

Revising Nonmonotonic Theories: The Case of Defeasible Logic

D. Billington, G. Antoniou, G. Governatori, and M. Maher

School of Computing and Information Technology
Griffith University, QLD 4111, Australia
{db,ga,guido,mjm}@cit.gu.edu.au

Abstract. The revision and transformation of knowledge is widely recognized as a key issue in knowledge representation and reasoning. Reasons for the importance of this topic are the fact that intelligent systems are gradually developed and refined, and that often the environment of an intelligent system is not static but changes over time. Traditionally belief revision has been concerned with revising first order theories. Nonmonotonic reasoning provides rigorous techniques for reasoning with incomplete information. Until recently the dynamics of nonmonotonic reasoning approaches has attracted little attention. This paper studies the dynamics of defeasible logic, a simple and efficient form of nonmonotonic reasoning based on defeasible rules and priorities. We define revision and contraction operators and propose postulates. Our postulates try to follow the ideas of AGM belief revision as far as possible, but some AGM postulates clearly contradict the nonmonotonic nature of defeasible logic, as we explain. Finally we verify that the operators satisfy the postulates.

1 Introduction

The revision and transformation of knowledge is widely recognized as a key problem in knowledge representation and reasoning. Reasons for the importance of this topic are the fact that intelligent systems are gradually developed and refined, and that often the environment of an intelligent system is not static but changes over time.

Belief revision [1, 7] studies reasoning with changing information. Traditionally belief revision techniques have been concerned with the revision of knowledge expressed in classical logic. The approach taken is to study *postulates for operators*, the most well-known operators being revision and contraction.

Until recently little attention was devoted to the revision of more complex kinds of knowledge. But in the past few years there has been an increasing amount of work on revising nonmonotonic knowledge, in particular default logic theories [17, 9, 18]. These works were motivated by the ability of default reasoning to maintain inconsistent knowledge, and the use of default reasoning in various application domains. For example, the use of default rules has been proposed for the maintenance of software [5, 14]. In requirements engineering the use of

default rules has been proposed [2, 19] to identify and trace inconsistencies among single requirements. One key issue in requirements engineering is the evolution of requirements, which technically translates to the evolution of default theories.

Default logics are known to be computationally complex [10, 12]. In this paper we will study a simple, efficient default reasoning approach, defeasible logic [16]. It is a sceptical reasoning approach based on the use of defeasible rules and priorities between them. Its usefulness has been demonstrated in several domains [3, 6].

In studying the revision of knowledge in defeasible logic, first we formulate postulates for revision and contraction operators in defeasible logic. We chose to be guided by the AGM postulates for classical belief revision [1]. Some of the AGM postulates can be readily adopted. Others need to be slightly modified, but we can demonstrate a close link to the *motivation* of the postulate as expressed, say, by Gardenfors [7]. But some AGM postulates contradict the nonmonotonic nature of defeasible logic. This contradiction is not surprising since AGM belief revision was designed for the revision of (monotonic) classical logical theories.

Once we establish the postulates we define concrete revision and contraction operators for defeasible logic, and show that they satisfy the proposed postulates.

2 Defeasible Logic

In this paper we use a simplified version of defeasible logic, in that strict rules are not considered; for the description of the full logic see [15]. We also consider only an essentially propositional version of the logic: the language does not contain function symbols and every expression with variables represents the finite set of its variable-free instances. A knowledge base consists of facts, rules (defeasible rules and defeaters), and a superiority relation among rules.

Facts denote simple pieces of information that are deemed to be true regardless of other knowledge items. Thus, facts are not revisable. A typical fact is that Tweety is a bird: $bird(tweety)$.

There are two kinds of rules, defeasible rules and defeaters. A rule r consists of its *antecedent* $A(r)$ (written on the left) which is a finite set of literals, an arrow, and its *consequent* $C(r)$ which is a literal. In examples we will often omit set notation for $A(r)$.

Defeasible rules are rules that can be defeated by contrary evidence. An example of such a rule is “Birds typically fly”; written formally:

$$bird(X) \Rightarrow flies(X).$$

The idea is that if we know that something is a bird, then we may conclude that it flies, unless there is other evidence suggesting that it may not fly. Defeasible rules with an empty antecedent are a little like facts, but they are defeatable and revisable.

Defeaters are rules that cannot be used to draw any conclusions. Their only use is to prevent some conclusions. In other words, they are used to defeat some

defeasible rules by producing evidence to the contrary. An example is “If an animal is heavy then it may not be able to fly”. Formally:

$$heavy(X) \rightsquigarrow \neg flies(X)$$

The main point is that the information that an animal is heavy is not sufficient evidence to conclude that it doesn't fly. It is only evidence that the animal *may* not be able to fly. In other words, we don't wish to conclude $\neg flies(X)$ if $heavy(X)$, we simply want to prevent a conclusion $flies(X)$.

The *superiority relation* among rules is used to define priorities among rules, that is, where one rule may override the conclusion of another rule. For example, given the defeasible rules

$$\begin{aligned} r &: republican \Rightarrow \neg pacifist \\ r' &: quaker \Rightarrow pacifist \end{aligned}$$

which contradict one another, no conclusive decision can be made about the pacifism of a person who is both a republican and a quaker. But if we introduce a superiority relation $>$ with $r > r'$, then we can indeed conclude $\neg pacifist$.

A *defeasible theory* T is a triple $(F, R, >)$ where F is a finite consistent set of literals (called *facts*), R a finite set of rules, and $>$ an acyclic superiority relation on R . $R[q]$ denotes the set of rules in R with head q , and $R_d[q]$ denotes the set of defeasible rules in R with head q .

A *conclusion* of T is a tagged literal and can have one of the following three forms: (i) $+\partial q$, which means that q is defeasibly provable in T ; (ii) $-\partial q$, which means that we have proved that q is not defeasibly provable in T ; and (iii) Σq , which means that there is a reasoning chain supporting q .

Provability is defined below. It is based on the concept of a proof in $T = (F, R, >)$. A *proof* or *derivation* is a finite sequence $P = (P(1), \dots, P(n))$ of tagged literals satisfying the following conditions ($P(1..i)$ denotes the initial part of the sequence P of length i , and $\sim p$ the complement of a literal p):

- $+\partial$: If $P(i+1) = +\partial q$ then either
- (1) $q \in F$ or
 - (2) (2.1) $\exists r \in R_d[q] \forall a \in A(r) : +\partial a \in P(1..i)$ and
 - (2.2) $\sim q \notin F$ and
 - (2.3) $\forall s \in R[\sim q]$ either
 - (2.3.1) $\exists a \in A(s) : -\partial a \in P(1..i)$ or
 - (2.3.2) $\exists t \in R_d[q]$
 - $\forall a \in A(t) : +\partial a \in P(1..i)$ and
 - $t > s$

- $-\partial$: If $P(i+1) = -\partial q$ then
- (1) $q \notin F$ and
 - (2) (2.1) $\forall r \in R_d[q] \exists a \in A(r) : -\partial a \in P(1..i)$ or
 - (2.2) $\sim q \in F$ or
 - (2.3) $\exists s \in R[\sim q]$ such that

- (2.3.1) $\forall a \in A(s) : +\partial a \in P(1..i)$ and
(2.3.2) $\forall t \in R_d[q]$ either
 $\exists a \in A(t) : -\partial a \in P(1..i)$ or
not $t > s$

- Σq : If $P(i+1) = \Sigma q$ then
(1) $q \in F$ or
(2) $\exists r \in R_d[q] \forall a \in A(r) : \Sigma a \in P(1..i)$

The elements of a proof are called *lines* of the proof. We say that a tagged literal L is *provable* in $T = (F, R, >)$, denoted $T \vdash L$, iff there is a proof P in T such that L is a line of P . We say that literal q is *provable in T* iff $T \vdash +\partial q$, and that q is *supported in T* iff $T \vdash \Sigma q$.

Even though the definition seems complicated, it follows ideas which are intuitively appealing. For example, the condition $+\partial$ states the following: One way of establishing that q is defeasibly provable is to show that it is a fact. The other way is to find a rule with conclusion q , all antecedents of which are defeasibly provable. In addition, it must be established that $\sim q$ is not a fact (to do otherwise would be counterintuitive – to derive q defeasibly, although there might be a definite reason against it), and for every rule s which might provide evidence for $\sim q$, either one of its antecedents is provably not derivable, or there is a rule with conclusion q which is stronger than s and can be applied (that is, all its antecedents are defeasibly provable). Essentially the defeasible rules with head q form a *team* which tries to counterattack any rule with head $\sim q$. If the rules for q win then q is derived defeasibly; otherwise q cannot be derived in this manner.

Support means simple forward chaining reasoning without considering counterarguments.

Finally let $T = (F, R, >)$ be a defeasible theory, M a set of literals and $X = M \cup \{\sim p \mid p \in M\}$. We define $T_M^\ominus = (F, R - R[X], >')$, where $R[X]$ denotes the set of rules in R with head a literal in X , and $>'$ is the reduct of $>$ on $R - R[X]$.

3 Postulates for Belief Change Operators

In this section we will formulate reasonable postulates for belief revision operators in a defeasible reasoning framework. We will be considering the classical AGM postulates and will be proposing necessary modifications.

3.1 Belief Bases and Belief Sets

First we need to specify the kinds of belief sets we will consider. In classical logic, a belief set is supposed to be a deductively closed set of formulas. In our framework it is natural to study sets of conclusions from a given defeasible theory T . We note that in the defeasible logic variant we are studying here the

conclusions are single literals. But we can easily reason with a somewhat more general set of conclusions: conjunctions of literals. Other types of conclusions can also be treated but we would need more complex technical means, so we defer their discussions to later papers. Technically we will consider a conjunction of literals $p_1 \wedge \dots \wedge p_n$ to be a conclusion of T iff $T \vdash +\partial p_i$ for all $i = 1, \dots, n$.

A set of literals BB is a *belief base* iff there is a defeasible theory T such that $BB = \{p \mid T \vdash +\partial p\}$. Equivalently, a belief base is any finite consistent set of literals. We say that BB is *generated by* T .

Now we define the belief set $B(T)$ generated by the defeasible theory T to be the conjunctive closure of the corresponding belief base. That is, a set of conjunctions of literals B is a *belief set* iff there is a defeasible theory T such that $B = \{p_1 \wedge \dots \wedge p_n \mid T \vdash +\partial p_1, \dots, T \vdash +\partial p_n\}$.

In the following we assume that $T = (F, R, >)$ is a defeasible theory, BB the belief base generated by T , and $B(T)$ the belief set generated by T . ΣT denotes the set of beliefs supported by T . Finally p_i and q_j denote literals, and c, c_k denote conjunctions of literals.

3.2 The Change Operators

As in classical belief revision, we will distinguish between three kinds of belief changes:

- (i) *expansion* $+$, which seeks to add a new formula φ to the belief set B if the negation of φ is not included in B . This is the original motivation of expansion as explained by Gardenfors in [7].
- (ii) *revision* $*$, which adds a formula φ to the belief set B even in cases where the negation of φ is in B . To achieve this outcome the operation may need to delete formulas in B .
- (iii) *contraction* $-$, which seeks to retract a sentence φ without adding new conclusions.

In practice, revision and contraction have attracted the greatest attention. Therefore we will not be addressing postulates for expansion.

In accordance with the defeasible logic we are studying in this paper, the sentences φ will be either literals or conjunctions of literals. In future work we will study change with full propositional formulas in variants of defeasible logic allowing for disjunction.

3.3 Postulates for Revision

The first postulate in classical belief revision states that B_p^* is closed under logical consequences. According to the discussion in section 3.1, its counterpart is the following which, under our definitions, holds trivially:

- (*1) $B(T_c^*)$ is a belief set.

The second AGM postulate in classical belief revision guarantees that the sentence φ is added to the belief set. Basically we adopt the same idea, with one difference: we forbid the addition of a contradiction to the belief base. There are two kinds of contradictions. The obvious one is to have a complementary pair of literals in the conjunction. For the other kind of contradiction we note that we consider facts in a defeasible theory to be undisputed information. Thus if we attempt to add a literal p_i but $\sim p_i$ is a fact, then the addition of p_i is rejected, as is the addition of any conjunction of literals that contains p_i .

(*2) If $\sim p_i \notin F$ and $p_i \neq \sim p_j$ (for all i, j) then $p_1 \wedge \dots \wedge p_n \in B(T_{p_1 \wedge \dots \wedge p_n}^*)$.

According to Gardenfors [7], the aim of the third and fourth AGM postulates is to identify revision with expansion in case the negation of the sentence φ to be added is not in B . The next two postulates achieve the same for defeasible logic. (The use of two postulates by [1] is for technical reasons [7], which do not seem to hold for defeasible logic. We use two postulates only to be consistent with the AGM numbering.)

(*3) If $\sim p_1, \dots, \sim p_n \notin B(T)$ then $B(T_{p_1 \wedge \dots \wedge p_n}^*) \subseteq B(T_{p_1 \wedge \dots \wedge p_n}^+)$.

(*4) If $\sim p_1, \dots, \sim p_n \notin B(T)$ then $B(T_{p_1 \wedge \dots \wedge p_n}^+) \subseteq B(T_{p_1 \wedge \dots \wedge p_n}^*)$.

The fifth AGM postulate states that the result of a revision by φ is the absurd belief set iff $\neg\varphi$ is logically valid. Since, by definition, there is no absurd belief set in defeasible logic (due to its sceptical nature, p and $\sim p$ cannot be proven together), establishing an exact counterpart to the fifth AGM postulate is not straightforward. We note, though, that in AGM postulates 2 and 5 are related: it is postulate 2 which admits a contradiction and creates the possibility of an absurd belief set, as specified in postulate 5. In our case we mentioned already that the addition of a contradictory sentence should be rejected. The types of contradiction were discussed before the definition of (*2). As in AGM, our fifth postulate defines the behaviour of revision in case we try to add a contradictory formula.

(*5) If $\sim p_i \in F$ or $\sim p_i = p_j$ (for some i, j) then $p_1 \wedge \dots \wedge p_n \notin B(T_{p_1 \wedge \dots \wedge p_n}^*)$.

In fact this postulate can be strengthened by providing more information on what happens if addition of a contradictory formula is rejected. We formulate two variations. (*5a) states that in case $p_1 \wedge \dots \wedge p_n$ is self-contradictory, or at least one $\sim p_i \in F$, then revision by $p_1 \wedge \dots \wedge p_n$ should not cause any change. Note that in case $\sim p_i \in F$ for some i , the entire conjunction contradicts the facts.

Condition (*5b) is weaker in that *all* p_1, \dots, p_n are required to contradict F for the revision of $p_1 \wedge \dots \wedge p_n$ to have no effect.

(*5a) If $\sim p_i \in F$ or $\sim p_i = p_j$ (for some i, j) then $B(T_{p_1 \wedge \dots \wedge p_n}^*) = B(T)$.

(*5b) If $\sim p_i \in F$ (for all i) or $\sim p_i = p_j$ (for some i, j) then $B(T_{p_1 \wedge \dots \wedge p_n}^*) = B(T)$.

The sixth AGM postulate expresses syntax-independence, and has a natural counterpart in defeasible logic.

(*6) If the set of literals in the conjunctions c_1 and c_2 is the same, then $B(T_{c_1}^*) = B(T_{c_2}^*)$.

This concludes the adaptation of the basic AGM postulates. Now we consider the composite postulates, the seventh and eighth, which regard the conjunction of sentences. According to [7] by the principle of minimal change revision with both φ and ψ ought to be the same as the expansion of $B(T_\varphi^*)$ by ψ , provided that ψ does not contradict the beliefs in $B(T_\varphi^*)$. Again for technical reasons, two AGM postulates were formulated. Below we state two postulates for defeasible logic revision to maintain the correspondence to the AGM postulates for classical belief revision.

- (*7) If $\sim q_1, \dots, \sim q_m \notin B(T_{p_1 \wedge \dots \wedge p_n}^*)$ then
 $B(T_{p_1 \wedge \dots \wedge p_n \wedge q_1 \wedge \dots \wedge q_m}^*) \subseteq B((T_{p_1 \wedge \dots \wedge p_n}^*)_{q_1 \wedge \dots \wedge q_m}^+)$.
- (*8) If $\sim q_1, \dots, \sim q_m \notin B(T_{p_1 \wedge \dots \wedge p_n}^*)$ then
 $B((T_{p_1 \wedge \dots \wedge p_n}^*)_{q_1 \wedge \dots \wedge q_m}^+) \subseteq B(T_{p_1 \wedge \dots \wedge p_n \wedge q_1 \wedge \dots \wedge q_m}^*)$.

3.4 Postulates for Contraction

The first postulate we formulate is a natural adaptation of the first AGM postulate.

(-1) $B(T_c^-)$ is a belief set.

The second AGM postulate states that we contract a formula *only* by deleting some formulas, but not by adding new ones. This postulate cannot be adopted in our framework because it contradicts the sceptical nonmonotonic nature of defeasible logic. To see this, suppose that we know a , and we have rules $\Rightarrow p$ and $a \Rightarrow \neg p$. Then a is sceptically provable and p is not. But if we decide to contract a then p becomes sceptically provable. We note that this behaviour is not confined to the specifics of defeasible logic but holds in any sceptical nonmonotonic formalism (e.g. in the sceptical interpretation of default logic).

But this example also suggests what proportion of the original idea of the second AGM postulate we can maintain: even though p was not in the original belief set, its appearance in the result of the contraction is not due to the addition of new rules. Stated differently, even though p was not *provable* in the original nonmonotonic knowledge base, it was nevertheless *supported*.

(-2) $B(T_c^-) \subseteq \Sigma T$.

The principle of minimal change requires that if the sentence to be retracted is not included in the belief set, no change is necessary.

(-3) If $c \notin B(T)$ then $B(T_c^-) = B(T)$.

The fourth AGM postulate states that the sentence to be retracted is not included in the outcome of the contraction operation, unless it is a tautology. In our framework this situation cannot arise since a conjunction of literals is never a tautology. But if we follow our earlier idea that the negations of facts are contradictions (thus facts may be viewed as being always true), it is natural to state the following:

(-4) If $p_1 \wedge \dots \wedge p_n \in B(T_{p_1 \wedge \dots \wedge p_n}^-)$ then $p_1, \dots, p_n \in F$.

The fifth AGM postulate expresses the possibility of recovery: if we retract a sentence and then add it again, then we do not lose any of the original beliefs. This idea is naturally expressed in defeasible logic below.

The condition $c \in B(T)$ is necessary: in case $c \notin B(T)$, the addition of c may remove a literal $\sim p_i$ which contradicts a conjunct p_i in c . Due to the nonmonotonic nature of defeasible logic, this may cause the removal of further literals from $B(T)$. Thus without the condition $c \in B(T)$ the postulate cannot be reasonably expected to hold.

(-5) If $c \in B(T)$ then $B(T) \subseteq B((T_c^-)^+)$.

Syntax independence is again straightforward.

(-6) If the set of literals in the conjunctions c_1 and c_2 is the same, then $B(T_{c_1}^-) = B(T_{c_2}^-)$.

The seventh and eighth AGM postulates are simply adopted.

(-7) $B(T_{c_1}^-) \cap B(T_{c_2}^-) \subseteq B(T_{c_1 \wedge c_2}^-)$

(-8) If $c_1 \notin B(T_{c_1 \wedge c_2}^-)$ then $B(T_{c_1 \wedge c_2}^-) \subseteq B(T_{c_1}^-)$

4 Revising Defeasible Theories

Having formulated postulates for change operators motivated by the original AGM postulates, the question remains what *concrete operators* might satisfy these sets of postulates, indeed whether such operators exist. In the following we define such concrete operators and prove that indeed they satisfy our postulates.

We will use a common notation for all operators. Let $c = p_1 \wedge \dots \wedge p_n$ be the formula to be added/deleted.

4.1 Expansion

Gärdenfors [7] states that expansion is meant to add a formula φ to $B(T)$ *only if* $\neg\varphi \notin B(T)$. In this sense, the case where $\neg\varphi \in B(T)$ is irrelevant. However AGM decided to also add φ in this case. We will keep T unchanged, following [7] rather than [1].

$$T_c^+ = \begin{cases} T & \text{if } \sim p_i \in B(T) \text{ for some } i \in \{1, \dots, n\} \\ T & \text{if } \sim p_i = p_j \text{ for some } i, j \in \{1, \dots, n\} \\ (F, R', >') & \text{otherwise} \end{cases}$$

where

$$\begin{aligned} R' &= R \cup \{\Rightarrow p_1, \dots, \Rightarrow p_n\} \\ >' &= (> \cup \{\Rightarrow p_i > r \mid i \in \{1, \dots, n\}, r \in R[\sim p_i]\}) - \\ &\quad \{r > \Rightarrow p_i \mid i \in \{1, \dots, n\}, r \in R[\sim p_i]\}. \end{aligned}$$

Thus we add rules that prove each of the literals p_i , and ensure that these are strictly stronger than any possibly contradicting rules.

4.2 Revision

AGM revision works in the same way as AGM expansion when the formula to be added does not cause an inconsistency, but revision adds a formula even if its negation is in the belief set. In our framework the definition of revision looks as follows:

$$T_c^* = \begin{cases} T & \text{if } p_1 \wedge \dots \wedge p_n \in B(T) \\ (F, R', >') & \text{otherwise} \end{cases}$$

where

$$\begin{aligned} R' &= R \cup \{\Rightarrow p_1, \dots, \Rightarrow p_n\} \\ >' &= (> \cup \{\Rightarrow p_i > r \mid i \in \{1, \dots, n\}, r \in R[\sim p_i]\}) - \\ &\quad \{r > \Rightarrow p_i \mid i \in \{1, \dots, n\}, r \in R[\sim p_i]\}. \end{aligned}$$

The difference to our previous definition of $+$ is that now we make the modifications to the rules and the superiority relation even if the negation of one of the p_i was in B . Note that due to the sceptical nature of defeasible logic, we still do not get a contradiction (in the sense of being able to prove both p_i and $\sim p_i$. Theorem 1 below (equivalently, the postulates (*1) – (*8)) specifies when we will be able to prove p_i and when $\sim p_i$.

4.3 Contraction

We define a concrete contraction operator as follows.

$$T_c^- = \begin{cases} T & \text{if } p_1 \wedge \dots \wedge p_n \notin B(T) \\ (F, R', >') & \text{otherwise} \end{cases}$$

where

$$\begin{aligned} R' &= R \cup \{p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_n \rightsquigarrow \sim p_i \mid i \in \{1, \dots, n\}\} \\ >' &= > - \{s > r \mid r \in R' - R\}. \end{aligned}$$

Intuitively we wish to prevent the proof of $p_1 \wedge \dots \wedge p_n$, that is, the proof of all the p_i . We achieve this by ensuring that at least one of the p_i will not be proven. The new rules in R' ensure that if all but one p_i have been proven, a defeater with head $\sim p_j$ will fire. Having made the defeaters not weaker than any other rules, the defeater cannot be “counterattacked” by another rule, and p_j will not be proven, as an inspection of the condition $+\partial$ in section 2 shows.

4.4 Results

Here we formulate some results on the revision concepts introduced in this paper. The main theme is to investigate which of the stated postulates are satisfied by the concrete change operations.

Neither of the two composite contraction AGM postulates (7 and 8) translate naturally into our framework, because they contradict the sceptical nonmonotonic nature of our underlying logical machinery. For example, the seventh AGM postulate would suggest $B(T_p^-) \cap B(T_q^-) \subseteq B(T_{p \wedge q}^-)$. Now consider a nonmonotonic knowledge base with defeasible rules $\Rightarrow p, \Rightarrow q, p \Rightarrow a$ and $q \Rightarrow a$. Then a is provable in both $B(T_p^-)$ and $B(T_q^-)$, but not in $B(T_{p \wedge q}^-)$. As for the eighth AGM postulate, it is motivated by the idea that if we need to remove more (say p and $p \wedge q$), then we will get a smaller belief set than just removing p . But already in our discussion of (-2) we saw that a similar property cannot be reasonably expected in our framework.

In our postulates we maintain the following idea of the last two AGM postulates: we try to express an upper and a lower bound for a composite contraction operation. For the lower bound we expect that $B(T_{p \wedge q}^-)$ contains at least the conclusions that can be proven if we remove from the underlying defeasible theory all rules with head $p, \sim p, q$, or $\sim q$. For the upper bound we use the same idea as for (-2) , that is, by contracting more information we may not add beliefs that are not *supported* in the result of contracting less.

$$(-7a) \quad B(T_{\{p_1, \dots, p_n, q_1, \dots, q_m\}}^\ominus) \subseteq B(T_{p_1 \wedge \dots \wedge p_n \wedge q_1 \wedge \dots \wedge q_m}^-).$$

$$(-8a) \quad \text{If } p_1 \wedge \dots \wedge p_n \notin B(T_{p_1 \wedge \dots \wedge p_n \wedge q_1 \wedge \dots \wedge q_m}^-) \text{ then } B(T_{p_1 \wedge \dots \wedge p_n \wedge q_1 \wedge \dots \wedge q_m}^-) \subseteq \Sigma T_{p_1 \wedge \dots \wedge p_n}^-.$$

The following result connects the concrete change operators defined above to the postulates established in the previous section.

Theorem 1. *Let $+, *$ and $-$ be the concrete change operators defined in the previous subsections. Then*

- Postulates $(*1) - (*8)$ are satisfied. Moreover, $(*5b)$ is satisfied but not $(*5a)$.
- Postulates $(-1) - (-6)$ are satisfied. Postulates (-7) and (-8) are not satisfied, but $(-7a)$ and $(-8a)$ are satisfied.

The proof can be found in the full version of this paper.

The Levi and Harper Identities [13, 11] have been proposed as ways to define revision and contraction in terms of the other two operators. These identities are consistent with the AGM postulates for classical belief revision. The Levi Identity holds for our concrete operators.

Theorem 2. *Let $+$, $*$ and $-$ be the change operators defined above. Then the following is true:*

$$B(T_p^{**}) = B((T_{\sim p}^-)_p^+). \quad \text{Levi Identity}$$

On the other hand, the counterpart of the Harper Identity does not hold, nor can it be reasonably expected:

$$B(T_p^-) = B(T) \cap B(T_{\sim p}^{**}). \quad \text{Harper Identity}$$

As we have already seen in the discussion of (-2) , $B(T_p^-)$ cannot be expected to be a subset of $B(T)$, but clearly the right hand side of the Harper Identity is a subset of $B(T)$.

Note that we are only able to formulate these identities for literals. For more complex formulas we need to incorporate disjunction in the belief sets and the deductive machinery of defeasible logic (to be able to express the negation of a conjunction).

5 Conclusion

In this paper we studied the revision of knowledge in defeasible logic. In particular, we formulated desirable postulates for revision and contraction operators. In some cases we followed the intuition and formalization of the original AGM postulates. In other cases we explained why certain AGM postulates are inappropriate because they contradict the nonmonotonic nature of defeasible logic, and proposed reasonable alternatives.

To our knowledge this is the first complete set of postulates for the revision of explicit nonmonotonic knowledge (of course we are aware of the close relationship between belief revision and nonmonotonic reasoning [8], as well as work showing that belief revision can be achieved using default reasoning [4]).

Then we defined concrete revision and contraction operators, and showed that they fulfill the postulates, as well as adaptations of the Levi Identity.

Our future work in the area will include the study of revision in defeasible logics with disjunction, and the examination of the relevance of revision to the evolution of knowledge in particular application domains.

References

1. C.E. Alchourron, P. Gardenfors and D. Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic* 50 (1985), 510–530.

2. G. Antoniou. The role of nonmonotonic representations in requirements engineering. *International Journal of Software Engineering and Knowledge Engineering* 8,3 (1998): 385–399.
3. G. Antoniou, D. Billington and M. Maher. On the analysis of regulations using defeasible rules. In *Proc. 32nd Hawaii International Conference on Systems Science*, IEEE Press 1999.
4. G. Brewka. Preferred Subtheories: An Extended Logical Framework for Default Reasoning. In *Proc. 11th IJCAI*, 1989, 1043–1048.
5. D.E. Cooke and Luqi. Logic Programming and Software Maintenance. *Annals of Mathematics and Artificial Intelligence* 21 (1997): 221–229.
6. M.A. Covington. Defeasible Logic on an Embedded Microcontroller. In *Proc. 10th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, June 1997.
7. P. Gardenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. The MIT Press 1988.
8. P. Gardenfors and D. Makinson. Nonmonotonic Inference Based on Expectations. *Artificial Intelligence* 65 (1994): 197–245.
9. A. Ghose and R. Goebel. Belief states as default theories: Studies in non-prioritized belief change. In *Proc. ECAI'98*.
10. G. Gottlob. Complexity Results for Nonmonotonic Logics. *Journal of Logic and Computation* 2,3 (1992): 397–425.
11. W.L. Harper. Rational Conceptual Change. In *PSA 1976*, Philosophy of Science Association, Vol. 2, 462–494.
12. H.A. Kautz and B. Selman. Hard problems for simple default logics. *Artificial Intelligence* 49 (1991):243–279.
13. I. Levi. Subjunctives, dispositions and chances. *Synthese* 34, 423–455, 1977.
14. Luqi and D.E. Cooke. How to Combine Nonmonotonic Logic and Rapid Prototyping to Help Maintain Software. *International Journal on Software Engineering and Knowledge Engineering* 5,1 (1995): 89–118.
15. M.J. Maher, G. Antoniou and D. Billington. A Study of Provability in Defeasible Logic. In *Proc. 11th Australian Joint Conference on Artificial Intelligence*, LNAI 1502, Springer 1998, 215–226.
16. D. Nute. Defeasible Reasoning. In *Proc. 20th Hawaii International Conference on Systems Science*, IEEE Press 1987, 470–477.
17. M.A. Williams and N. Foo. Nonmonotonic Dynamics of Default Logics. In *Proc. 9th European Conference on Artificial Intelligence*, 1990, 702–707.
18. M.A. Williams and G. Antoniou. A Strategy for Revising Default Theory Extensions. in *Proc. 6th International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann 1998.
19. D. Zowghi and R. Offen. A Logical Framework for Modeling and Reasoning about the Evolution of Requirements. In *Proc. Third International Symposium on Requirements Engineering*, IEEE Press 1997.