Published in J. A. Allen, R. Fikes, and E. Sandewall (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, p. 417–428, Cambridge, MA, April 1991. Morgan Kaufmann.

Belief Revision and Default Reasoning: Syntax-Based Approaches

Bernhard Nebel
German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3
D-6600 Saarbrücken 11, Germany
e-mail: nebel@dfki.uni-sb.de

Abstract

Belief revision leads to *temporal nonmonotonicity*, i.e., the set of beliefs does not grow monotonically with time. Default reasoning leads to *logical nonmonotonicity*, i.e., the set of consequences does not grow monotonically with the set of premises. The connection between these forms of nonmonotonicity will be studied in this paper focusing on syntax-based approaches. It is shown that a general form of syntax-based belief revision corresponds to a special kind of partial meet revision in the sense of the theory of epistemic change, which in turn is expressively equivalent to some variants of logics for default reasoning. Additionally, the computational complexity of the membership problem in revised belief sets and of the equivalent problem of derivability in default logics is analyzed, which turns out to be located at the lower end of the polynomial hierarchy.

1 INTRODUCTION

Belief revision is the process of incorporating new information into a knowledge base while preserving consistency. Recently, belief revision has received a lot of attention in AI,[1] which led to a number of different proposals for different applications [Ginsberg, 1986; Ginsberg and Smith, 1987; Dalal, 1988; Gärdenfors and Makinson, 1988; Winslett, 1988; Myers and Smith, 1988; Rao and Foo, 1989; Nebel, 1989; Winslett, 1989; Katsuno and Mendelzon, 1989; Katsuno and Mendelzon, 1990; Doyle, 1990]. Most of this research has been considerably influenced by approaches in philosophical logic, in particular by Gärdenfors and his colleagues [Alchourrón *et al.*, 1985; Gärdenfors, 1988], who developed the *logic of theory change*, also called *theory of epistemic change*, which will be briefly sketched in Section 2.

Syntax-based approaches to belief revision to be introduced in Section 3 have been very popular because of their conceptual simplicity. However, there also has been criticisms since the outcome of a revision operation relies an arbitrary syntactic distinctions (see, e.g., [Dalal, 1988; Winslett, 1988; Katsuno and Mendelzon, 1989])—and for this reason such operations cannot be analyzed on the *knowledge level*. In [Nebel, 1989] we showed that syntax-based approaches can be interpreted as assigning higher relevance to explicitly represented sentences. Based on that view, one particular kind of syntax-based revision, called *base revision*, was shown to fit into the theory of epistemic change. In Section 4 we generalize this result to *prioritized bases* by employing the notion of *epistemic relevance* [Nebel, 1990]—a complete preorder over the set of all derivable sentences.

The resulting revision operations do not satisfy all *rationality postulates* belief revision operations should obey, however. In Section 5 some interesting special cases of epistemic relevance are analyzed that satisfy all rationality postulates. In particular, we show that *epistemic entrenchment* as introduced in [Gärdenfors and Makinson, 1988] is a special case of *epistemic relevance.*

Makinson and Gärdenfors [1990] showed that there is a tight connection between belief revision and nonmonotonic logics. In Section 6 we will strengthen this result. First, we show that the form of *logical nonmonotonicity* observable when revising beliefs is a *necessary* consequence of *temporal nonmonotonicity* induced by belief revision. Second, we will prove that this similarity can be strengthened to equivalence of expressiveness for particular nonmonotonic logics and belief revision operations in the case of propositional logic. Poole's [1988] and Brewka's [1989; 1990] approaches are shown to be expressively equivalent to some forms of syntax-based belief revision approaches. An interesting consequence of this result is that the "absurd belief state"

[1] See also [Brachman, 1990], in which "practical and well-founded theories of belief revision" are called for.

that is inconsistent turns out to be more important than assumed to be in the theory of epistemic change.

Additionally to the logical properties of belief revision and default reasoning, in Section 7 the computational properties are analyzed. As it turns out, propositional syntax-based belief revision and default reasoning is not very much harder than propositional derivability.

2 THE THEORY OF EPISTEMIC CHANGE

In this section we will briefly survey some of the results of the theory of epistemic change in a setting of propositional logic.[2] Throughout this paper, a propositional language $\mathcal{L}$ with the usual logical connectives ($\neg$, $\vee$, $\wedge$, $\rightarrow$ and $\leftrightarrow$) is assumed. The countable alphabet of propositional variables $a, b, c \ldots$ is denoted by Σ, propositional sentences by $v, w, x, y, z, \ldots$, constant truth by $\top$, its negation by $\bot$, and countable sets of propositional sentences by $A, B, C, \ldots$ and $X, Y, Z, \ldots$

The symbol $\vdash$ denotes derivability and Cn the corresponding closure operation, i.e.,

$$Cn(A) \stackrel{\text{def}}{=} \{x \in \mathcal{L} | A \vdash x\}. \qquad (1)$$

Instead of $Cn(\{x\})$, we will also write $Cn(x)$. Deductively closed sets of propositional sentences, i.e., $A = Cn(A)$, are denoted by capital letters from the beginning of the alphabet and are called *belief sets*. Arbitrary sets of sentences are called *belief bases* and are denoted by capital letters from the end of the alphabet. Systems of belief bases and belief sets are denoted by S. Finite belief bases Z are often identified with the conjunction of all propositions $\bigwedge Z$. If $S = \{X_1, \ldots X_n\}$ is a finite family of finite belief bases, then $\bigvee S$ shall denote a proposition logically equivalent to $(\bigwedge X_1) \vee \ldots \vee (\bigwedge X_n)$. As usual, we set $\bigvee \emptyset = \bot$.

In [Gärdenfors, 1988] three operations on belief sets are analyzed, namely, *expansion*, *contraction*, and *revision*. *Expansion* is the addition of a sentence x to a belief set A, written $A + x$, resulting in a new (possibly inconsistent) belief set, defined by

$$A + x \stackrel{\text{def}}{=} Cn(A \cup \{x\}). \qquad (2)$$

Contraction is the removal of a sentence x from a belief set A resulting in a new belief set, denoted by $A \dot{-} x$, that does not contain x (if x is not a tautology), and *revision* is the addition of a sentence x to A, denoted by $A \dot{+} x$, such that $Cn(\bot) \neq A \dot{+} x$ whenever $\not\vdash \neg x$. Although *contraction* and *revision* are not uniquely determined operations—the only commonly agreed criterion is that the changes to the original belief sets have to be *minimal*—it is possible to constrain the space of reasonable change operations. Gärdenfors proposed sets of *rationality postulates*[3] change operations on belief sets should satisfy. The *Gärdenfors postulates for revision* look as follows (A a belief set, x, y propositional sentences):

($\dot{+}$1) $A \dot{+} x$ is a belief set;

($\dot{+}$2) $x \in A \dot{+} x$;

($\dot{+}$3) $A \dot{+} x \subseteq A + x$;

($\dot{+}$4) If $\neg x \notin A$, then $A + x \subseteq A \dot{+} x$;

($\dot{+}$5) $A \dot{+} x = Cn(\bot)$ only if $\vdash \neg x$;

($\dot{+}$6) If $\vdash x \leftrightarrow y$ then $A \dot{+} x = A \dot{+} y$;

($\dot{+}$7) $A \dot{+} (x \wedge y) \subseteq (A \dot{+} x) + y\}$;

($\dot{+}$8) If $\neg y \notin A \dot{+} x$,
then $(A \dot{+} x) + y \subseteq A \dot{+} (x \wedge y)$.

These postulates intend to capture the intuitive meaning of minimal change—from a logical point of view [Alchourrón *et al.*, 1985; Gärdenfors, 1988]. ($\dot{+}$1) states that revision of belief set always results in a belief set. ($\dot{+}$2) formalizes the requirement that revision is always successful. ($\dot{+}$3) gives an upper bound for a revised belief set. It should at most contain the consequences of the original belief set and the new sentence. ($\dot{+}$4) is the conditional converse of ($\dot{+}$3). In case when x is consistent with the original belief set, the revised belief set shall at least contain the original belief base and the new sentence. ($\dot{+}$5) states that inconsistency should be avoided when possible, and ($\dot{+}$6) formalizes the requirement that revision shall be independent from the syntactic form of the sentence the belief set is revised by. While the first six postulates, also called *basic postulates*, are straightforward, the last two postulates are less obvious. They can be interpreted as generalizations of ($\dot{+}$3) and ($\dot{+}$4).

Based on this framework, it is possible to analyze different ways of defining revision operations. In [Alchourrón *et al.*, 1985], so-called partial meet revisions are investigated. This notion is based on systems of maximal (w.r.t. to set-inclusion) subsets of a given belief set A that do not allow the derivation of x, called the *removal of* x and written $A \downarrow x$:

$$A \downarrow x \stackrel{\text{def}}{=} \{B \subseteq A | B \not\vdash x, \forall C : B \subset C \subseteq A \Rightarrow C \vdash x\}. (3)$$

A *partial meet revision* (on A for all x) is defined by a *selection function* $\mathcal{S}$ that selects a nonempty subset of $A \downarrow \neg x$ (provided $A \downarrow \neg x$ is nonempty, $\emptyset$ otherwise) in

[2] The formulation in [Alchourrón *et al.*, 1985; Gärdenfors, 1988] is more general in that only some specific properties are required for the underlying logic.

[3] In order to avoid confusion, one should note that *rationality* in the sense of the theory of epistemic change means an *idealization*: "In this way the rationality criteria serve as regulative ideals. Actual psychological states of belief normally fail to be ideally rational in this sense" [Gärdenfors, 1988, Section 1.2]. Further, this notion of *rationality* is quite different from the notion of economic rationality [Doyle, 1990].

the following way:[4]

$$A \dot{+} x \stackrel{\text{def}}{=} \Big(\bigcap \mathcal{S}(A\downarrow\neg x)\Big) + x. \quad (4)$$

Such partial meet revisions satisfy unconditionally the first six postulates, also called *basic postulates*. Furthermore, it is possible to show that all revision operations satisfying the basic postulates are partial meet revisions [Gärdenfors, 1988, Theorem 4.13]. Actually, this and the other results cited below were proven for contraction. However, if contraction and revision satisfy the basic postulates, they are interdefinable by the Harper (5) and Levi (6) identity:

$$A \dot{-} x = (A \dot{+} \neg x) \cap A, \quad (5)$$
$$A \dot{+} x = (A \dot{-} \neg x) + x. \quad (6)$$

Further, the eight *Gärdenfors postulates for contraction* (see, e.g. [Alchourrón *et al.*, 1985; Gärdenfors, 1988]) are equivalent to the revision postulates under these definitions in the following sense. The first six, seven, or eight contraction postulates are satisfied if and only if the first six, seven or eight, revision postulates for the corresponding revision operation are satisfied, respectively [Gärdenfors, 1988, Theorem 3.2–3.5].

It should be noted that two special cases of partial meet revisions are unreasonable [Alchourrón *et al.*, 1985; Gärdenfors, 1988]. The first special case is that $\mathcal{S}$ always selects all of the elements of $A\downarrow\neg x$—leading to the so-called *full meet revision*. In this case $A \dot{+} x = Cn(x)$ if $\neg x \in A$. Although unreasonable, full meet revision is "fully rational" in the sense that it satisfies all the Gärdenfors postulates, as is easy to verify.

The second special case is that $\mathcal{S}$ always selects *singletons* from $(A\downarrow\neg x)$—resulting in the class of so-called *maxi-choice revisions*. These revision operations have the property that $A \dot{+} x$ is a *complete* belief set—provided that $\neg x \in A$. This means $y \in A \dot{+} x$ or $\neg y \in A \dot{+} x$ for every $y \in \mathcal{L}$. In other words maxi-choice revisions lead to an unmotivated inflation of beliefs.

3 SYNTAX-BASED REVISION APPROACHES

The theory sketched above captures the *logical* portion of minimal change giving us a kind of yardstick to evaluate approaches to belief revision. However, it still leaves open the problem of how to specify additional restrictions so that a revision operation also satisfies a "pragmatic" measure of minimal change.

Two principal points of departure are conceivable. Starting with a belief base as the representation of a belief set, either the syntactic form of the belief base [Fagin *et al.*, 1983; Ginsberg, 1986; Nebel, 1989] or the possible states of the world described by the belief base—the models of the belief base—could be changed minimally [Dalal, 1988; Winslett, 1988; Katsuno and Mendelzon, 1989; Katsuno and Mendelzon, 1990]. The former approach seems to be more reasonable if the belief base corresponds to a body of explicit beliefs that has some relevance, such as a code of norms or a scientific or naive theory which is almost correct. The latter view seems plausible if the application is oriented towards minimal change of the state of the world described by a belief set.

The idea of *changing a description minimally* could be formalized by selecting maximal subsets of the belief base not implying a given sentence. If there is more than one such maximal subset, the intersection of the consequences of these subsets is used as the result. Thus, using $(Z\downarrow x)$ as the set of maximal subsets of Z not implying x as above, *simple base revision*, written as $Z \oplus x$, could be defined as follows [Fagin *et al.*, 1983; Ginsberg, 1986; Nebel, 1989]:

$$Z \oplus x \stackrel{\text{def}}{=} \Big(\bigcap_{Y\in(Z\downarrow\neg x)} Cn(Y)\Big) + x. \quad (7)$$

The operation $\oplus$ considers all sentences in a base as equally relevant. In most applications, however, we want to distinguish between the relevance of sentences (see, e.g., [Fagin *et al.*, 1983; Ginsberg, 1986]). For this purpose, we assume that Z is partitioned into disjoint *priority classes* Z_i, $i \geq 1$, and define the *prioritized removal of* x, written $Z \Downarrow x$, in a way such that sentences in Z_i have higher priority than those in Z_j iff $i < j$:

$$\begin{aligned} Z \Downarrow x \stackrel{\text{def}}{=} \Big\{ & Y \subseteq Z \,|\, Y \not\vdash x, Y = \textstyle\bigcup_{i\geq 1} Y_i, \\ & \forall i \geq 1\colon \Big(Y_i \subseteq Z_i, \\ & \quad \forall X\colon Y_i \subset X \subseteq Z_i \Rightarrow \\ & \quad\quad (\textstyle\bigcup_{j=1}^{i-1} Y_j \cup X) \vdash x\Big)\Big\}. \end{aligned} \quad (8)$$

Intuitively, the elements of $Z \Downarrow x$ are constructed in a stepwise manner starting with Z_1 and selecting as many sentences from Z_i as possible such that the selected sentences from $Z_1, \ldots, Z_i$ do not lead to the derivation of x.

Using $\Downarrow$ instead of $\downarrow$ in the definition (7) leads then to a *prioritized base revision* operation, denoted by $\hat{\oplus}$, with the special case of only one priority class that is identical with simple base revision.

In the interesting special case when we are dealing with *finite* belief bases, the result of a base revision can be finitely represented.

Proposition 1 *If Z is a finite belief base then*

$$Z \hat{\oplus} x = Cn\Big(\big(\bigvee(Z \Downarrow \neg x)\big) \wedge x\Big), \quad (9)$$

[4] Note that all elements of $A\downarrow\neg x$ are belief sets and that the intersection of belief sets is a belief set again.

for every prioritized base revision $\hat{\oplus}$ on Z.[5]

In order to demonstrate how base revision works, let us assume the following scenario. Assume that a suspect tells you that he went to the beach for swimming and assume that you have observed that the sun was shining. Further, you firmly believe that going to the beach for swimming when the sun is shining implies a sun tan. If you then discover that the suspect is not tanned, there is an inconsistency to resolve. Supposing the following propositions:

$$\begin{aligned} b &= \text{“going to the beach for swimming”}, \\ s &= \text{“the sun is shining”}, \\ t &= \text{“sun tan”}, \end{aligned}$$

the situation can be modeled formally by a prioritized base Z:

$$\begin{aligned} Z_1 &= \{(b \wedge s \rightarrow t)\}, \\ Z_2 &= \{s\}, \\ Z_3 &= \{b\}, \\ Z &= Z_1 \cup Z_2 \cup Z_3. \end{aligned}$$

From this belief base t can be derived. If we later observe that $\neg t$, the belief base has to be revised:

$$\begin{aligned} Z \hat{\oplus} \neg t &= \bigcap \Big(Cn(Z \Downarrow t) \Big) + \neg t \\ &= Cn\Big(\bigvee \Big\{ \{(b \wedge s \rightarrow t), s\} \Big\} \Big) + \neg t \\ &= Cn(\{(b \wedge s \rightarrow t), s, \neg t\}). \end{aligned}$$

In particular, we would conclude that b was a lie.

A consequence of the definition of (simple and prioritized) base revision is that for two different belief bases X and Y that have the same meaning, i.e., $Cn(X) = Cn(Y)$, base revision can lead to different results, i.e., $Cn(X \hat{\oplus} x) \neq Cn(Y \hat{\oplus} x)$. Base revision has a "morbid sensitivity to the syntax of the description of the world" [Winslett, 1988], which is considered as an undesirable property. It is argued that revision shall be independent from the syntactical representation of a belief set, that they should be specified on the *knowledge level* [Newell, 1982]. Dalal [1988] formulated the *principle of irrelevance of syntax* which states that a revision shall be independent of the syntactic form of the belief base representing a belief set and of the syntactic form of the sentence that has to be incorporated into the belief set (see also [Katsuno and Mendelzon, 1989]).

Obviously, base revision does not satisfy the principle of irrelevance of syntax—and is not a revision operation in the sense of the theory of epistemic change for this reason. Worse yet, abstracting from the syntactic representation of a belief base and considering the logical equivalent belief set leads to nowhere. Simple base revision applied to belief sets is equivalent to full meet revision, thus, useless. For these reasons, it is argued in [Dalal, 1988; Winslett, 1988; Katsuno and Mendelzon, 1989] that revision shall be performed on the *model-theoretic* level, i.e., by viewing a belief set as the set of models that satisfy a given belief base and by performing revision in a way that selects models that satisfy the new sentence and *differ minimally* from the models of the original belief base. In order to define what the term *minimal difference* means, we have to say something about how models are to be compared, though. In Dalal [1988], for instance, the "distance" between models is measured by the number of propositional variables that have different truth values. Katsuno and Mendelzon [1989] generalize this approach by considering particular orderings over models. In any case, it is impossible to define a revision operation by referring only to logical properties. Some inherently extra-logical, pragmatic preferences are necessary to guide the revision process.

As argued above, for some applications it does not seem to be a bad idea to derive preferences from the syntactic form of the representation of a belief set. Actually, from a more abstract point of view, it is not the particular syntactic form a belief base we are interested in, but it is the fact that we believe that a particular set of sentences is more valuable or justified than another logically equivalent set, and we want to preserve as many of the "valuable" sentences as possible. Using this idea it is possible to reconstruct base revision in the framework of the theory of epistemic change by employing the notion of *epistemic relevance* [Nebel, 1990].

4 EPISTEMIC RELEVANCE AND BASE REVISION

The intention behind base revision is that all the sentences in a belief base X are considered as *relevant*—some perhaps more so than others. For this reason we want to give up as few sentences from X as possible, while with sentences that are only derivable we are more liberal. This idea can be formalized by employing a *complete preorder* with maximal elements, written $x \preceq y$, on the elements of a belief set. In other words, we consider a *reflexive* and *transitive* relation on $Cn(X)$ such that for all $x, y \in Cn(X)$ we have $x \preceq y$ or $y \preceq x$. Further, there exists at least one maximal element x, i.e., for no element y: $x \preceq y$ and $y \not\preceq x$. This relation will be called *epistemic relevance ordering*. It induces an equivalence relation, written $x \simeq y$, as follows:

$$x \simeq y \quad \text{iff} \quad (x \preceq y \text{ and } y \preceq x). \tag{10}$$

The corresponding equivalence classes are denoted by $\overline{x}$ and are called *degrees of epistemic relevance*. Since the preorder is complete, $\preceq$ is a linear order on the

[5] Full proofs for this and the following propositions can be found in [Nebel, 1991].

degrees of epistemic relevance. Further, there exists a maximal such degree because the preorder contains maximal elements. Using the degrees of epistemic relevance, we define a *strict partial ordering* expressing preferences on subsets $X, Y \in 2^A$, written as $X \ll Y$, by

$$X \ll Y \quad \text{iff} \quad \exists \overline{v}\colon \Big((X \cap \overline{v} \subset Y \cap \overline{v}) \text{ and } \forall \overline{w} \not\preceq \overline{v}\colon (X \cap \overline{w} = Y \cap \overline{w}) \Big), \tag{11}$$

which in turn can be used to define a selection function $\mathcal{S}_{\preceq}$ that selects all maximally preferred elements of $A \downarrow x$:

$$\mathcal{S}_{\preceq}(A \downarrow x) \stackrel{\text{def}}{=} \{B \in (A \downarrow x) \mid \forall C \in (A \downarrow x)\colon B \not\ll C\}. \tag{12}$$

Such a selection function may then be used to define a revision operation as done in equation (4). Revisions defined in this way will be called *revisions based on epistemic relevance*. Analyzing the properties of such revisions, we note that they satisfy most of the Gärdenfors postulates.

Theorem 2 *Revisions based on epistemic relevance satisfy* ($\dot{+}$*1*)–($\dot{+}$*7*).

Proof Sketch: First of all, note that any revision based on epistemic relevance is a partial meet revision because $\mathcal{S}_{\preceq}$ selects always a nonempty subset of $(A \downarrow x)$, provided this set is nonempty. By Theorem 4.13 of [Gärdenfors, 1988], any partial meet revision satisfies ($\dot{+}$1)–($\dot{+}$6). Satisfaction of ($\dot{+}$7) follows from the fact that $\mathcal{S}_{\preceq}$ is a *relational* selection function [Gärdenfors, 1988, Lemma 4.14]. ■

Note that revision based on epistemic relevance does not satisfy ($\dot{+}$8) in general. The interesting point about such partial meet revisions is that they correspond to prioritized base revision as defined in Section 3. The deeper reason for this correspondence is that a selection function can be constructed in a way such that it selects the intersection of the consequential closure of a system of sets, as spelled out below.

Lemma 3 *Let A be a belief set, and let x be a sentence such that $\neg x \in A$. Let S be a system of subsets of A, where $Z \not\vdash \neg x$, for all $Z \in S$. Then*

$$\Big(\bigcap\{C \in A \downarrow \neg x \mid \exists Z \in S\colon Z \subseteq C\}\Big) + x = \Big(\bigcap_{Z \in S} Cn(Z)\Big) + x. \tag{13}$$

Using this lemma, the correspondence between revision based on epistemic relevance and prioritized base revision can be easily shown.

Theorem 4 *For any revision operation $\dot{+}$ defined on A and based on epistemic relevance there exists a corresponding prioritized base revision $\hat{\oplus}$ on some base Z and vice versa such that*

$$A \dot{+} x \quad = \quad Z \hat{\oplus} x. \tag{14}$$

Proof Sketch: For the limiting cases $\neg x \notin A$ or $\vdash \neg x$ the theorem holds trivially.

$\Rightarrow$: First, any belief set is a belief base by definition. Second, any epistemic entrenchment ordering defines priority classes on such a base. Third, for any deductively closed prioritized base A, a set X is a member of $(A \Downarrow x)$ if and only if X is maximal in $(A \downarrow x)$ w.r.t. $\ll$. Applying Lemma 3, the conclusion follows.

$\Leftarrow$: Given a prioritized base Z, define an epistemic entrenchment ordering on $Cn(Z)$ such that $v \preceq w$ iff $v \in Cn(Z) - Z$ or $v \in Z_i$, $w \in Z_j$, and $i \geq j$. Now it is easy to show that $B \in \mathcal{S}_{\preceq}(Cn(Z) \downarrow x)$ iff there is a set $X \in (Z \Downarrow x)$ and $X \subseteq B$. Applying again Lemma 3, shows that a revision based on the relevance ordering defined above leads to the same results as the prioritized base revision. ■

This means that prioritized base revision coincides with revision based on epistemic relevance. This abstract view on syntax-based revision may also answer some of the questions raised by Myers and Smith [1988]. They observed that sometimes base revision does not seem to be the appropriate operation because some derived information turns out to be more relevant than the syntactically represented sentences in a belief base, and we get the wrong results when using base revision. However, there is no magic involved here. Base revision leads to the right results only if the syntactic representation really reflects the epistemic relevance. For this reason, the notion of revision based on epistemic relevance seems to be preferable over base revision because it avoids the confusion between surface-level syntactic representation and the intended relevance of propositions.

The question of whether the above defined correspondence can be exploited computationally cannot be answered positively in the general case. For the case of belief sets that are finite modulo logical equivalence, however, revision based on epistemic relevance can be performed by a prioritized base revision on a finite base.

Proposition 5 *Let A be a belief set finite modulo logical equivalence. If $\dot{+}$ is a revision based on epistemic relevance defined on A, then there exists a finite prioritized base Z, such that for all x:*

$$A \dot{+} x \quad = \quad Z \hat{\oplus} x. \tag{15}$$

Although revisions based on epistemic relevance do not satisfy all Gärdenfors postulates, there are special cases that do so. A trivial special case is a revision based on only one degree of epistemic relevance, which is equivalent to full meet revision. There are more interesting cases, however, we will investigate in the following section.

5 EPISTEMIC RELEVANCE AND ENTRENCHMENT

Gärdenfors and Makinson claim that the notion of *epistemic entrenchment* introduced in [Gärdenfors and Makinson, 1988] is closely related to the notion of *database priorities* as proposed in [Fagin *et al.*, 1983]. Since the notion of database priorities is identical to the notion of *priority classes* introduced in Section 3, which in turn coincides with *degrees of epistemic relevance*, one would expect that epistemic entrenchment orderings are closely related to epistemic relevance orderings. Although the intuitions are clearly similar, the question is whether the different formalizations lead to identical result of revision operations.

Epistemic entrenchment orderings, written as $x \preceq_\epsilon y$, are defined over the entire set of sentences $\mathcal{L}$ and have to satisfy the following properties:

($\preceq_\epsilon$1) If $x \preceq_\epsilon y$ and $y \preceq_\epsilon z$, then $x \preceq_\epsilon z$.

($\preceq_\epsilon$2) If $x \vdash y$, then $x \preceq_\epsilon y$.

($\preceq_\epsilon$3) For any x, y, $x \preceq_\epsilon (x \wedge y)$ or $y \preceq_\epsilon (x \wedge y)$.

($\preceq_\epsilon$4) When $A \neq Cn(\bot)$, then $x \notin A$ iff $x \preceq_\epsilon y$ for all $y \in \mathcal{L}$.

($\preceq_\epsilon$5) If $y \preceq_\epsilon x$ for all $y \in \mathcal{L}$, then $\vdash x$.

Using such a relation, Gärdenfors and Makinson define *contraction based on epistemic entrenchment*, written $A \overset{\epsilon}{-} x$, by

$$y \in A \overset{\epsilon}{-} x \text{ iff } y \in A \text{ and } ((x \vee y) \not\preceq_\epsilon x \text{ or } \vdash x) \quad (16)$$

and show that a contraction based on epistemic entrenchment satisfies all rationality postulates for contraction as well as the following condition [Gärdenfors and Makinson, 1988, Theorem 4]:

$$x \preceq_\epsilon y \quad \text{iff} \quad x \notin A \overset{\epsilon}{-} (x \wedge y) \text{ or } \vdash (x \wedge y). \quad (17)$$

Further, they show that any contraction operation satisfying all of the rationality postulates is generated by some epistemic entrenchment ordering [Gärdenfors and Makinson, 1988, Theorem 5].

The question is now how to interpret these results in the framework of epistemic relevance. It follows straightforwardly that the restriction of $\preceq_\epsilon$ to the sentences in a belief set can be considered as an epistemic relevance ordering as defined in the previous section. Further, in this case, using interdefinability of revision and contraction, definition (16) coincides with a contraction operation that is defined by using the Harper identity (5) and a revision operation based on a epistemic relevance ordering that satisfies ($\preceq_\epsilon$1)–($\preceq_\epsilon$5).

Theorem 6 *Suppose a belief set A, an epistemic entrenchment ordering $\preceq_\epsilon$, and a contraction operation $\overset{\epsilon}{-}$ based on $\preceq_\epsilon$. Let $\preceq$ be an epistemic relevance ordering that is the restriction of $\preceq_\epsilon$ to A, and let $\dot{+}$ be a revision based on the epistemic relevance ordering $\preceq$. Then*

$$A \overset{\epsilon}{-} x \quad = \quad (A \dot{+} \neg x) \cap A. \quad (18)$$

Proof Sketch: For the limiting cases $\vdash x$ and $x \notin A$ the theorem holds trivially. For the principal case, one straightforwardly verifies that

$$y \in \bigcap \mathcal{S}_{\preceq}(A \downarrow x) \text{ iff } y \in A \text{ and } (x \vee y) \not\preceq_\epsilon x. \quad (19)$$

From this the conclusion follows immediately. ■

Thus, the notion of epistemic entrenchment can indeed be viewed as a refined special case of database priorities or epistemic relevance. It is not obvious, however, how to arrive at such epistemic entrenchment orderings. While epistemic relevance can be easily derived from given priority classes, it is not clear whether there are natural ways to generate epistemic entrenchment orderings. In [Gärdenfors and Makinson, 1988] it is proposed to start with a complete ordering over the maximal disjunctions derivable from a belief set. Despite the fact that this does not sound very "natural", it also implies that a large amount of information has to be supplied, sometimes too much (see Section 7), in order to change a belief set.

Interestingly, there is another special case of epistemic relevance that leads to a revision operation that satisfies all postulates. When all priority classes of a prioritized belief base are singletons, then the prioritized base revision (as well as the corresponding partial meet revision and the epistemic relevance ordering) are called *unambiguous*.

Proposition 7 *Let Z be a prioritized belief base such that all priority classes are singletons. Then $(Z \Downarrow x)$ is a singleton iff $\not\vdash x$.*

Clearly, the corresponding epistemic relevance ordering is not necessarily an epistemic entrenchment ordering. Nevertheless, unambiguous partial meet revisions satisfy all rationality postulates.

Theorem 8 *Let $\preceq$ be an unambiguous epistemic relevance ordering. Then the revision based on this ordering satisfies all Gärdenfors postulates.*

Proof Sketch: By Theorem 2 ($\dot{+}$1)–($\dot{+}$7) are satisfied. If $\neg x \notin A$, then ($\dot{+}$8) holds because of ($\dot{+}$4). If $\vdash \neg x$, then ($\dot{+}$8) holds because the "if" part of ($\dot{+}$8) is never true. For the principal case we make use of Proposition 7 and Theorem 4 and show that for the unambiguously prioritized base Z:

$$\begin{aligned} &\text{If } \{X\} = (Z \Downarrow \neg x) \text{ and} \\ &\quad X \cup \{x\} \not\vdash \neg y \text{ and} \\ &\quad \{Y\} = (Z \Downarrow (\neg x \vee \neg y)) \\ &\text{then } X \subseteq Y. \end{aligned} \quad (20)$$

Assume that (20) is violated, i.e., there exists a smallest index j such that $X \cap \bigcup_{i=1}^{j} Z_i \supset Y \cap \bigcup_{i=1}^{j} Z_i$. This, however, would mean that $X \supseteq ((Y \cap \bigcup_{i=1}^{j} Z_i) \cup Z_j) \vdash \neg x \vee \neg y$. Thus, $X \vdash x \rightarrow \neg y$, which is equivalent to $X \cup \{x\} \vdash \neg y$ contradicting the precondition in (20). ∎

Although an unambiguous relevance ordering is not necessarily an entrenchment ordering, using (17) it is possible to specify an epistemic entrenchment ordering that leads to an identical revision operation since unambiguous revisions are fully rational.

6 BELIEF REVISION AND DEFAULT REASONING

Doyle remarked in [Doyle, 1990, App. A] that "the adjective 'nonmonotonic' has suffered much careless usage recently in artificial intelligence, and the only thing common to many of its uses is the term 'nonmonotonic' itself." Doyle identified two principal ideas behind the use of this term, namely,

> [...] that attitudes are gained and lost over time, that reasoning is nonmonotonic—this we call *temporal* nonmonotonicity—and that unsound assumptions can be the deliberate product of sound reasoning, incomplete information, and a "will to believe"—which we call *logical* nonmonotonicity.

Although these two forms of nonmonotonicity should not be confused, they are intimately connected. In particular, the temporal nonmonotonicity induced by belief revision, i.e., the fact that in general we do not have $A \subseteq A \dot{+} x$, is connected with logical nonmonotonicity induced by some forms of default reasoning.

When reasoning with defaults in a setting as described in [Poole, 1988; Brewka, 1989], we are prepared to "drop" some of the defaults if they are inconsistent with the facts. This, however, is quite similar to what we are doing when revising beliefs in the theory of epistemic change. Propositions of a theory are given up when they are inconsistent with new facts. Since default reasoning leads to logical nonmonotonicity, one would expect that belief revision is nonmonotonic in the facts to be added, i.e., we would expect that $Cn(x) \subseteq Cn(y)$ does not imply $A \dot{+} x \subseteq A \dot{+} y$. Conversely, requiring monotony in the second operand of a belief revision operation is impossible in the general case.

Proposition 9 *Let $\dot{+}$ be a belief revision operation defined on a belief set A. If for all x, y*

$$A \dot{+} x \subseteq A \dot{+} y \quad \textit{if} \quad Cn(x) \subseteq Cn(y), \tag{21}$$

then

1. *$A = Cn(\emptyset)$ and $A \dot{+} x = Cn(x)$, or*
2. *$A = Cn(\bot)$ and $A \dot{+} x = Cn(x)$, or*
3. *$\dot{+}$ violates at least one of the basic Gärdenfors postulates.*

Makinson and Gärdenfors [1990] use this similarity of logical nonmonotonicity and the nonmonotonicity of belief revision in the second operand as a starting point to investigate the relationship between nonmonotonic logics and belief revision on an very general and abstract level. They compare various general conditions on nonmonotonic provability relations with the Gärdenfors postulates.

For the approaches to belief revision described in the previous section there is an even stronger connection to nonmonotonic logic. Prioritized base revision, and hence partial meet revision based on epistemic relevance, is expressively equivalent to *skeptical* provability[6] in Poole's [1988] *theory formation* approach and Brewka's [1989] *level default theories* (LDT)—for the case of finitary propositional logic.

A common generalization of both approaches are *ranked default theories* (RDT). A RDT Δ is a pair $\Delta = (\mathcal{D}, \mathcal{F})$, where $\mathcal{D}$ is a finite sequence $\langle \mathcal{D}_1, \ldots, \mathcal{D}_n \rangle$ of finite sets of sentences (propositional, in our case) interpreted as ranked defaults and $\mathcal{F}$ is a finite set of sentences interpreted as hard facts.

An *extension* of Δ is a deductively closed set of propositions $E = Cn((\bigcup_{i=1}^{n} \mathcal{R}_i) \cup \mathcal{F})$ such that for all i with $1 \leq i \leq n$:

1. $\mathcal{R}_i \subseteq \mathcal{D}_i$,
2. $(\bigcup_{j=1}^{i} \mathcal{R}_j) \cup \mathcal{F}$ is consistent,[7] and
3. $(\bigcup_{j=1}^{i} \mathcal{R}_j)$ is set-inclusion maximal.

A sentence x is *strongly provable* in Δ, written $\Delta \mid\!\sim x$, iff for all extensions E of Δ: $x \in E$.

Poole's approach is a special case of RDT's where $\mathcal{D} = \langle \mathcal{D}_1 \rangle$, and Brewka's LDT's are RDT's with $\mathcal{F} = \emptyset$. Note, however, that the expressive difference between RDT's and LDT's is actually very small and shows up only if $\mathcal{F}$ is inconsistent. In this case, RDT's allow the derivation of $\bot$ while this is impossible in LDT's.

Theorem 10 *Let $\Delta = (\langle \mathcal{D}_1, \ldots, \mathcal{D}_n \rangle, \mathcal{F})$ be a* RDT. *Let $Z = \bigcup_{i=1}^{n} \mathcal{D}_i$ be a prioritized base with priority classes $\mathcal{D}_1, \ldots, \mathcal{D}_n$. Then*

$$\Delta \mid\!\sim x \quad \textit{iff} \quad x \in (Z \hat{\oplus} \mathcal{F}). \tag{22}$$

[6] A correspondence to *credulous* derivability can be achieved if a notion of *nondeterministic* revision as proposed in [Doyle, 1990] is adopted.

[7] Note that this definition, which is similar to the definition of an extension in [Poole, 1988], excludes inconsistent extensions. Nevertheless, the definition of strong provability implies that $\bot$ can be derived iff $\mathcal{F}$ is inconsistent.

Proof Sketch: In the limiting case when $\mathcal{F} \vdash \bot$, $Z \hat{\oplus} \mathcal{F} = Cn(\bot)$. Further, in this case there is no extension of Δ, hence $\Delta \mid\!\sim x$ for all $x \in \mathcal{L}$ by the definition of strong provability.

When $\mathcal{F}$ is consistent, it is easy to see that E is an extension of Δ if and only if there is a belief base $X \in (Z \Downarrow \neg(\bigwedge \mathcal{F}))$ such that $E = Cn(X \cup \mathcal{F})$. ■

This means that finite ranked default theories have the same expressive power as prioritized base revision operations, which coincide with revisions based on epistemic relevance by Theorem 4.

It should be noted that in ranked default theories there is no requirement on the *internal* consistency of defaults. This means that the set $\bigcup_i \mathcal{D}_i$ may very well be inconsistent. In Theorem 10 that may lead to $\bot \in Cn(Z)$, i.e., the belief set to be revised is inconsistent. Although this might sound unreasonable in the context of modeling (idealized) epistemic states—in fact, inconsistency is indeed explicitly excluded by requirement (2.2.1) in [Gärdenfors, 1988]—it does not lead to technical problems in the theory of epistemic change. Additionally, it is possible to give a transformation between reasoning in RDT's and prioritized base revision using only consistent belief sets.

Corollary 11 *Let Δ be a* RDT *as above. Then there exists a consistent prioritized base Z and a proposition y such that*

$$\Delta \mid\!\sim x \quad \textit{iff} \quad x \in \Big(Z \hat{\oplus} (y \wedge \mathcal{F})\Big). \tag{23}$$

Proof Sketch: Define Z as in Theorem 10. Transform every sentence in Z into negation normal form (i.e., into a formula such that negation signs appear only in front of propositional variables), replace any negative literal $\neg a$ in all sentences of Z by a fresh variable a' and define (assuming w.l.g. that Σ is finite):

$$y \quad \stackrel{\text{def}}{=} \quad \bigwedge_{a \in \Sigma} (\neg a \leftrightarrow a'). \tag{24}$$

This ensures that Z is consistent and that the result of the revision is the same as in Theorem 10. ■

From the results above and the translation of ($\dotplus$8) to a condition on nonmonotonic derivability relations in [Makinson and Gärdenfors, 1990], it follows that the derivability relation of RDT's w.r.t. the set of hard facts $\mathcal{F}$ does not satisfy the condition of *rational monotony* (see [Makinson and Gärdenfors, 1990]). Note that this result depends on the *exact* correspondence between RDT's and belief revision based on epistemic relevance. In [Makinson and Gärdenfors, 1990; Gärdenfors, 1990] the correspondence between Poole's logic and belief revision was only approximate because the defaults were assumed to be deductively closed.

Another interesting observation in this context is that the addition of *constraints* to RDT's is similar but not identical to a contraction operation as defined in Section 2. Poole [1988] introduced *constraints*—another set of sentences—as a means to restrict the applicability of defaults. A *ranked default theory with constraints* is a triple $\Delta = (\mathcal{D}, \mathcal{F}, \mathcal{C})$, where $\mathcal{D}$ and $\mathcal{F}$ are defined as above and $\mathcal{C}$ is a finite set of sentences interpreted as constraints. The notion of an extension is modified as follows. Instead of condition 2. it is required that

2. $(\bigcup_{i=1}^{n} \mathcal{R}_i) \cup \mathcal{F} \cup \mathcal{C}$ is consistent.

Provided the set $\mathcal{F} \cup \mathcal{C}$ is consistent, which is the interesting case, skeptical derivability can be modeled as a form of contraction on *belief bases* (see [Nebel, 1989]), followed by an expansion.

Theorem 12 *Let $\Delta = (\langle \mathcal{D}_1, \ldots, \mathcal{D}_n \rangle, \mathcal{F}, \mathcal{C})$ be an* RDT *with constraints such that $\mathcal{F} \cup \mathcal{C}$ is consistent. Let $Z = \bigcup_{i=1}^{n} \mathcal{D}_i$ be a prioritized base with $\mathcal{D}_1, \ldots, \mathcal{D}_n$ the priority classes of Z. Then*

$$\Delta \mid\!\sim x \quad \textit{iff} \quad \bigvee \Big(Z \Downarrow \neg(\mathcal{F} \wedge \mathcal{C})\Big) \wedge \mathcal{F} \vdash x \tag{25}$$

Proof Sketch: Assuming that $\mathcal{F} \cup \mathcal{C}$ is consistent, any extension E of Δ is the set of consequences of $\mathcal{F}$ and a set $X = \bigcup_{i=1}^{n} \mathcal{R}_i$ such that for all $j \leq n$: $\bigcup_{i=1}^{j} \mathcal{R}_i$ is consistent with $\mathcal{F} \cup \mathcal{C}$ and maximal. This, however, is by definition equivalent to $X \in (Z \Downarrow \neg(\mathcal{F} \wedge \mathcal{C}))$. ■

This means that contrary to the opinion that there is no counter-part to contraction in nonmonotonic logics as spelled out in [Makinson and Gärdenfors, 1990], default reasoning with constraints in Poole's theory formation approach can be modeled by using an operation similar to contraction. However, this similarity does not apply to contraction on belief sets. Assuming that the defaults are internally inconsistent, i.e., $Z \vdash \bot$, application of the Harper identity (5) leads to

$$A \dot{-} \neg x \quad = \quad A \dotplus x, \tag{26}$$

which means that the constraints become part of every extension violating the intention behind introducing constraints. This does not happen when contracting a base because base contraction removes more beliefs than contraction on belief sets (see also [Nebel, 1989]).

7 COMPUTATIONAL COMPLEXITY

For the investigation of the computational complexity of belief revision, we consider the problem of determining membership of a sentence y in a belief set $A = Cn(Z)$ revised by x, i.e.,

$$y \in A \dotplus x. \tag{27}$$

As the input size we use the sum of the size $|Z|$ of the belief base Z that represents A and the sizes $|x|$ and $|y|$ of the sentences x and y, respectively.

This assumption implies that the representation of the preference relation used to guide the revision process should be polynomially bounded by $|Z|+|x|+|y|$. Although this sounds like a reasonable restriction, it is not met by all belief revision approaches. Belief revision based on *epistemic entrenchment* orderings [Gärdenfors and Makinson, 1988], for instance, requires more preference information in the general case. An epistemic entrenchment ordering over all elements of a belief set can be uniquely characterized by an *initial* complete order over the set of all derivable *maximal disjunctions* (over all literals) [Gärdenfors and Makinson, 1988, Theorem 7]. This set is logarithmic in the size of the set of formulas (modulo logical equivalence) in a *belief set*. However, the number of maximal disjunctions may still be very large, namely, exponential in the size of a *belief base*.

A similar statement could be made about revisions based on epistemic relevance. However, if we consider only complete preorders over Z with the understanding that the degree of least relevant sentences is just $Cn(Z) - Z$, then the ordering is represented in a way that is polynomially bounded by $|Z|$ and $\dot{+}$ can be computed by using the corresponding prioritized base revision.

Analyzing the computational complexity of the belief revision problems, the first thing one notes that deciding the trivial case $y \in Cn(\emptyset) \dot{+} x$ is already co-NP-complete,[8] and we might give up immediately. However, finding a characterization of the complexity that is more fine grained than just saying it is NP-hard can help to understand the structure of the problem better. In particular, we may be able to compare the inherent complexity of different approaches and, most importantly, we may say something about feasible implementations, which most likely will make compromises along the line that the expressiveness of the logical language is restricted and/or incompleteness is tolerated at some point. For this purpose we have to know, however, what the sources of complexities are.

The belief revision problems considered in this paper fall into complexity classes located at the lower end of the *polynomial hierarchy*. Since this notion is not as common as the central complexity classes, it will be briefly sketched [Garey and Johnson, 1979, Sect. 7.2]. Let X be a class of decision problems. Then P^{X} denotes the class of decision problems $L \in \mathsf{P}^{\mathsf{X}}$ such that there is a decision problem $L' \in \mathsf{X}$ and a polynomial Turing-reduction from L to L', i.e., all instances of L can be solved in polynomial time on a Turing machine that employs an oracle for L'. Similarly, NP^{X} denotes the class of decision problems $L \in \mathsf{NP}^{\mathsf{X}}$ such that there is nondeterministic Turing-machine that solves all instances of L in polynomial time using an oracle for $L' \in \mathsf{X}$. Based on these notions, the sets Δ_k^p, Σ_k^p, and Π_k^p are defined as follows:[9]

$$\Delta_0^p = \Sigma_0^p = \Pi_0^p = \mathsf{P}, \quad (28)$$

$$\Delta_{k+1}^p = \mathsf{P}^{\Sigma_k^p}, \quad (29)$$

$$\Sigma_{k+1}^p = \mathsf{NP}^{\Sigma_k^p}, \quad (30)$$

$$\Pi_{k+1}^p = \text{co-}\Sigma_{k+1}^p. \quad (31)$$

Thus, $\Sigma_1^p = \mathsf{NP}$, $\Pi_1^p = \text{co-}\mathsf{NP}$, and Δ_2^p is the set of NP-easy problems. Further note that $\bigcup_{k\geq 0} \Delta_k^p = \bigcup_{k\geq 0} \Sigma_k^p = \bigcup_{k\geq 0} \Pi_k^p \subseteq \mathsf{PSPACE}$.

The role of the "canonical" complete problem (w.r.t. polynomial transformability), which is played by SAT for Σ_1^p, is played by k-QBF for Σ_k^p. k-QBF is the problem of deciding whether the following quantified boolean formula is true:

$$\underbrace{\exists \vec{a}\, \forall \vec{b} \ldots}_{k \text{ alternating quantifiers starting with } \exists} F(\vec{a}, \vec{b}, \ldots). \quad (32)$$

Turning now to the revision operations discussed in this paper, we first of all notice that the special belief revision problem of determining membership for a full meet revision, called FMR-problem, is comparably easy. With respect to Turing-reducibility, there is actually no difference to the complexity of ordinary propositional derivability, i.e., the FMR-problem is NP-equivalent.

Proposition 13 FMR $\in \Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$ *provided* $\Sigma_1^p \neq \Pi_1^p$.

Proof Sketch: If $\dot{+}$ is a full meet revision, $x \in Cn(Z) \dot{+} y$ can be solved by the following algorithm:

if $Z \not\vdash \neg x$
  then $Z \cup \{x\} \vdash y$
  else $x \vdash y$

From this, membership in Δ_2^p follows.

Further, SAT can be polynomially transformed to FMR by solving $x \in Cn(x) \dot{+} \top$, and unsatisfiability ($\overline{\mathsf{SAT}}$) can be polynomially transformed to FMR by solving $\bot \in Cn(\emptyset) \dot{+} x$. Hence, assuming FMR $\in$ NP $\cup$ co-NP would lead to NP = co-NP. ∎

The membership problem for simple base revision will be called SBR-problem. This problem is obviously more complicated than the FMR-problem. However, the added complexity is not overwhelming—from a theoretical point of view.

Theorem 14 SBR *is* Π_2^p*-complete.*

[8] We assume some familiarity with the basic notions of the theory of NP-completeness as presented in the first few chapters of [Garey and Johnson, 1979].

[9] The superscript p is only used to distinguish these sets from the analogous sets in the Kleene hierarchy.

Proof Sketch: We will prove that the complementary problem $Z \oplus x \not\vdash y$, which is called $\overline{\mathsf{SBR}}$, is Σ_2^p-complete. Hardness is shown by a polynomial transformation from 2-QBF to $\overline{\mathsf{SBR}}$. Let $\vec{a} = a_1, \ldots, a_n$, let $\vec{b} = b_1, \ldots, b_m$, and let $\exists \vec{a}\, \forall \vec{b}\, F(\vec{a}, \vec{b})$ be an instance of 2-QBF. Now set

$$Z = \{a_1, \ldots, a_n, \neg a_1, \ldots, \neg a_n, \neg F(\vec{a}, \vec{b})\}. \quad (33)$$

Now it is easy to see that

$$Z \oplus \top \not\vdash \neg F(\vec{a}, \vec{b}) \quad \text{iff} \quad \exists \vec{a}\, \forall \vec{b}\, F(\vec{a}, \vec{b}) \text{ is true.} \quad (34)$$

Membership in Σ_2^p follows from the following algorithm that needs nondeterministic polynomial time using an oracle for SAT:

1. Guess a set $Y \subseteq Z$.
2. Verify $Y \cup \{x\} \not\vdash x$.
3. Verify that there is no $z \in Z - Y$ such that $Y \cup \{z\} \not\vdash \neg x$.

■

This means that SBR is, on one hand, not much more difficult than FMR, and, on the other hand, apparently easier than derivability in the modal logics (e.g., K, T, and $S4$), which is a PSPACE-complete problem [Garey and Johnson, 1979, p. 262]. Asking for the computational significance of this result, the answer is somewhat unsatisfying. All problems in the polynomial hierarchy have the same property as the NP-complete problems, namely, that they can be solved in polynomial time if and only if P = NP. Further, all problems in the polynomial hierarchy can be solved by an exhaustive search that takes exponential time. However, from the structure of the algorithm used in the proof one sees that even if we restrict ourselves to polynomial methods for computing propositional satisfiability, there would still be the problem of determining the maximal consitent subsets Y.

Having now a very precise idea of the complexity of the SBR-problem, we may ask what the computational costs of introducing priorities are. In other words whether the membership problem for prioritized base revision, called PBR-problem, is more difficult than SBR.

Theorem 15 PBR *is* Π_2^p*-complete.*

Proof Sketch: Π_2^p-hardness is immediate by Theorem 14. Membership in Π_2^p follows from the fact that the algorithm used in the proof above can be easily adapted by guessing and verifying maximality for every priority class. ■

This means that we do not have to pay for introducing priority classes. In the case of default logics, the generalization from Poole's logic to RDT's does not increase the computational costs. Note also, that the computational complexity of derivability for Brewka's LDT's is not easier because the reduction in the proof of Theorem 14 applies to the special case $\mathcal{F} = \emptyset$, as well.

The membership problem for unambiguous prioritized base revision, the UBR-problem, turns out to be easier than SBR and PBR.

Theorem 16 $\mathsf{UBR} \in \Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$, *provided* $\Sigma_1^p \neq \Pi_1^p$.

Proof Sketch: The nondeterministic algorithm used in the proof above can be obviously modified so that it runs in deterministic polynomial time if the priority classes are singletons. Hence, $\mathsf{UBR} \in \Delta_2^p$. Further, $\mathsf{UBR} \notin \mathsf{NP} \cup \mathsf{co\text{-}NP}$ if $\mathsf{NP} \neq \mathsf{co\text{-}NP}$ follows from Proposition 13 because FMR is a special case of UBR. ■

From the proof, we can infer that if we can come up with a polynomial algorithm for satisfiability (by restricting the language, for instance), then unambiguous base revision will be itself polynomial. This result gives a formal justification for the claim made in [Nebel, 1989] that this form of revision is similar to the functionality the RUP system [McAllester, 1982] offers—in an abstract sense, though.[10] The important point to note is that a feasible implementation of belief revision is possible if we restrict ourselves to polynomial methods for satisfiability by restricting the language or by tolerating incompleteness *and* by using a polynomial method for selecting among competing alternatives.

Finally, it may be interesting to compare syntax-based revision approaches with model-based approaches, such as the one proposed by Dalal [1988]. In order to do so, we first need some definitions. A *truth-assignment* $\mathcal{I}$ is a function $\mathcal{I}: \Sigma \to \{\mathsf{T}, \mathsf{F}\}$. A *model* $\mathcal{I}$ of a belief base Z is a truth-assignment that satisfies all propositions in Z in the classical sense, written $\models_{\mathcal{I}} Z$. $mod(Z)$ denotes the set of all models of Z. $\delta(\mathcal{I}, \mathcal{J})$ denotes the number of propositional variables such that $\mathcal{I}$ and $\mathcal{J}$ map them to different truth-values. $g^m(\mathcal{M})$ is the set of truth assignments $\mathcal{J}$ such that there is a truth-assignment $\mathcal{I} \in \mathcal{M}$ with $\delta(\mathcal{J}, \mathcal{I}) \leq m$. If Z is a finite belief base, then $G^m(Z)$ is some belief base such that $mod(G^m(Z)) = (g^m(mod(Z)))$. Although G^m is not a deterministic function, all possible results are obviously logically equivalent.

Now, model based revision, written $Z \circ x$ is defined

[10]The RUP systems provides the possibility to put premises into different likelihood classes. However, it seems to be the case that in resolving inconsistencies it could select non-maximal sets w.r.t. $\ll$ [McAllester, 1990, personal communication].

by:[11]

$$Z \circ x \stackrel{\text{def}}{=} \begin{cases} G^m(Z) \cup \{x\} & \text{for the least } m \text{ s.t.} \\ & G^m(Z) \cup \{x\} \not\vdash \bot \\ \{x\} & \text{if } Z \vdash \bot \text{ or } x \vdash \bot. \end{cases} \quad (35)$$

Interestingly, the membership problem for model-based revision, called MBR-problem, has the same complexity as UBR and FMR. However, it is not obvious whether a restriction of the expressiveness of the logical language would lead to a polynomial algorithm in this case.

Theorem 17 MBR $\in \Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$, *provided* $\Sigma_1^p \neq \Pi_1^p$.

Proof Sketch: Note that for any n, $G^n(Z) \not\vdash x$ is a problem that can be solved in nondeterministic polynomial time by guessing two truth assignment $\mathcal{I}, \mathcal{J}$ and verifying in polynomial time that

1. $\models_{\mathcal{I}} Z$,
2. $\not\models_{\mathcal{J}} x$, and
3. $\delta(\mathcal{I}, \mathcal{J}) \leq n$.

Further, solving $G^n(Z) \cup \{x\} \not\vdash y$ for fixed n also requires only nondeterministic polynomial time. Hence, MBR $\in \Delta_2^p$. Since SAT and $\overline{\text{SAT}}$ can be polynomially transformed to MBR, MBR $\in$ NP $\cup$ co-NP only if NP = co-NP. ■

Reconsidering the complexity results, there appears to be an interesting pattern. Note that the best result for a belief revision problem we can hope for is membership in Δ_2^p because the problem involves consistency and inconsistency problems. While, the revision operations satisfying all the Gärdenfors postulates, namely, FMR, UBR, and MBR (for the latter see [Dalal, 1988]) turn out to be in this class, the revision operations that do not satisfy ($\dot{+}$8) cannot be shown to be in this class. An interesting question is what conditions are in fact responsible for membership in Δ_2^p.

8 SUMMARY AND OUTLOOK

Syntax-based approaches to belief revision, the class of partial meet revisions based on epistemic relevance, and ranked default theories (RDT's)—a generalization of Poole's and Brewka's approaches to default reasoning—turn out to be strictly equivalent in the case of finitary propositional logic. One of the consequences is that RDT's do not satisfy rational monotony w.r.t. the set of hard facts. Further, we are able to apply the complexity results for belief revision directly to reasoning in RDT's.

[11] This definition is a slight extension of the definition given in [Dalal, 1988] that takes also care of the limiting cases when Z or x is inconsistent.

The complexity results for revision and default reasoning confirm the intuition that unambiguous prioritized base revision is not harder but apparently less complex than general prioritized base revision [Doyle, 1990, Sect. 3.2], which in turn is not harder than simple base revision. An interesting point is that model-based revision as proposed by Dalal is still NP-easy.

One of the open questions is, whether the correspondence between belief revision and nonmonotonic logic holds for the infinite case as well. However, for this purpose the theory of epistemic change has to be extended so that belief sets cannot only be revised by sentences but also by other belief sets. Another interesting question in this context is whether there is a natural condition on belief revision and nonmonotonic consequence operations that characterizes syntax-based approaches completely.

Finally, the observation that all "fully rational" revision operations analyzed in this paper share the property of being NP-easy suggests to analyze this class of revision operations in more detail in order to detect interesting tractable special cases.

Acknowledgement

I would like to thank Gerd Brewka, Jon Doyle, David Makinson, Peter Gärdenfors, and David McAllester for discussions about the subject of this paper and Gerd and David Makinson for comments on an earlier draft.

This work was supported by the German Ministry for Research and Technology BMFT under contract ITW 8901 8 as part of the WIP project.

References

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

[Brachman, 1990] Ronald J. Brachman. The future of knowledge representation. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 1082–1092, Boston, Mass., August 1990.

[Brewka, 1989] Gerhard Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1043–1048, Detroit, Mich., August 1989.

[Brewka, 1990] Gerhard Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense.* Cambridge University Press, Cambridge, England, 1990. To appear.

[Dalal, 1988] Mukesh Dalal. Investigations into a theory of knowledge base revision: Preliminary report.

In *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, pages 475–479, Saint Paul, Minn., August 1988.

[Doyle, 1990] Jon Doyle. Rational belief revision. Presented at the Third International Workshop on Nonmonotonic Reasoning, Stanford Sierra Camp, Cal., June 1990.

[Fagin *et al.*, 1983] Ronald Fagin, Jeffrey D. Ullman, and Moshe Y. Vardi. On the semantics of updates in databases. In *2nd ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, pages 352–365, Atlanta, Ga., 1983.

[Gärdenfors and Makinson, 1988] Peter Gärdenfors and David Makinson. Revision of knowledge systems using epistemic entrenchment. In M. Vardi, editor, *Proceedings of the 2nd Workshop on Theoretical Aspects of Reasoning about Knowledge*. Morgan Kaufmann, Los Altos, Cal., 1988.

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux—Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Mass., 1988.

[Gärdenfors, 1990] Peter Gärdenfors. Belief revision and nonmonotonic logic: Two sides of the same coin? In L. C. Aiello, editor, *Proceedings of the 9th European Conference on Artificial Intelligence*, pages 768–773, Stockholm, Sweden, August 1990.

[Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, Cal., 1979.

[Ginsberg and Smith, 1987] Matthew L. Ginsberg and David E. Smith. Reasoning about action I: A possible worlds approach. In F. M. Brown, editor, *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*, pages 233–258. Morgan Kaufmann, Los Altos, Cal., 1987.

[Ginsberg, 1986] Matthew L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30(1):35–79, October 1986.

[Katsuno and Mendelzon, 1989] Hirofumi Katsuno and Alberto O. Mendelzon. A unified view of propositional knowledge base updates. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1413–1419, Detroit, Mich., August 1989.

[Katsuno and Mendelzon, 1990] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. Technical Report KRR-TR-90-6, University of Toronto, Computer Science Department, Toronto, Ont., August 1990.

[Makinson and Gärdenfors, 1990] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In A. Fuhrmann and M. Morreau, editors, *Proceeding of the Konstanz Workshop on Belief Revision*. Springer-Verlag, Berlin, Germany, 1990. To appear.

[McAllester, 1982] David A. McAllester. Reasoning utility package user's manual. AI Memo 667, AI Laboratory, Massachusetts Institute of Technology, Cambridge, Mass., April 1982.

[Myers and Smith, 1988] Karen L. Myers and David E. Smith. The persistence of derived information. In *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, pages 496–500, Saint Paul, Minn., August 1988.

[Nebel, 1989] Bernhard Nebel. A knowledge level analysis of belief revision. In R. J. Brachman, H. J. Levesque, and R. Reiter, editors, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 301–311, Toronto, Ont., May 1989.

[Nebel, 1990] Bernhard Nebel. *Reasoning and Revision in Hybrid Representation Systems*, volume 422 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany, 1990.

[Nebel, 1991] Bernhard Nebel. Belief revision and default reasoning: syntax-based approaches. DFKI Report, German Research Center for Artificial Intelligence (DFKI), Saarbrücken, West Germany, 1991. To appear.

[Newell, 1982] Allen Newell. The knowledge level. *Artificial Intelligence*, 18(1):87–127, 1982.

[Poole, 1988] David Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

[Rao and Foo, 1989] Anand S. Rao and Norman Y. Foo. Formal theories of belief revision. In R. J. Brachman, H. J. Levesque, and R. Reiter, editors, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 369–380, Toronto, Ont., May 1989.

[Reiter, 1980] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1): 81–132, April 1980.

[Winslett, 1988] Marianne S. Winslett. Reasoning about action using a possible models approach. In *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, pages 89–93, Saint Paul, Minn., August 1988.

[Winslett, 1989] Marianne S. Winslett. Sometimes updates are circumscription. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 859–863, Detroit, Mich., 1989.