Declarative Representation of Revision Strategies

Gerhard Brewka [1]

Abstract. We introduce a nonmonotonic framework for belief revision in which reasoning about the reliability of different pieces of information based on meta-knowledge about the information is possible. The approach is based on a Poole-style system for default reasoning in which entrenchment information is represented in the logical language. A notion of inference based on the least fixed point of a monotone operator is used to make sure that all theories possess a consistent set of conclusions.

1 INTRODUCTION

Formal models of belief revision differ in what they consider as representations of the epistemic states of an agent. In the AGM approach epistemic states are identified with logically closed theories. Other approaches like those discussed in [11] consider finite sets of formulas, sometimes called belief bases, as epistemic states. A smaller fraction of work in belief revision has studied an obvious alternative: the revision of epistemic states expressed as nonmonotonic theories [4, 9, 16, 2, 6].

This is somewhat surprising since close relationships between properties of nonmonotonic inference relations and postulates for belief revision have been established [10]. Indeed, one of the reasons why nonmonotonic logics were invented is their ability to handle conflicts and inconsistencies, one of the major issues in belief revision. If this is the case, shouldn't it be possible to use the power of nonmonotonic inference to simplify revision? In fact, in this paper we will put the burden of revising beliefs entirely on the inference relation. We will revise nonmonotonic theories simply by adding new information, and leave everything else to nonmonotonic inference.

An early approach in this spirit was [4] where an extension of Poole-systems [12], the so-called preferred subtheory approach, was used. New information, possibly equipped with information about the reliability level of this information, was simply added to the available information. The nonmonotonic inference relation determined the acceptable beliefs. This approach is not fully satisfactory for several reasons. Existing theories of belief revision, including the one presented in that paper, have difficulties to model the way real agents revise their beliefs. In particular, they do not represent information which is commonly used by agents for this purpose. New information always comes together with certain meta-information: Where does the information come from? Was it an observation? Did you read it in the newspaper? Did someone tell you, and if so, who? Did the person who gave you the information have a motive to lie? and so on. In most cases we reason with and about this meta-information when revising our beliefs. We believe that realistic models of revision should provide the necessary means to represent this kind of information.

The meta-information is used to determine the entrenchment of pieces of information. The less entrenched the information is, the more willing we are to give it up. Again, entrechment relations are not just there, they result from reasoning processes. To model this kind of reasoning, entrenchment should be expressible in the logical language. Once we have the possibility to express entrenchment (or plausibility, or preference) in the language, it will also become possible to represent revision strategies declaratively. This in turn makes it possible to revise the revision strategies themselves.

The following example illustrates what we have in mind. Assume Peter tells you that your girl-friend Anne went out for dinner with another man yesterday. Peter even knows his name: it was John, a highly attractive person known for having numerous affairs. You are concerned and talk to Anne about this. She tells you she was at home yesterday evening waiting for you to call. Peter insists that he saw Anne with that man. You are not sure what to believe. Luckily, you find out that Anne has a twin sister Mary. Mary indeed went out with her new boy-friend John. This explains why Peter got mixed up. You now believe Anne and happily continue your relationship.

What this example illustrates is the way we reason about the reliability of information. There is no given fixed entrenchment ordering to start with. In the example there is also, at least in the beginning, no reason to trust Peter more than the girl-friend, or vice versa. And obviously, it is not the new information that is accepted in each situation. It is the additional context information which is relevant here: it gives us an explanation for Peter's mistake and decreases the reliability of Peter's observation enough to break the tie.

To be able to formalize examples of this kind we propose in this paper an approach to belief revision where 1) nonmonotonic belief bases represent epistemic states and nonmonotonic inference fully solves the revision problem, 2) it is possible to express and reason about meta-information, including the reliability of formulas, and 3) revision strategies can be represented declaratively.

The outline of the paper is as follows. In Sect. 2 we introduce the nonmonotonic formalism used to represent epistemic states. In Sect. 3 we show how to use this formalism for representing revision strategies. Sect. 4 shows that almost all of the AGM postulates are not valid in our approach. In Sect. 5 we briefly discuss contraction. Sect. 6 discusses related work and concludes.

2 REPRESENTING RELIABILITY RELATIONS

In this section we introduce the formalism used in this paper. One of the distinguishing features of our approach is the ablility to reason about the reliability of the available information in the logical language. In the AGM approach [7, 8] entrenchment relations are used to represent how strongly an agent sticks to his beliefs: the more entrenched a formula, the less willing to give it up the agent is. Entrenchment relations have several properties which are based on the

[1] Institut für Informatik, Universität Leipzig, Augustusplatz 10-11, 04109 Leipzig, Germany, brewka@informatik.uni-leipzig.de

logical strength of the formulas. For instance, logically weaker formulas are not less entrenched than logically stronger ones. The intuition is that if a weaker formula has to be given up, the stronger formula has to be given up anyway.

In our approach we do not require such properties. We may even have identical formulas p and p' with different reliability. This may happen when, for instance, p and p' come from different sources s and s' with different reliability. Note that although the less reliable information does not add to the accepted beliefs as long as the more reliable equivalent information is in force, the situation may change when new information about the reliability of s is obtained. Should s turn out to be highly unreliable later (of course, beliefs about the reliability of sources may be revised as any other beliefs) then it becomes important to have p' with, say, somewhat lower reliability available.

All we require, therefore, is the existence of a strict partial order $\prec$ between formulas which tells us how to resolve potential conflicts. To avoid misunderstandings we will not call $\prec$ an entrenchment relation. Instead, we speak of reliability, or simply priority among formulas. Since we want to represent $\prec$ in the logical language we need to be able to refer to formulas. We will use named formulas, that is pairs consisting of a formula and a name for the formula. Technically, names are just ground terms that can be used everywhere in the language.

Our formalism extends the well-known Poole systems [12]. Such systems consist of a consistent set of (first order) formulas F, the facts, and a possibly inconsistent set of formulas D, the defaults. A set of formulas E is an extension of a Poole-system (F, D) iff $E = Th(F \cup D')$ where D' is a maximal F-consistent subset of D.

Our formalism differs from this approach in the following respects:

1. We do not want to consider some information as absolutely certain and unrevisable. We therefore do not use F. Instead, we have a single set T containing all the information.
2. We represent preference and other meta-information in the language. We therefore introduce names for formulas and a special symbol $<$. $d < d'$ intuitively says that in case of a conflict d' should be given up rather than d since the latter is more reliable. We require that $<$ represents a strict partial order.[2]
3. We introduce a new notion of extension which takes the preference information into account adequately.

We now present the formal definitions. For simplicity, we only consider finite default theories in this paper. A generalization to the infinite case would have to be based on well-orderings rather than total orders.

Definition 1 *A named formula is a structure of the form $d{:}p$, where p is a first order formula and d a ground term representing the name of the formula.*

We use the functions $name$ and $form$ to extract the name respectively formula of a named formula, that is $name(d{:}p) = d$ and $form(d{:}p) = p$. We will also apply both functions to sets of named formulas with the obvious meaning.

Definition 2 *A preference default theory T is a finite set of named formulas such that*

- *$form(T)$ is a set of first order formulas whose logical language contains a reserved symbol $<$ representing a strict partial order, and*
- *$d_1{:}p \in T$, $d_2{:}q \in T$ and $p \neq q$ implies $d_1 \neq d_2$.*

The second clause of the definition above guarantees that syntactically different formulas have different names.

Definition 3 *Let T be a preference default theory, $\prec$ a total order on T. The extension of T generated by $\prec$, denoted $E_T^\prec$, is the set $E_T^\prec = Th(\bigcup_{i=0}^{|T|} E_i)$ where*

- *$E_0 = \emptyset$, and for $0 < i \leq |T|$*
- *$E_i = E_{i-1} \cup \{form(d_i{:}p)\}$ if this set is consistent, E_{i-1} otherwise.*
  Here $d_i{:}p$ is the i-th element of T according to the total order $\prec$.

The set $\bigcup_{i=0}^{|T|} E_i$ is called the extension base of $E_T^\prec$.

We say E is an extension of T if there is some total order $\prec$ such that $E = E_T^\prec$. Obviously, all maximal consistent subsets of $form(T)$ are extension bases. We now consider the general case of partial orders.

Definition 4 *Let T be a preference default theory, $\prec$ a strict partial order on T. The set of extensions of T generated by $\prec$ is*

$$Ext_T^\prec = \{E_T^{\prec'} \mid \prec' \text{ is a total order extending } \prec\}.$$

We next define two notions of compatibility:

Definition 5 *Let T be a preference default theory, $\prec$ a strict partial ordering of T, S a set of formulas. We say $\prec$ is compatible with S iff*

$$S \cup \{d < d' \mid d{:}p \prec d'{:}q\} \cup \{\neg(d < d') \mid d{:}p \not\prec d'{:}q\}$$

is consistent.
An extension E of T is compatible with S iff there is a strict partial ordering $\prec$ of T compatible with S such that $E \in Ext_T^\prec$.

The set of extensions of T compatible with S is denoted Ext_T^S.

Definition 6 *Let T be a preference default theory. A set of formulas E is called a preferred extension of T iff $E \in Ext_T^E$.*

Intuitively, E is a preferred extension if it is the deductive closure of a maximal consistent subset of T which can be generated through a preference ordering compatible with the preference information in E itself.

Here is a simple example illustrating preference default theories:

$d_1(x) : bird(x) \rightarrow flies(x) \mid x$ is a ground object term
$d_2 : \forall x. penguin(x) \rightarrow \neg flies(x)$
$d_3 : bird(tweety) \wedge penguin(tweety)$
$d_4 : \forall x. d_3 < d_1(x)$
$d_5 : \forall x. d_2 < d_1(x)$

As is common in Poole systems, rules with exceptions, that is, formulas whose instances can be defeated without defeating the formula as a whole (here d_1), are represented as schemata used as abbreviations for all of their ground instances. As above we will make the intended instances explicit in all examples. To make sure that the different ground instances can be distinguished by name we have to parameterize the names also. We assume that terms used as names can be distinguished from other terms which we call object terms.[3] In our case, $d_1(tweety)$ is a proper rule name, $d_1(d_1)$ is not. Since

[2] We assume that the properties of $<$, like those of equality, are part of the underlying logic and need not be represented through explicit axioms in our default theories.

[3] A more elaborate formalization would be based on sorted logic with sorts for names and other types of objects from the beginning. We do not pursue this here since we want to keep things as simple as possible.

we only consider finite theories we must also assume that the set of object terms is finite.

In our example we obtain 3 extensions E_1, E_2 and E_3. In E_1 the instance of $d_1(x)$ with $x = tweety$ is rejected, in E_2 d_2 is rejected, and E_3 rejects d_3. All extensions contain d_4 and d_5. It is not difficult to see that only E_1 can be constructed using an ordering of T which is compatible with this information. E_1 is thus the single preferred extension of this preference default theory.

Preference default theories under extension semantics are very flexible and expressive. Unfortunately, they can express unsatisfiable preference information: there are theories which do not possess any preferred extensions. The simplest example is as follows:

$d_1 : d_2 < d_1$
$d_2 : d_1 < d_2$

Accepting the first of the two contradictory formulas requires to give preference to the second, and vice versa. No preferred extension exists for this theory.

This means that preference default theories together with the standard notion of nonmonotonic inference where a formula is considered derivable whenever it is contained in all (preferred) extensions do not seem fully adequate for representing epistemic states of rational agents.

We will therefore introduce another, somewhat less standard notion of nonmonotonic consequence based on the least fixed point of a monotone operator. We consider an extension as acceptable if it is contained in the biggest set of extensions $\mathcal{E}$ satisfying the condition $E \in \mathcal{E}$ implies E is compatible with $\bigcap \mathcal{E}$. To obtain this set we first compute all extensions taking no preferences into account. We then eliminate extensions not compatible with the intersection of all extensions obtained so far and continue like this until no further extension can be eliminated, that is, until a fixed point is reached.

To formalize this idea we define an operator whose least fixed point is the intersection of the extensions which are acceptable in the sense just described. The least fixed point is computed by iterating the operator on the empty set. In each step, the argument of the operator corresponds to the preference information that needs to be taken into account, and the result of the operator corresponds to the intersection of those extensions which are still under consideration.

Definition 7 *Let T be a preference default theory, S a set of formulas. We define an operator C_T as follows:*

$$C_T(S) = \bigcap Ext_T^S$$

Proposition 8 *The operator C_T is monotone.*

Proof: $S \subseteq S'$ implies that an ordering $\prec$ is compatible with S whenever it is compatible with S'. We thus have $Ext_T^{S'} \subseteq Ext_T^S$ and therefore $\bigcap Ext_T^S \subseteq \bigcap Ext_T^{S'}$. □

Monotone operators, according to the well-known Knaster-Tarski theorem [15], possess a least fixed point. We, therefore, can define the accepted conclusions of a preference default theory as follows:

Definition 9 *Let T be a preference default theory. A formula p is an accepted conclusion of T iff $p \in lfp(C_T)$, where $lfp(C_T)$ is the least fixed point of the operator C_T.*

We call extensions which are compatible with $lfp(C_T)$ accepted extensions.

Several illustrative examples will be given in the next section. Here we just show how the theory without preferred extension is handled in this approach. We have $T = \{d_1{:}(d_2 < d_1), d_2{:}(d_1 < d_2)\}$. We first compute $C_T(\emptyset)$. Since no preference information is available in the empty set we obtain $Th(\{d_2 < d_1\}) \cap Th(\{d_1 < d_2\})$ which is equivalent to $Th(\{d_2 < d_1 \vee d_1 < d_2\})$. This set is already the least fixed point.

Proposition 10 *Let T be a preference default theory, p an accepted conclusion of T. Then p is contained in all preferred extensions of T.*

Proof: If T has no preferred extension the proposition is trivially true. So assume T possesses preferred extension(s). A simple induction shows that each preferred extension is among the extensions compatible with the formulas computed in each step of the iteration of C_T. Therefore each preferred extension is also an accepted extension. □

Proposition 11 *Let T be a preference default theory. The set of accepted conclusions of T is consistent.*

Proof: We can show by induction that the set of formulas obtained after an arbitrary number of applications of C_T is consistent. If S is consistent, then Ext_T^S is nonempty since an S-compatible partial ordering $\prec$ exists and each partial ordering generates at least one extension. Moreover, since extensions are consistent the intersection of a nonempty set of extensions is also consistent. □

3 REVISING EPISTEMIC STATES

From now on we identify an agent's epistemic state with a preference default theory as introduced in the last section. It is natural, then, to identify the set of beliefs accepted by the agent with the accepted conclusions of this theory. We therefore define belief sets as follows:

Definition 12 *Let T be an epistemic state. $Bel(T)$, the belief set induced by T, is the set of accepted conclusions of T.*

It is a basic assumption of our approach that belief sets cannot be revised directly. Revision of belief sets is always indirect, through the revision of the epistemic state inducing the belief set. Note that since two different epistemic states may induce the same belief set, the revision function which takes an epistemic state and a formula and produces a new epistemic state does not induce a corresponding function on belief sets.

Given an epistemic state T, revising it with new information simply means generating a new name for it and adding the corresponding named formula.

Definition 13 *Let T be an epistemic state, p a formula. The revision of T with p, denoted $T{*}p$, is the epistemic state $(T \cup \{n{:}p\})$ where n is a new name not appearing in T.*

Notation: in the rest of the paper we assume that names are of the form d_j where j is a numbering of the formulas. If T has j elements and a new formula is added, then its new name is d_{j+1}.

We now show how the revision strategies of an agent can be represented in our approach. We first discuss an example where the strategy is based on the type of the available information. We distinguish between strict rules, observations and defaults. Strict rules have highest priority because they represent well-established or terminological information. Observations can be wrong, but they are considered more reliable than default information. Consider the following epistemic state T:

$d_1 : penguin(tweety)$
$d_2 : \forall x.penguin(x) \rightarrow bird(x)$
$d_3 : \forall x.penguin(x) \rightarrow \neg flies(x)$
$d_4(x) : bird(x) \rightarrow flies(x) \mid x$ is a ground object term
$d_5 : observation(d_1)$
$d_6 : strictrule(d_2)$
$d_7 : strictrule(d_3)$
$d_8 : \forall x.default(d_4(x))$
$d_9 : \forall n, n'.strictrule(n) \wedge observation(n') \rightarrow n < n'$
$d_{10} : \forall n, n'.observation(n) \wedge default(n') \rightarrow n < n'$

The set of accepted conclusions of T contains all formulas in T except the instance of $d_4(x)$ with $x = tweety$. $Bel(T)$ thus does not contain $flies(tweety)$.

The next example formalizes the revision strategy of an agent who prefers newer information over older information and information from a more reliable source over information from a less reliable source. In case of a conflict between the two criteria the latter one wins. Assume the following specific scenario: At time 10 Peter informs you that p holds. At time 11 John tells you this is not true. Although you normally prefer later information, you also have reason to prefer what Peter told you since you believe Peter is more reliable than John. Since you consider reliability of your sources even more important than the temporal order you believe p.

Here is the formal representation of this scenario. We use $X < d$ where X is a finite set of names as an abbreviation for $\bigwedge_{x \in X} x < d$:

$d_1 : p$
$d_2 : \neg p$
$d_3 : time(d_1) = 10$
$d_4 : time(d_2) = 11$
$d_5 : source(d_1) = Peter$
$d_6 : source(d_2) = John$
$d_7 : more\text{-}rel(Peter, John)$
$d_8(n, n') : more\text{-}rel(source(n), source(n')) \rightarrow n < n'$
$d_9(n, n') : time(n) < time(n') \rightarrow n' < n$
$d_{10} : \forall n, n'.\{d_3, \ldots, d_7\} < d_8(n, n') < d_9(n, n')$

The schemata $d_8(n, n')$ and $d_9(n, n')$ represent ground instances with $n, n' \in \{d_1, \ldots, d_7\}$. Note that we have to make sure that the rules representing our revision strategy cannot be used - via contraposition - to defeat our meta-knowledge about d_1 and d_2. This is what the first inequality in d_{10} achieves.

We next present the example from the introduction. This time we use categories low, $medium$ and $high$[4] to express reliability: the reliability of a formula with name n is $rel(n)$. We have the following information:

$d_1 : date(A, J)$
$d_2 : \neg date(A, J)$
$d_3 : rel(d_1) = medium$
$d_4 : rel(d_2) = medium$
$d_5 : date(M, J)$
$d_6 : twins(M, A)$
$d_7 : date(M, J) \wedge twins(M, A) \rightarrow rel(d_1) = low$
$d_8 : \forall n, n'.rel(n) = high \wedge rel(n') = medium \rightarrow n < n'$
$d_9 : \forall n, n'.rel(n) = medium \wedge rel(n') = low \rightarrow n < n'$
$d_{10} : rel(d_5) = rel(d_6) = rel(d_7) = rel(d_8) = rel(d_9) = high$
$d_{11} : rel(d_3) = rel(d_4) = medium$

[4] We assume uniqueness of names for the categories. Otherwise the set $\{d_3, d_5, d_6, d_7\}$ would be consistent and could be used to defeat d_9 which, obviously, is unintended.

Although the agent initially considers d_1 and d_2 as equally reliable, the information that Anne has a twin sister Mary who is dating John decreases the reliability of d_1 to low. d_8 and d_9 say how the reliability categories are to be translated to preferences. d_{10} and d_{11} make sure that meta-information is preferred, and that d_7 can defeat d_3.

4 POSTULATES

We now discuss the postulates for revision which are at the heart of the AGM approach [7]. Since our approach uses epistemic states rather than deductively closed sets of formulas (belief sets) as substrate of revision, some of the postulates need reformulation. In particular, AGM use the expansion operator $+$ in some postulates. Expansion of a belief set K with a formula p means adding p to the belief set and closing under deduction, that is $K + p = Th(K \cup \{p\})$. Since epistemic states always induce consistent belief sets the distinction between revising and expanding an epistemic state does not make much sense in our context. We therefore translate expansion to expansion of the induced belief set. The following reformulations of the postulates are obtained from the original AGM postulates by replacing $K{*}A$ with $Bel(T{*}A)$ and $K + A$ with $Bel(T) + A$.

(T*1) $Bel(T{*}A)$ is belief set.

Obviously satisfied.

(T*2) $A \in Bel(T{*}A)$

Not satisfied. New information is not necessarily accepted in our approach. We see this as an advantage since otherwise belief sets would always depend on the order in which information was obtained.

(T*3) $Bel(T{*}A) \subseteq Bel(T) + A$

Not satisfied. Assume we have $T = \{d_1{:}p, d_2{:}\neg p\}$, that is $Bel(T)$ is the set of tautologies. Let $A = d_1 < d_2$. Now $Bel(T{*}A)$ contains p which is not contained in $Bel(T) + A$.

(T*4) if $\neg A \notin Bel(T)$ then $Bel(T) + A \subseteq Bel(T{*}A)$

Not satisfied. It may be the case that $\neg A$, although not in the belief set, is contained in one of the accepted extensions. Adding A to the epistemic state does not necessarily lead to a situation where this extension disappears.

(T*5) $Bel(T{*}A) \vdash \perp$ iff $\vdash \neg A$

Not satisfied. Revising an epistemic state with logically inconsistent information has no effect whatsoever. The information is simply disregarded. Inconsistent belief sets are impossible in our approach, so the right to left implication does not hold.

(T*6) If $A \leftrightarrow B$ then $Bel(T{*}A) = Bel(T{*}B)$

Satisfied under the condition that A and B are given the same name, or the names of A and B do not yet appear in S. But note that logically equivalent information may have different impact on the belief sets when different meta-information is available. For instance, $d_1{:}p$ and $d_2{:}p$ may have different effects if different meta-information about the sources of d_1 and d_2, respectively, is available.

(T*7) $Bel(T{*}(A \wedge B)) \subseteq Bel(T{*}A) + B$

Not satisfied. Here is a counterexample. Assume we have $T = \{d_1{:}p, d_2{:}\neg p, d_3{:}\neg p\}$. Now let $A = d_1 < d_2$ and $B = d_1 < d_3$.

Clearly, revising the epistemic state with $A \wedge B$ leads to a single accepted extension containing p since the two conflicting formulas are less preferred. p is thus in the belief set induced by the revised state. On the other hand, revising the epistemic state with A leads to two extensions, one containing p, the other $\neg p$. p is thus not in the belief set induced by the new state. This does not change when we expand the belief set with $d_1 < d_3$.

(T*8) If $\neg B \notin Bel(T*A)$
then $Bel(T*A) + B \subseteq Bel(T*(A \wedge B))$

Not satisfied. This is immediate from the fact that $Bel(T*(A \wedge B))$ does not necessarily contain B, that is from the failure of (T*2).

This analysis shows that the intuitions captured by the AGM postulates are indeed very different from those underlying our approach.

5 CONTRACTION

Contraction means making a formula underivable without assuming its negation. In the context of AGM-style approaches the contraction operator $-$ can be defined through revision on the basis of the so-called Harper identity: $K - A = (K*\neg A) \cap K$. The intuition here is that revision with $\neg A$ removes the formulas used to derive A, and the intersection with K guarantees that no new information is derived from $\neg A$. This intuition can, to a certain extent, be captured using Poole's constraints [12]. Constraints are formulas used in the construction of maximal consistent subsets of the premises, but not used for derivations. To model contraction of epistemic states we must distinguish between these two types of formulas, premises and constraints. Extension bases consist of both types and also the compatibility of preference orderings is checked against premises and constraints. Extensions, however, are generated only from the premises. Constraints, as regular formulas, have names and may come with meta-information, e.g., information about their reliability.

We cannot go into technical detail. Instead, we illustrate contraction using an example. We indicate constraints by choosing names of the form c_j for them. Assume the epistemic state is as follows:

$d_1 : peng(tw)$
$d_2(x) : peng(x) \rightarrow \neg flies(x)$

The agent receives the information "do not believe $\neg flies(tw)$", that is, we add the constraint $c_1 : flies(tw)$. Let $inst(d_2)$ denote the set of all ground instances of d_2. We obtain three extension bases

$E_1 = \{peng(tw)\} \cup inst(d_2)$
$E_2 = \{peng(tw), flies(tw)\} \cup inst(d_2) \setminus \{peng(tw) \rightarrow \neg flies(tw)\}$
$E_3 = \{flies(tw)\} \cup inst(d_2)$

Although E_2 and E_3 contain $flies(tw)$ this formula is not in the extensions generated from these extension bases, and for this reason not in the belief set, since it is a constraint.

Note that constraints do not necessarily prohibit formulas from being in the belief set since they may have low reliability. For example, if we revise the epistemic state obtained above with $d_1 < c_1$ and $\forall x. d_2(x) < c_1$ then the belief set contains $\neg flies(tweety)$.

Although we used the Harper identity above to motivate the use of constraints for contraction, its natural reformulation

$$Bel(T - A) = Bel(T*\neg A) \cap Bel(T)$$

is not valid in our approach. Assume $T = \{d_1{:}p, d_2{:}\neg p\}$. Obviously, $Bel(T)$ is the set of tautologies. Now let $A = \neg(d_1 < d_2)$. We contract by adding the constraint $c_1{:}(d_1 < d_2)$. Now the single accepted extension of the new epistemic state and thus its belief set is $Th(p)$, a strict superset of $Bel(T)$.

6 RELATED WORK AND DISCUSSION

We proposed a framework for belief revision where preference default theories together with a corresponding nonmonotonic inference relation are used to represent epistemic states and belief sets, respectively. Our underlying formalism draws upon ideas developed in [3] and [5], the notion of accepted conclusions introduced to guarantee consistency of belief sets and its application to belief revision is new. The framework is expressive enough to represent and reason about reliability and other properties of information.

In [4] nonmonotonic belief bases in the preferred subtheories framework were used to model revision. This approach, however, did not represent reliability information explicitly. Williams and Antoniou [16] investigated revision of Reiter default theories. In a similar spirit, Antoniou et al [2] discuss revision of theories expressed in Nute's defeasible logic. Also these approaches do not reason about the reliability of information. This is also true for existing work in revising logic programs, see [1] for an example. These approaches are thus very different from ours.

Acknowledgements

I thank R. Booth, S. Lange, H. Sturm and F. Wolter for helpful comments.

References

[1] Alferes, J.J., Pereira, L.M., Reasoning with Logic Programming, LNAI 1111, Springer, 1996

[2] Antoniou, G, Billington, D., Governatori, G. and M. Maher, Revising Nonmonotonic Belief Sets: The Case of Defeasible Logic. In Proc. 23rd German Conference on Artificial Intelligence, LNAI 1701, Springer 1999, 101-112

[3] Brewka, G., Preferred Subtheories - An Extended Logical Framework for Default Reasoning, Proc. IJCAI-89, 1989

[4] Brewka, G., Belief Revision in a Framework for Default Reasoning, in: A. Fuhrmann and M. Morreau (eds.), The Logic of Theory Change, Springer LNAI 465, Berlin, 1991

[5] Brewka, G., Eiter, T., Prioritizing Default Logic, to appear in Festschrift 60th Anniversary of W. Bibel, Kluwer, 2000

[6] Chopra, S., Parikh, R., An Inconsistency Tolerant Model for Belief Representation and Belief Revision, Proc. IJCAI-99, Morgan Kaufmann, 1999

[7] Gärdenfors, P., Knowledge in Flux, MIT Press, Cambridge, MA, 1988

[8] Gärdenfors, P., Makinson, D., Revisions of Knowledge Systems Using Epistemic Entrenchment. In: Vardi, M. (ed): Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge, Morgan Kaufmann, Los Altos, 1988

[9] Ghose, A., Goebel R., Belief States as Default Theories: Studies in Non-Prioritized Belief Change, Proc. ECAI-98, 1998

[10] Makinson, D., Gärdenfors, P., Relations Between the Logic of Theory Change and Nonmonotonic Logic, in: A. Fuhrmann and M. Morreau (eds.), The Logic of Theory Change, Springer LNAI 465, Berlin, 1991

[11] Nebel, B., Syntax Based Approaches to Belief Revision, in: P. Gärdenfors (ed.), Belief Revision, Cambridge, 1992

[12] Poole, D., A Logical Framework for Default Reasoning, Artificial Intelligence 36, 1988

[13] Reiter, R., A Logic for Default Reasoning, Artificial Intelligence 13, 1980

[14] Rescher, N., Hypothetical Reasoning, North-Holland Publ., Amsterdam. 1964

[15] Tarski, A., A lattice-theoretical fixpoint theorem and its applications, Pacific Journal of Mathematics, 5:285-309, 1955

[16] Williams, M.-A., Antoniou, G., A Strategy for Revising Default Theory Extensions, Proc. KR-98, 1998