be changed) and facts (sentences that cannot be changed).[2]

This paper deals with approaches to finite base revision, looking for a unique solution to the revision, and using, as sole extra-logical property, the distinction between rules and facts. When we use the word *revision* we intend this type of revision, unless otherwise stated. We describe three approaches to revision, and show that they are equivalent, in the sense that they result in the same bases. These approaches fail to satisfy the principle of extensionality: equivalent

[1] It may be argued that syntax is far from irrelevant, at least in some context. This is an example from a normative context.

Let D be a system containing the following norms: 1) "adults may vote and drive cars" 2) "people become adults at 18".

Of course, 3) "people adults 18 may vote", and 4) "people adults 18 may drive cars" are consequences of D . If we accept that systems bearing the same consequences are essentially the same (irrelevance of syntax), then D is equivalent to D′ explicitly made of: 1) "adults may vote and drive cars", 2) "people become adults at 18", and 3) "people adults 18 may vote". 4), of course, is a consequence of D′. Now, a bill changing the age in 2) to 21 passes. D is revised into E: 1) "adults may vote and drive cars", 2) "people become adults at 21"; and D′ is revised into E′: 1) "adults may vote and drive cars", 2) "people become adults at 21", and 3) "people adults 18 may vote". In both, driving requires to be 21, but only in E′ it is possible to vote at 18.

[2] Another view, originating from the similarities between belief revision and the evolution of normative systems, holds that sentences which cannot be modified are "intended consequences" of regulations (norms): if a regulation conflicts with an intended consequence, it is the norm that should be changed.

sets of rules may result, after revision, in non-equivalent set of rules. Then, we present a new approach to revision enjoying the following properties: the set of models of the resulting base is the same set of models of the base obtained with the other approaches, and the set of models of the resulting rule set is not larger than the set of models of the rule sets resulting from the other approaches. Moreover, the approach is fully algorithmic, and lends itself to a straightforward implementation using tableaux systems.

In section 1, following AGM, we quickly revise the common ground of revision of closed belief sets. In section 2 we describe three approaches to base revision, and we show that they essentially lead to the same result. In section 3 we present a new approach, in which sentences are modified instead of being deleted, and we discuss its relationship with AGM postulates. In section 4 we show, with the help of some examples, the differences between the new method and the approaches described in section 2. Theorems are stated without proof, and the mathematical apparatus is kept to a minimum.

1 Revision for Belief Sets

The baseline for the modern treatment of belief revision is usually taken to be the result of the joint efforts of Alchourrón, Makinson and Gärdenfors. The theory they eventually arrived at is know as the AGM approach to belief revision (or, simply, AGM). This approach is fully described in [7]; later work in the same line includes [13], [8].

The central notion in AGM is that of *belief set.* A belief set is a set of sentences (of some propositional language) such that it may be rationally held by an individual, that is ([7], ch.2.2), a consistent set closed under logical consequence (i.e., a *theory*).

Definition 1. A set K of sentences is a non-absurd belief set iff (i) $\perp$ is not a logic consequence of the sentences in K and (ii) if $\mathrm{K} \vdash b$ then $b \in \mathrm{K}$.

Let $\mathrm{Cn}(S)$ denote the set of the consequences of a set S, then $\mathrm{K} = \mathrm{Cn}(\mathrm{K})$ holds for belief sets. The set of all sentences in the language is also a belief set, namely the *absurd belief set* $\mathrm{K}_\perp$. Belief sets may be infinite, but, usually, they are not maximal, i.e., given a sentence of the language a, it is possible that neither $a \in \mathrm{K}$ nor $\neg a \in \mathrm{K}$. If $a \in \mathrm{K}$, then a is accepted in K, if $\neg a \in \mathrm{K}$, then a is rejected, otherwise a is undetermined. Three basic kind of changes of belief are identified: expansion, contraction, revision. The table below summarizes the meaning and notation for them.

previous state	expansion K^+a	contraction K^-a	revision K^*a
a accepted in K	a accepted in K	a indetermined in K	a accepted in K
a indetermined in K	a accepted in K	a indetermined in K	a accepted in K
a rejected in K	$\mathrm{K}_\perp$	a rejected in K	accepted in K

Instead of giving operational definitions for the three operations, AGM give postulates for them intended to constrain the class of all possible operational definitions ([7], ch.3). Here we recall only the postulates for revision. For the postulates for expansion and contraction the reader is referred to [7]. The postulates for expansion define a unique operation, $\mathrm{K}^{+}a = \mathrm{Cn}(\mathrm{K} \cup \{a\})$; on the other hand the postulates for contraction and revision do not define a unique operation. Among the postulates the most controversial is the so-called postulate of recovery: if $a \in K$, then $K = (K^{-}a)^{+}a$. Revision and contraction are usually linked by the so called Levi Identity: $\mathrm{K}^{*}a = (\mathrm{K}^{-}\neg a)^{+}a$.

Definition 2. The AGM revision of a belief set K by a proposition a is a set $\mathrm{K}^{*}a$ such that:

(K*1) K^{*} is a belief set;
(K*2) $a \in \mathrm{K}^{*}a$;
(K*3) $\mathrm{K}^{*}a \subseteq \mathrm{K}^{+}a$;
(K*4) if $\neg a \notin \mathrm{K}$, then $\mathrm{K}^{+}a \subseteq \mathrm{K}^{*}a$;
(K*5) $\mathrm{K}^{*}a = \mathrm{K}_{\perp}$ iff $\neg a$ is logically true;
(K*6) if $a \leftrightarrow b$ is logically valid, then $\mathrm{K}^{*}a = \mathrm{K}^{*}b$;
(K*7) $\mathrm{K}^{*}(a \wedge b) \subseteq (\mathrm{K}^{*}a)^{+}b$;
(K*8) if $\neg b \notin \mathrm{K}^{*}a$, then $(\mathrm{K}^{*}a)^{+}b \subseteq \mathrm{K}^{*}(a \wedge b)$.

Notice that from K*3 and K*4 we obtain that if $\neg a \notin \mathrm{K}$, then $\mathrm{K}^{*}a = \mathrm{K}^{+}a$; moreover if $a \in \mathrm{K}$, then $\mathrm{K}^{*}a = \mathrm{K}$. Two properties of revision that can be derived from K*1–K*8 are: $\mathrm{K}^{*}a = \mathrm{K}^{*}b$ iff $b \in \mathrm{K}^{*}a$ and $a \in \mathrm{K}^{*}b$ and

$$\mathrm{K}^{*}(a \vee b) = \begin{cases} \mathrm{K}^{*}a & \text{or} \\ \mathrm{K}^{*}b & \text{or} \\ \mathrm{K}^{*}a \cap \mathrm{K}^{*}b \end{cases} \tag{1}$$

Thanks to the Levi identity, the discussion may be limited to contraction, as we may derive revision accordingly.

The result of contraction is a subset of the original set, so we have to look for a belief set K' which is a subset of K such that $a \notin \mathrm{K}'$; the principle of *minimal change* suggests that we look for the unique largest of such subsets. It is easy to see that, in general, there is not such a unique set. In general, we may find a family of sets, which is called $\mathrm{K}\perp a$ (see the next section for a formal definition). The most natural candidate for contraction is *Full Meet Contraction.*

$$\mathrm{K}^{-}a = \begin{cases} \bigcap(\mathrm{K}\perp a) & \text{if } \mathrm{K}\perp a \text{ is non-empty} \\ \mathrm{K} & \text{otherwise} \end{cases} \qquad \text{(Full Meet Contraction)}$$

Unfortunately it suffers from a major drawback:

Theorem 3. *If $a \in \mathrm{K}$ and $\mathrm{K}^{-}a$ is defined as Full Meet Contraction, then $b \in \mathrm{K}^{-}a$ iff (a) $b \in \mathrm{K}$; (b) $\neg a \rightarrow b$ is logically true.*

In other words, in this case $\mathrm{K}^{-}a = \mathrm{K} \cap \mathrm{Cn}(\neg a)$. Contraction results in a very small set and minimal change is not respected. On the other hand, if we take only one of the sets in $\mathrm{K}\bot a$ (*maxichoice* contraction), we are stuck in a non-deterministic choice. To overcome this shortcoming AGM defines contraction as the intersection of some *preferred* elements of $\mathrm{K}\bot a$. To this end we stipulate an order relation $\preceq$ over 2^{K}, thus the intersection is defined only over the maximal elements of $\preceq$: as we intersect less sets, we may get a larger result.

2 Revision for Belief Bases

AGM applies to sets which are closed under the consequence operator, that is, sets K such that K = Cn(K). Although some of the original papers in that tradition deals with sets which are not closed under Cn (see, e.g., [12]), the theory was developed mainly in the direction of closed sets. A different approach maintains that the assumption of closed sets is too unrealistic to be really significant as a model of how belief works. A belief base B is a set of propositions. The set of consequences of B, under a suitable consequence operator Cn, is $\mathrm{Cn}(B) \supseteq \mathrm{B}$. The idea is that a realistic approach to belief representation must take into account that any agents has only a finite set of beliefs, possibly relying on them in order to derive an infinite set of consequences. Any normative system, for instance, is finite, and revising a normative system means revising a finite set of norms actually comprising it. This approach is discussed, among other, in [4], [15], [10], [11], [16]. First we describe bases in which all sentences may be changed, then bases partitioned in rules and facts.

2.1 Bases Unpartitioned

We start from the definition of three families of interesting subsets of B.

Definition 4. $\mathrm{B}\bot a$, $\mathrm{B}_{\Rightarrow}a$, $\mathrm{B} \ddagger a$ are defined as follows.

- $\mathrm{B}\bot a$ is the family of maximal subsets of B not implying a:

$$\mathrm{C} \in \mathrm{B}\bot a \iff \begin{cases} \mathrm{C} \subseteq \mathrm{B} \\ a \notin \mathrm{Cn}(\mathrm{C}) \\ \text{if } b \notin \mathrm{C} \text{ and } b \in \mathrm{B}, \text{ then } a \in \mathrm{Cn}(\mathrm{C} \cup \{b\}) \end{cases}$$

- $\mathrm{B}_{\Rightarrow}a$ is the family of minimal subsets of B implying a:

$$\mathrm{C} \in \mathrm{B}_{\Rightarrow}a \iff \begin{cases} \mathrm{C} \subseteq \mathrm{B} \\ a \in \mathrm{Cn}(\mathrm{C}) \\ \text{if } \mathrm{D} \subset \mathrm{C}, \text{ then } \mathrm{D} \notin \mathrm{B}_{\Rightarrow}a \end{cases}$$

- $\mathrm{B} \ddagger a$ is the family of minimal incisions of a from B such that:

$$\mathrm{I} \in \mathrm{B} \ddagger a \iff \begin{cases} \mathrm{I} \subseteq \mathrm{B} \\ \mathrm{C} \cap \mathrm{I} \neq \emptyset \text{ for each } \mathrm{C} \in \mathrm{B}_{\Rightarrow}a \\ \text{if } \mathrm{D} \subset \mathrm{I}, \text{ there is } \mathrm{C} \in \mathrm{B}_{\Rightarrow}a \text{ such that } \mathrm{C} \cap \mathrm{D} = \emptyset \end{cases}$$

It is worth noting that if a is logically true then $\mathrm{B}\perp a = \emptyset$, $\mathrm{B}_{\Rightarrow} a = \{\emptyset\}$, $\mathrm{B} \ddagger a = \emptyset$. Moreover, if B is finite any element C of one of the families above is finite, and there is a finite number of them. Different definitions of revision are obtained as functions of the above defined families. For simplicity, we define at first the contraction functions, then, according to Levi identity, the revision functions are obtained by adding the new data to the output of the contraction functions. The main idea is that of simple base contraction, using $\mathrm{B}\perp a$. Contraction is defined as the intersection of all subsets in $\mathrm{B}\perp a$. (We modify the original notion given in [3], [9], [14] in order to retain the finiteness of the base).

A different approach was introduced in [1] for belief sets, but may be easily adapted to belief bases. The safe contraction is the base obtained deleting all possible minimal subsets implying a. Deleting the elements of an incision on $\mathrm{B}_{\Rightarrow} a$ transforms B into a set B′ such that $a \notin \mathrm{B}'$. Deleting the elements of all possible incisions (all possible choices of elements from $\mathrm{B}_{\Rightarrow} a$) we get the excided contraction.

Definition 5. Simple base, safe, and excided contraction and revision are defined as follows:

type	*contraction*	*symbol*
simple base	$\bigcap_{\mathrm{C}\in \mathrm{B}\perp a} \mathrm{C}$	$\mathrm{B}^{\ominus_1} a$
safe	$\left\{b \,\middle|\, b \in \mathrm{B} \text{ and } b \notin \bigcup_{\mathrm{C}\in \mathrm{B}_{\Rightarrow} a} \mathrm{C}\right\}$	$\mathrm{B}^{\ominus_2} a$
excided	$\left\{b \,\middle|\, b \in \mathrm{B} \text{ and } b \notin \bigcup_{\mathrm{I}\in \mathrm{B}\ddagger a} \mathrm{I}\right\}$	$\mathrm{B}^{\ominus_3} a$

revision
$\mathrm{B}^{\oplus_i} a = \mathrm{B}^{\ominus_i} \neg a \cup \{a\}$ $i = 1, 2, 3$

The three contraction (and revision) functions share a common flavour. First of all, all of them use a set-theoretic operation on all the sets of a certain family. They might be called *full* operations, in the sense in which the adjective "full" pertains to full meet contraction according to AGM.

We study now the relationships between the sets defined by these operations.

Theorem 6. *If* $a \in \mathrm{Cn}(\mathrm{B})$ *and* B *is finite, then* $\mathrm{B}^{\ominus_1} a = \mathrm{B}^{\ominus_2} a = \mathrm{B}^{\ominus_3} a$.

The same is not true, however, when only one set is considered. We are guaranteed that

- $a \notin \mathrm{Cn}(\{b | b \in \mathrm{B} \text{ and } b \notin \mathrm{I}\})$, for $\mathrm{I} \in \mathrm{B} \ddagger a$, and that
- $a \notin \mathrm{Cn}(\{b | b \in \mathrm{C}\})$, for $\mathrm{C} \in \mathrm{B}\perp a$,

but we are not guaranteed that

- $a \notin \mathrm{Cn}(\{b | b \in \mathrm{B} \text{ and } b \notin \mathrm{C}\})$, for $\mathrm{C} \in \mathrm{B}_{\Rightarrow} a$.

Following what is defined for belief sets, we call this type of operations *maxichoice* operations. In a sense, the notions of $\mathrm{B}\perp a$ and $\mathrm{B} \ddagger a$ are more robust than that of $\mathrm{B}_{\Rightarrow} a$, because they allow not only full operations, but also maxichoice operations which satisfy the minimal requirement of success. Maxichoice operations, however, even if they satisfy the success postulate, do not satisfy the uniqueness requirement, that is, the requirement that the result of contraction or revision is well defined. This is usually considered unacceptable.

The problems with full approaches is essentially the same as those with full meet contraction: they, usually, contract too much: the resulting set is too small.

2.2 Bases Partitioned in Rules and Facts

Now we introduce the distinction among facts and rules. Facts are sentences which are not subject to change in the process of revision, while rules may be changed. If B is a base, the sets of facts and rules will be denoted B_φ and B_ρ. The operations studied in the previous section are extended to this case.

We first examine an extension to the $\ominus_1$ operation, called prioritized base revision, the idea of which is due to [15]. In a prioritized revision, we first take all facts not implying a, then we add as many rules as we can without implying a, in all different ways.

Definition 7. $\mathrm{B} \Downarrow a$ is the family of the sets $\mathrm{C} = (\mathrm{C}_\rho, \mathrm{C}_\varphi)$, such that:

(a) $\mathrm{C}_\rho \subseteq \mathrm{B}_\rho$ and $\mathrm{C}_\varphi \subseteq \mathrm{B}_\varphi$;
(b) $a \notin \mathrm{Cn}(\mathrm{C}_\varphi)$ and if $\mathrm{C}_\varphi \subset \mathrm{D} \subseteq \mathrm{B}_\varphi$, then $a \in \mathrm{Cn}(\mathrm{D})$;
(c) $a \notin \mathrm{Cn}(\mathrm{C}_\varphi \cup \mathrm{C}_\rho)$, and if $\mathrm{C}_\rho \subset \mathrm{E} \subseteq \mathrm{B}_\rho$, then $a \in \mathrm{Cn}(\mathrm{C}_\varphi \cup \mathrm{E})$.

Definition 8. Let B be a finite base, and let $\Phi(\mathrm{C}) = \bigwedge_{b_i \in \mathrm{C}_\rho \cup \mathrm{C}_\varphi} b_i$; then the prioritized base contraction of B by a, is the set $\mathrm{B}\hat{\ominus}_1 a = \{\bigvee_{\mathrm{C}\in \mathrm{B}\Downarrow a} \Phi(\mathrm{C})\}$.

If we take the assumption that the sentence we want to retract is not implied by facts alone, that is $a \notin \mathrm{Cn}(\mathrm{B}_\varphi)$, then the elements of $\mathrm{B} \Downarrow a$ have the form $(\mathrm{C}_\rho, \mathrm{B}_\varphi)$, $\Phi(C) = \bigwedge_{b_i \in \mathrm{C}_\rho} b_i \wedge \bigwedge_{\varphi_i \in \mathrm{B}_\varphi} \varphi_i = \Phi(\mathrm{C}_\rho) \wedge \bigwedge_{\varphi_i \in \mathrm{B}_\varphi} \varphi_i$ and the following holds:

$$\text{if } a \notin \mathrm{Cn}(\mathrm{B}_\varphi),\ \text{then } \mathrm{B}\hat{\ominus}_1 a = \left(\left\{ \bigvee_{\mathrm{C}\in\mathrm{B}\Downarrow a} \Phi(\mathrm{C}_\rho)\right\}, \mathrm{B}_\varphi\right) \tag{2}$$

Another approach is related to the $\ominus_2$ operation, whose idea is due to [1]. We assume that only rules have to be blamed for a, so only rules involved in deriving a are deleted from the base; no set in $\mathrm{B}_\Rightarrow a$ contains only facts.

Definition 9. For each $\mathrm{C} \in \mathrm{B}_\Rightarrow a$, let $\mathrm{C}_\rho = \mathrm{B}_\rho \cap \mathrm{C}$, such that $\mathrm{C} \not\subseteq \mathrm{B}_\varphi$. The safe contraction of a from B, $\mathrm{B}\hat{\ominus}_2 a = \left(\left\{b | b \in \mathrm{B}_\rho \text{ and } b \notin \bigcup_{\mathrm{C}\in \mathrm{B}_\Rightarrow a} \mathrm{C}_\rho\right\}, \mathrm{B}_\varphi\right)$.

A third approach may be derived extending the $\ominus_3$ operation. This uses prioritized incisions. We assume, as usual, that a cannot be derived by facts alone. As above, for $\mathrm{C} \in \mathrm{B}_\Rightarrow a$, $\mathrm{C}_\rho = \mathrm{B}_\rho \cap \mathrm{C}$ and $\mathrm{C}_\varphi = \mathrm{B}_\varphi \cap \mathrm{C}$.

Definition 10. A prioritized incision (p-incision) on $\mathrm{B}_\Rightarrow a$ is a set $\mathrm{I} \subseteq \mathrm{B}_\rho$ such that $\mathrm{C}_\rho \cap \mathrm{I} \neq \emptyset$ for each $\mathrm{C} \in \mathrm{B}_\Rightarrow a$. A p-incision is minimal if for all $\mathrm{D} \subset \mathrm{I}$, D is not a p-incision. The family of all possible p-incisions is denoted by $\mathrm{B} \dagger a$.

Definition 11. The excided p-contraction of B by a, $\mathrm{B}\hat{\ominus}_3 a$, is the set $\{b | b \in \mathrm{B} \text{ and } b \notin \bigcup_{\mathrm{I}\in\mathrm{B}\dagger a} \mathrm{I}\}$.

As above, we examine the relations among these three types of operations. A result like the one obtained in the unpartitioned case holds, that is:

Theorem 12. *If $a \in \mathrm{Cn}(\mathrm{B})$ and B is finite, then $\mathrm{B}\hat{\ominus}_1 a = \mathrm{B}\hat{\ominus}_2 a = \mathrm{B}\hat{\ominus}_3 a$.*

As before, it is easy to see that $\hat{\ominus}_2$ and $\hat{\ominus}_3$ are equivalent, as p-incisions contain all and only the rules contained in at least one minimal subset, only differently arranged: $\bigcup_{\mathrm{C}\in\mathrm{B}_\Rightarrow a} \mathrm{C}_\rho = \bigcup_{\mathrm{I}\in\mathrm{B}\dagger a} \mathrm{I}$.

3 Modifying Is Better Than Deleting

3.1 The Proposed Procedure

In this section, we show how to modify rules in order to obtain a revised base. The revised base is equivalent to that we would obtain from the previously described methods, but the rule part, taken alone, has a larger set of consequences than its counterparts. That is, if RB is the revised base obtained through our procedure, and RB1 is the revised base obtained through one of the previously described procedures, $\mathrm{Cn(RB)} = \mathrm{Cn(RB1)}$ but $\mathrm{Cn(RB}_\rho) \supseteq \mathrm{Cn(RB1}_\rho)$.

Admittedly, taking this as an advantage is a matter of taste. In order to defend our view, we would like to compare rules to *programs*, and facts to *data*. Even if two programs are equivalent with respect to a certain set of data, it is fully justified to think that one is superior to the other if the former has a broader field of applicability. In a sense, if information is defined as constraints on the set of possible worlds, we want to get a set of rules that is as informative as possible. Another view is that we want to get the most from our base in the unlucky case that we are forced to retract all the facts we believe in. In Section 3.2 the concept of retraction, as opposed to contraction, is discussed at some length.

Dealing with belief bases instead of belief sets is difficult because it is not clear how much we have to compromise with syntax. We try to take a third way between irrelevance of syntax (reducing belief bases to belief sets) and full dependence on how propositions are expressed. In order to reach this deal, we use from one side a semantic approach (relating bases to their consequences) and from the other side a sort of syntactical normal form (see below). We want to stress that we want to describe a viable way of revising effectively a finite set of beliefs. By the way, the approach is fully algorithmic and can be implemented in the framework of well known procedures for tableaux-based theorem provers. We hope to report about it shortly.

Bases, Evaluations, Models Each fact φ_i may be rewritten as a disjunction of possibly negated atoms: $\varphi_i = \bigvee_k \gamma_{ik}$; similarly each rule ρ_i may be rewritten as

$$\bigvee_j \neg\alpha_{ij} \vee \bigvee_k \beta_{ik} = \alpha_i \rightarrow \beta_i$$

where $\alpha_i = \bigwedge_j \alpha_{ij}$ and $\beta_i = \bigvee_k \beta_{ik}$. Facts and rules in this form are said to be in normal form (or, simply, normal).

Let X be a set of atoms, and $\Sigma(X)$ be the set of sentences built upon X. Let v be an evaluation function for X, $v : X \rightarrow \{true, false\}$. We extend v to $\Sigma(X) \rightarrow \{true, false\}$ in the standard way. Let V be the set of evaluation functions for $\Sigma(X)$.

Definition 13. $v \in V$ is a model of $a \in \Sigma(A)$ iff $v(a) = true$. v is a model of $A = \{a_i\}$ iff v is a model of each a_i. The set of models of a base B is denoted by $V(\mathrm{B})$.

Proposition 14. *If $V(a1)$ and $V(a2)$ are models of $a1$ and $a2$, respectively, then $V(a1 \wedge a2) = V(\{a1, a2\}) = V(a1) \cap V(a2)$; and $V(a1 \vee a2) = V(a1) \cup V(a2)$.*

$V(\mathrm{B}) = V(\mathrm{B}_\rho) \cap V(\mathrm{B}_\varphi)$*;* $\mathrm{B1} \subseteq \mathrm{B2} \Rightarrow V(\mathrm{B1}) \supseteq V(\mathrm{B2})$*;* $\mathrm{Cn(B1)} \subseteq \mathrm{Cn(B2)} \Leftrightarrow V(\mathrm{B1}) \supseteq V(\mathrm{B2})$

Adding formulae to a base means getting a smaller set of models. In particular, if a is a formula inconsistent with a base B, $V(B \cup \{a\}) = \emptyset$. The revision of $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi)$ by a should be some base B' such that $V(\mathrm{B}'_\rho) \cap V(\mathrm{B}'_\varphi) \neq \emptyset$. In the process of revision, only the rule part of the base is changed, that is, $\mathrm{B}'_\varphi = \mathrm{B}_\varphi \cup \{a\}$. So, finding a revision amounts to finding a new set of rules B'_ρ.

Models and Revisions It principle, nothing is said about the relationships between $V(\mathrm{B}'_\rho)$ and $V(\mathrm{B}_\rho)$. Some properties are however desirable. First of all, we want the new formula a to be independent from the new rules $V(\mathrm{B}'_\rho)$, that is, $V(\mathrm{B}'_\rho) \cap V(a) \neq V(\mathrm{B}'_\rho)$; then we want that the new rules do not discard any of the models for the original rules, that is, $V(\mathrm{B}'_\rho) \supseteq V(\mathrm{B}_\rho)$. This is in line with Levi's idea that a revision loses something and then adds something else; in this case, the set of models gets larger (from $V(\mathrm{B}_\rho)$ to $V(\mathrm{B}'_\rho)$) and then is (effectively) intersected with $V(a)$. This covers the main case. If a is consistent with B, that is $V(\mathrm{B}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \neq \emptyset$, we simply add a to the set of facts, leaving the rules unchanged.

So the complete definition is:

Definition 15. A revision of $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi)$ by a is a (finite) base $\mathrm{B}' = (\mathrm{B}'_\rho, \mathrm{B}_\varphi \cup \{a\})$ such that:

If $V(\mathrm{B}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \neq \emptyset$, then $\mathrm{B}'_\rho = \mathrm{B}_\rho$, else:
(a) $V(\mathrm{B}'_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \neq \emptyset$;
(b) $V(\mathrm{B}'_\rho) \supseteq V(\mathrm{B}_\rho)$;
(c) $V(\mathrm{B}'_\rho) \cap V(a) \neq V(\mathrm{B}'_\rho)$;

Now we have to define a minimal inconsistent sub-base; this may be easily done using the definition of $\mathrm{B}_{\Rightarrow} a$. A minimal inconsistent sub-base for $(\mathrm{B}_\rho, \mathrm{B}_\varphi \cup \{a\})$ is simply an element of $\mathrm{B}_{\Rightarrow} \neg a$. When recast in the present semantic framework, the definition becomes:

Definition 16. Let $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi)$ be an inconsistent base. A minimal inconsistent sub-base of B is a base $\mathrm{C} = (\mathrm{C}_\rho, \mathrm{C}_\varphi)$ such that

1. $V(\mathrm{C}_\rho) \cap V(\mathrm{C}_\varphi) = \emptyset$;
2. $\mathrm{C}_\rho \subseteq \mathrm{B}_r$ and $\mathrm{C}_\varphi \subseteq \mathrm{B}_\varphi$;
3. if $\mathrm{D}_\rho \subseteq \mathrm{C}_\rho$, $\mathrm{D}_\varphi \subseteq \mathrm{C}_\rho$, and either $\mathrm{D}_\rho \subset \mathrm{C}_\rho$ or $\mathrm{D}_\varphi \subset \mathrm{C}_\varphi$, then $V(\mathrm{D}_\rho) \cap V(\mathrm{D}_\varphi) \neq \emptyset$.

The set of rules not included in the rule part of any minimal inconsistent sub-base will be denoted $\hat{\mathrm{B}}_\rho$. When clear from the context, $\hat{\mathrm{B}}_\rho$ will denote the set of rules not included in the rule part of any minimal inconsistent sub-base for $(\mathrm{B}_\rho, \mathrm{B}_\varphi \cup \{a\})$. In need, we shall use the notation $\hat{\mathrm{B}}_{\rho,a}$.

Orderly Revision and Minimax Revision Saying that rules are not ordered by importance means, roughly, that all the rules (involved in inconsistency) should be modified, their sets of model getting larger. We say that a revision is orderly if it modifies all and only the rules involved in inconsistency. Rules in the new base correspond to rules in the old one, with the possible exception that some old rules are eliminated. There is no reason to modify rules that do not belong to any minimal inconsistent subbase: so a "good" revision $\mathrm{B}' = (\mathrm{B}'_\rho, \mathrm{B}_\varphi \cup \{a\})$ should have a set of rules B'_ρ such that $\hat{\mathrm{B}}_\rho \subseteq \mathrm{B}'_\rho$, that is, $V(\mathrm{B}'_\rho) \subseteq V(\hat{\mathrm{B}}_\rho)$. This means that $V(\mathrm{B}') = V(\mathrm{B}'_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \subseteq V(\hat{\mathrm{B}}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a)$ and sets an upper bound on the size of $V(\mathrm{B}')$. It is immediate to see that $\hat{\mathrm{B}}_\rho$ is the set of rules obtained through the operation of safe revision: $\mathrm{B}\hat{\ominus}_2 a$.

Definition 17. Let $\mathrm{B}\top a$ be the family of sets of rules B'_ρ such that:

1. $V(\mathrm{B}_\rho) \subset V(\mathrm{B}'_\rho) \subseteq V(\hat{\mathrm{B}}_\rho)$;
2. $V(\mathrm{B}'_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \neq \emptyset$.

Let $m(\mathrm{B}\top a)$ be the set of $\subseteq$-minimal elements of $\mathrm{B}\top a$: $\mathrm{B}'_\rho \in m(\mathrm{B}\top a)$ iff there is no $\mathrm{B}''_\rho \in m(\mathrm{B}\top a)$ such that $V(\mathrm{B}''_\rho) \subseteq V(\mathrm{B}'_\rho)$. A minimax revision of B by a is $\check{B} = (\check{B}_\rho, \mathrm{B}_\varphi \cup \{a\})$ such that $V(\check{B}_\rho) = \bigcup_{\mathrm{B}'_\rho \in m(\mathrm{B}\top a)} V(\mathrm{B}'_\rho)$

Theorem 18. $V(\check{\mathrm{B}}) = V(\hat{\mathrm{B}}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a)$

Theorem 19. *There is an orderly revision* B^*a *such that* $V(\mathrm{B}^*a) = V(\hat{\mathrm{B}}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a)$.

Corollary 20. *Let* $\{\mathrm{C}^s\}_{s=1,\ldots,t} = \{(\mathrm{C}^s_\rho, \mathrm{C}^s_\varphi \cup \{at\})\}_{s=1,\ldots,t}$ *be the set of minimal inconsistent subbases of* $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi \cup \{a\})$. *Let be* $\mathrm{B}^{*a}_\rho = \{\hat{\rho}_1, \ldots, \hat{\rho}_k, \check{\rho}_1, \ldots, \check{\rho}_n\}$ *where* $\{\hat{\rho}_1, \ldots, \hat{\rho}_k\} = \hat{\mathrm{B}}_\rho$ *and* $\{\check{\rho}_1, \ldots, \check{\rho}_n\}$ *are rules such that* $V(\check{\rho}_i) = V(\rho_i) \cup (V(\mathrm{B}_\varphi) \cap V(a))$, *where* $\rho_i \in \mathrm{C}^s_\rho$. *The revision* $\mathrm{B}^*a = (B^{*a}_\rho, \mathrm{B}_\varphi \cup \{a\})$ *is orderly and minimax.*

From Theorems 18 and 19 $V(\mathrm{B}^{*a}_\rho) = \bigcup_{\mathrm{B}'_\rho \in m(\mathrm{B}\top a)} V(\mathrm{B}'_\rho) \subseteq V(\hat{\mathrm{B}}_\rho)$.

Theorem 21. $\mathrm{B}^*a = (\{\bar{\rho}_i\}, \mathrm{B}_\varphi \cup \{a\})$ *is a minimax revision where:*

$$\bar{\rho}_i = \begin{cases} \alpha_i \to \beta_i & \text{if } \alpha_i \to \beta_i \in \hat{\mathrm{B}}_\rho \\ \{(\alpha_i \wedge \neg\varphi) \to \beta_i)\} & \text{where } \varphi \in \mathrm{B}_\varphi \cup a \text{ otherwise} \end{cases}$$

3.2 Minimax Revision and AGM Postulates

We are now going to examine relations between minimax base revision and the AGM theory of belief set revision. Let us say that a sentence a belongs to the belief set generated by B iff a holds in any model in $V(\mathrm{B})$, that is:

Definition 22. Given a base $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi)$, the belief set generated by B is denoted K(B). A sentence $a \in \mathrm{K(B)}$ iff $V(\mathrm{B}) \subseteq V(a)$.

It follows from the definition that $\mathrm{K}(\mathrm{B}_1) \subseteq \mathrm{K}(\mathrm{B}_2)$ iff $V(\mathrm{B}_2) \subseteq V(\mathrm{B}_1)$. From Theorem 19 we know that $V(\mathrm{B}^*a) = V(\hat{\mathrm{B}}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a)$, that is, $V(\mathrm{B}^*a) = V((\hat{\mathrm{B}}_\rho, \mathrm{B}_\varphi \cup \{a\}))$. It follows that $\mathrm{K}(\mathrm{B}_\rho^{*a}, \mathrm{B}_\varphi \cup \{a\}) = K(\hat{\mathrm{B}}_\rho, \mathrm{B}_\varphi \cup \{a\})$, $V(\mathrm{B}_\rho^{*a}) \subseteq V(\hat{\mathrm{B}}_\rho)$, and $\mathrm{K}(\hat{\mathrm{B}}_\rho) \subseteq \mathrm{K}(\mathrm{B}_\rho^{*a})$.

The importance of this fact becomes apparent if we introduce another operation, which we name *retraction* to distinguish it from contraction. Retraction applies only to (sentences implied only by) facts and modifies the fact part of a base.[3]

Definition 23. The retraction of φ_j from $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi)$ is the base $\mathrm{B}_{-\varphi_j} = (\mathrm{B}_\rho, \{\varphi_i\}_{i \neq j})$.

Proposition 24. *For any* $\varphi_j \in \mathrm{B}_\varphi \cup \{a\}$, $\mathrm{K}((\hat{\mathrm{B}}_\rho, \mathrm{B}_\varphi \cup \{a\})_{-\varphi_j}) \subseteq \mathrm{K}((\mathrm{B}_\rho^{*a}, \mathrm{B}_\varphi \cup \{a\})_{-\varphi_j})$.

To show the result it suffices to show that $V((\mathrm{B}_\rho^{*a}, \mathrm{B}_\varphi \cup \{a\})_{-\varphi_j}) \subseteq V((\hat{\mathrm{B}}_\rho, \mathrm{B}_\varphi \cup \{a\})_{-\varphi_j})$. This holds because $V(\mathrm{B}_\rho^{*a}) \subseteq V(\hat{\mathrm{B}}_\rho)$.

We may now define contraction from revision and retraction.

Definition 25. The contraction of $\mathrm{B} = (\mathrm{B}_\rho, \mathrm{B}_\varphi)$ by a is $\mathrm{B}^- a = (\mathrm{B}_\rho^{*\neg a}, \mathrm{B}_\varphi) = \mathrm{B}^- a = (\mathrm{B}_\rho^{*\neg a}, \mathrm{B}_\varphi \cup \{\neg a\})_{-\neg a}$.

Postulates for Revision In the following table we show the status of minimax revision with respect to axioms K*1–K*8. Axioms are interpreted with K meaning $\mathrm{K}(\mathrm{B}_\rho, \mathrm{B}_\varphi)$, K^*a meaning $\mathrm{K}(\mathrm{B}^*a)$ and K^+a meaning $\mathrm{K}(\mathrm{B}_\rho, \mathrm{B}_\varphi \cup \{a\})$.

axiom	*holds?*	*notes*
K*1	*yes*	
K*2	*yes*	
K*3	*yes*	if a is consistent with B then $\mathrm{K}^*a = \mathrm{K}^+a$ else $\mathrm{K}^+a = \emptyset$
K*4	*yes*	if a is consistent with B, then $\mathrm{K}^*a = \mathrm{K}^+a$
K*5	*no*	it holds: $\mathrm{K}^*a = K_\perp$ iff $\neg a$ is logically true or $\neg a \in \mathrm{Cn}(\mathrm{B}_\varphi)$
K*6	*yes*	if $a \leftrightarrow b$ is logically true, then $V(\mathrm{B}_\varphi \cup \{a\}) = V(\mathrm{B}_\varphi \cup \{b\})$
K*7	*yes*	$V(\mathrm{B}^*(a \wedge b)) = \left(V(\hat{\mathrm{B}}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \cap V(b)\right)$ $V((\mathrm{B}^*a)^+b) = \left(V(\hat{\mathrm{B}}_\rho) \cap V(\mathrm{B}_\varphi) \cap V(a) \cap V(b)\right)$
K*8	?	it holds: if $\mathrm{B}_{\rho,a\wedge b} = \mathrm{B}_{\rho,a}$, then $(\mathrm{K}^*a)^+b \subseteq \mathrm{K}^*(a \wedge b)$

In conclusion, our revision operation on bases defines a corresponding revision operation on belief sets which satisfies the AGM postulates, with the exception of K*5 and K*8 which are substituted by weaker ones.

[3] In order to keep the definition simple, we suppose that the fact part of the base is irreducible that is, no fact can be deduced from the rest of the base. For technical reasons, we say that in this case B_φ is irreducible w.r.t. B_ρ. In the following, bases will be assumed to be irreducible.

The Axiom of Recovery As for contraction, defined through retraction, we focus on the most controversial of the postulates, that is, recovery:

$$\text{if } a \in \mathrm{K},\ \text{then } \mathrm{K} = (\mathrm{K}^{-}a)^{+}a$$

Then $(\mathrm{B}^{-}a)^{+}a = ((\mathrm{B}_{\rho}^{*\neg a}, \mathrm{B}_{\varphi}))^{+}a$; moreover, a belongs to the consequences of B, that is, $V(\mathrm{B}) \subseteq V(a)$.

Proposition 26. *The contraction defined from minimax revision and retraction satisfies the postulate of recovery.*

3.3 Discussion

Theorem 21 gives us a simple recipe for transforming rules. Not only is the procedure fully algorithmic, but it can also be decomposed: the revision is built piecewise, starting from individual rules and individual facts.

The result of the revision is itself in normal form, so that iterated revision is well-defined, contrary to what results for standard AGM belief set revision. (For ways of modifying the AGM theory in order to allow iterated revision, see:[2],[6]).

Minimax revision offers a clear definition of so-called *multiple* revision, that is revision by means of a set of sentences. When contraction is considered as the main operation, using a set of sentences instead of one often results in unclear intuitions (see [5]).

Extensionality does not hold in general. Adding to a set a derivative rule does not change the set of consequences of the rule, but it may change the output of a revision operation. This suggests a stronger (not algorithmic) notion of equivalence between two bases:

Definition 27. (Strong equivalence). $\mathrm{B1} = (\mathrm{B1}_{\rho}, \mathrm{B1}_{\varphi})$ is strongly equivalent to $\mathrm{B2} = (\mathrm{B2}_{\rho}, \mathrm{B2}_{\varphi})$ iff

1. $V(\mathrm{B1}) = V(\mathrm{B2})$ and
2. $\forall a,\ V(\mathrm{B1}^{*}a) = V(\mathrm{B2}^{*}a)$.

Equivalent, but not strongly equivalent, bases are given in example 32.

4 Examples

Five examples are described. The first two are closely related each other, and show how the procedure described in the preceding sections fares when compared to standard procedures (e.g. safe revision or prioritized base revision). The third example illustrates the need for using rules in normal form. The fourth one shows that the procedure, notwithstanding its simplicity, may deal with contradiction resulting from chains of implications.

Example 28. Let $\mathrm{B} = (\{a \to c, b \to c\}, \{a \vee b\})$. Let the new fact be $\neg c$.

$$\begin{aligned}\mathrm{B}^* \neg c &= (\{(a \wedge \neg((a \vee b) \wedge \neg c)) \to c, (b \wedge \neg((a \vee b) \wedge \neg c)) \to c\}, \{a \vee b, \neg c\}) \\ &= (\{(a \wedge c) \to c, (b \wedge c) \to c\}, \{a \vee b, \neg c\}) = (\emptyset, \{a \vee b, \neg c\})\end{aligned}$$

The same result would obtain if standard procedures had been used.

Example 29. Let $\mathrm{B} = (\{a \to c, b \to c\}, \{a \wedge b\})$. Let the new fact be $\neg c$.

$$\begin{aligned}\mathrm{B}^* \neg c &= (\{(a \wedge \neg((a \wedge b) \wedge \neg c)) \to c, (b \wedge \neg((a \wedge b) \wedge \neg c)) \to c\}, \{a \wedge b, \neg c\}) \\ &= (\{(a \wedge \neg b) \to c, (b \wedge \neg a) \to c\}, \{a \wedge b, \neg c\})\end{aligned}$$

The only difference between this example and the preceding one is that the fact component is $\{a \wedge b\}$ instead of $\{a \vee b\}$.

Standard procedures do not differentiate between the two aspects: when rules are deleted, we always get the revision $(\emptyset, \mathrm{B}_\varphi \cup \{\neg c\})$. Our procedure, which modifies rule on the ground of the new set of facts, produces instead a different result.

It could be argued that the models of the two revisions are exactly the same, so that the difference is only at a surface level. The answer is that the two revisions are indeed different as far as the rule part is concerned, because $V(\{(a \wedge \neg b) \to c, (b \wedge \neg a) \to c\}) \subset V(\emptyset)$. Were we to retract all facts, we would end up with a different set of evaluations. Our procedure is sensitive to differences which do not influence the behavior of other procedures.

These two examples shed some light on a sort of monotonicity property of the procedure. The (fact component of the) base of the first example is, intuitively, weaker than the (fact component of the) base of the second one, and this is reflected by the corresponding sets of evaluations. The two revisions show the same trend, in the sense that the rule set resulting from the first revision (the empty set) is of course weaker than the rule set resulting from the second revision.

Example 30. Let $\mathrm{B} = (\{a \vee b \to c\}, \{a\})$. (the rule is not in normal form). Let the new fact be $\neg c$. If we did not care about the normal form, the revision would be:

$$(\{((a \vee b) \wedge \neg(a \wedge \neg c)) \to c\}, \{a, \neg c\}) = (\{(b \wedge \neg a) \to c\}, \{a, \neg c\})$$

A base equivalent to B and in normal form is the the base $\mathrm{B}' = (\{a \to c, b \to c\},$ $\{a\})$. As $b \to c$ is not included in any minimal inconsistent sub-base,

$$\mathrm{B}'^* \neg c = (\{(a \wedge \neg(a \wedge \neg c)) \to c, b \to c\}, \{a, \neg c\}) = (\{b \to c\}, \{a, \neg c\})$$

Here again, the difference between the two resulting bases is appreciated only if we imagine the retraction of facts; in this case, it results that $V(b \to c) \subset V((b \wedge \neg a) \to c)$. This is a general result: rules in normal form result in smaller sets of evaluations, that is, more specific rule sets.

Example 31. Let $\mathrm{B} = (\{a \to b, b \to c, b \to d\}, \{a, d\})$. Let the new fact be $\neg c$.
The only minimal inconsistent sub-base is $(\{a \to b, b \to c\}, \{a, d, \neg c\})$.

$$\mathrm{B}^{*}\neg c = (\{a \wedge \neg d \to b, a \wedge c \to b, b \wedge \neg a \to c, b \wedge \neg d \to c, b \to d\}, \{a, d, \neg c\})$$

This should be contrasted with the result of standard procedures, which delete rules involved in contradiction: $(\{b \to d\}, \{a, d, \neg c\})$. Our procedure produces a more specific rule set, or, from another point of view, a rule set which leaves open the possibility of more inferences. This example shows also that the simple procedure automatically deals with contradictions deriving from chains of conditionals.

Example 32. Let

$$\mathrm{B1} = (\{a \to b, b \to c\}, \{b\}) \qquad \mathrm{B2} = (\{a \to b, b \to c, a \to c\}, \{b\})$$

be two equivalent bases. It is easy to see that $V(\mathrm{B1}_{\rho}) = V(\mathrm{B2}_{\rho})$. Let the new fact be $\neg c$, therefore the revisions are

$$\mathrm{B1}^{*}\neg c = (\{a \to b\}, \{b, \neg c\}) \qquad \mathrm{B2}^{*}\neg c = (\{a \to b, a \to c\}, \{b, \neg c\})$$

and $V(\mathrm{B2}^{*}\neg c) \subset V(\mathrm{B1}^{*}\neg c)$. This example shows that in dealing with bases, *syntax matters.*

5 Directions for future work

In this paper we have described a simple procedure for modifying knowledge bases expressed by finite sets of formulae of a propositional language, where each formula is a rule (which can be changed) or a fact (which cannot), to accommodate new facts; the main difference between our approach and other approaches is that rules are modified, instead of being deleted.

This procedure may be extended in many directions. The first possibility is the extension to bases with more than two degrees of importance for sentences. This is the case with normative systems, where the hierarchy of importance stems directly from the hierarchy of normative sources. The second direction deals with the meaning of the rules. While the belief revision literature usually employs a propositional language, speaking about rules and facts suggests that rules might be rather seen as axiom schemes, describing universally quantified relationships between variables. How do belief revision procedures behave in this new setting? This is, we believe, a general questions, to which few answers are given in the literature. The third field for extending the framework is that of (propositional) modal logic. The motivation is again that of representing some features of the evolution of normative systems, which are usually represented by means of systems of modal (deontic) logic. It is indeed the mechanism of derogation (specifying exceptions) which originally suggested to modify rules instead of deleting them.

Acknowledgements

This research was partially supported by the Australia Research Council under Large Grant No. A49803544.

References

1. Alchourrón, C.E., Makinson, D., On the logic of theory change: safe contraction. *Studia Logica*, 44, pp.405-422 (1985)
2. Boutilier, C., Iterated Revision and Minimal Change of Conditional Beliefs. *Journal of Philosophical Logic*, 25, pp.263-305 (1996)
3. Fagin, R., Ullmann, J.D., Vardi, M.Y., On the semantics of Updates in Databases. In *2nd ACM SIGACT-SIGMOD Symposium on Principles of database Systems*, pp. 352-365, (1983)
4. Fuhrmann, A., Theory contraction through base contraction. *Journal of Philosophical Logic*, 20, pp.175-203, (1991)
5. Fuhrmann, A. and Hansson, S.O.: A Survey of Multiple Contractions, *The Journal of Logic, Language and Information 3*, pp. 39-75, (1994)
6. Gabbay, D.M., Rodrigues, O., A Methodology for Iterated Theory Change. In *Gabbay D.M., Ohlbach H.J. (eds.): Practical Reasoning,* LNAI 1085, Springer-Verlag Berlin, pp.193-207 (1996)
7. Gärdenfors, P., Knowledge in Flux. MIT Press, Cambridge (MA), (1988)
8. Gärdenfors, P., Makinson, D., Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65, pp. 197-245, (1994)
9. Ginsberg, M.L., Counterfactuals. *Artificial Intelligence* 30(1), pp. 35-79, (1986)
10. Hansson, S.O., In Defence of Base Contraction. Synthese, 91 (1992), pp.239-245
11. Hansson, S.O., A Dyadic Representation of Belief. In *Gärdenfors, P. (ed): Belief Revision*, Oxford University Press, 1993
12. Makinson, D., How to give up: a survey of some formal aspects of the logic of theory change. *Synthese*, 62, pp.347-363, (1985)
13. Makinson, D., Gärdenfors, P., Relations between the logic of Theory change and Nonmonotonic Logic. In *Fuhrmann, A., Morreau, M. (eds): The Logic of Theory Change*, Springer-Verlag, Berlin, (1991)
14. Nebel, B., A Knowledge Level Analysis of Belief Revision. In *Brachman, R.J., Levesque, H.J., Reiter, R. (eds.), Principles of Knowledge Representation and Reasoning: Proceedings of the 1st Int. Conference*, pp.301-311,(1989)
15. Nebel, B., Syntax-based Approaches to Belief Revision. In *Gärdenfors, P. (ed): Belief Revision*, Oxford University Press, 1993
16. Tennant, N., Changing the Theory of Theory Change: Towards a Computational Approach. *British Journal for Philosophy of Science*, 45, pp. 865-897, (1994)