Computer Science, p. 52–88. Cambridge University Press, Cambridge, UK, 1992

Syntax-Based Approaches to Belief Revision*

BERNHARD NEBEL

German Research Center for Artificial Intelligence (DFKI),
D-6600 Saarbrücken, Germany

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, Smith 1987; Dalal 1988; Gärdenfors, Makinson 1988; Winslett 1988; Myers, Smith 1988; Rao, Foo 1989; Nebel 1989; Winslett 1989; Katsuno, Mendelzon 1989; Katsuno, 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, Gärdenfors, Makinson 1985; Gärdenfors 1988), who developed the *logic of theory change*, also called *theory of epistemic change*. This theory formalizes *epistemic states* as deductively closed theories and defines different change operations on such epistemic states.

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, 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*. It will be shown that the class of prioritized base revisions is identical with the class of belief revision operations generated by *epistemic relevance orderings* (Nebel 1990).

The belief revision operations generated by epistemic relevance orderings do not sa-

---

*This chapter is a revised and extended version of a paper with the title "Belief Revision and Default Reasoning: Syntax-Based Approaches" in J. A. Allen, R. Fikes, and E. Sandewall (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, Morgan Kaufmann, San Mateo, CA, 1991, pp. 417–428.

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

tisfy all *AGM postulates* belief revision operations should obey, however (see (Gärdenfors, this book)). In Section 5 some interesting special cases of epistemic relevance are analyzed that lead to the satisfaction of all AGM postulates. In particular, we show that *epistemic entrenchment* as introduced in (Gärdenfors, 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" 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, the complexity of propositional syntax-based belief revision is located at the lower end of the polynomial hierarchy.

2 FORMAL PRELIMINARIES

Throughout this chapter, a propositional language $\mathcal{L}$ with the usual logical connectives ($\neg$, $\vee$, $\wedge$, $\rightarrow$ and $\leftrightarrow$) is assumed. The countable alphabet of propositional variables $p, q, r \ldots$ is denoted by Σ, propositional sentences by $\tau, \phi, \psi, \chi, \omega, \ldots$, constant truth by $\top$, its negation by $\perp$, and countable sets of propositional sentences by $K, L, M, \ldots$ and $A, B, C, \ldots$

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

$$Cn(K) \stackrel{\text{def}}{=} \{\phi \in \mathcal{L} \mid K \vdash \phi\}. \tag{1}$$

Instead of $Cn(\{\phi\})$, we will also write $Cn(\phi)$. Deductively closed sets of propositional sentences, i.e., $K = Cn(K)$, are denoted by $K, L, M, \ldots$ and are called *belief sets*. Arbitrary sets of sentences are called *belief bases* and are denoted by the letters A, B, C. Systems of belief bases and belief sets are denoted by S. Finite belief bases C are often identified with the conjunction of all propositions $\bigwedge C$. If $S = \{A_1, \ldots, A_n\}$ is a finite family of finite belief bases, then $\bigvee S$ shall denote a proposition logically equivalent to $(\bigwedge A_1) \vee \ldots \vee (\bigwedge A_n)$. As usual, we set $\bigvee \emptyset = \perp$.

Sometimes, we will also talk about truth assignments and models of propositions and belief bases. A *truth assignment* is a function $\mathcal{I}: \Sigma \rightarrow \{\mathsf{T}, \mathsf{F}\}$. A *model* $\mathcal{I}$ of a proposition ϕ is a truth assignment that satisfies ϕ in the classical sense, written $\models_{\mathcal{I}} \phi$. A *model of a belief base* C is a truth assignment that satisfies all propositions in C, written $\models_{\mathcal{I}} C$.

As usual (see (Gärdenfors, this book)), $K + \phi$ is the *expansion* of K by ϕ, $K \dotplus \phi$ is the *revision* of the belief set K by ϕ, and $K \dot{-} \phi$ is the *contraction* of K by ϕ. $(K \perp \phi)$ denotes the system of set-inclusion maximal subsets of K that do not imply ϕ, and γ denotes a *selection function* that selects a subset of $(K \perp \phi)$.

3 SYNTAX-BASED REVISION APPROACHES: BASE REVISIONS

The logic of theory change 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, Ullman, Vardi 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, Mendelzon 1989; Katsuno, 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. In this paper, we adopt the former perspective. In order to distinguish operations on syntactic descriptions – on belief bases – from operations on belief sets, belief base changes are called *base revision* and *base contraction*.

The idea of *changing a belief base 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 $(C \perp \phi)$ as the set of maximal subsets of C not implying ϕ, *simple base revision*, written as $C \oplus \phi$, could be defined as follows (Fagin, Ullman, Vardi 1983; Ginsberg 1986; Nebel 1989):

$$C \oplus \phi \stackrel{\text{def}}{=} \Big(\bigcap_{B \in (C \perp \neg\phi)} Cn(B) \Big) + \phi. \tag{2}$$

The operation $\oplus$ considers all sentences in a base as equally relevant. In most applications, however, we want to distinguish between the importance or relevance of different sentences. In (Fagin, Ullman, Vardi 1983) database priorities are assigned to propositions in order to reflect the distinction between facts and integrity rules. Ginsberg (1986) and Ginsberg and Smith (1987) make a distinction between facts that can change and those that are "protected."[2]

This idea of assigning different priorities to sentences can be formalized by employing a *complete preorder* with maximal elements, written $\phi \preceq \psi$, on the elements of a belief base C. In other words, we consider a *reflexive* and *transitive* relation such that for all $\phi, \psi \in C$ we have $\phi \preceq \psi$ or $\psi \preceq \phi$. For $\phi \preceq \psi$ and $\psi \not\preceq \phi$, we will also write $\phi \prec \psi$. Further, there exists at least one maximal element ϕ, i.e., for no element ψ: $\phi \prec \psi$. This relation will be called *epistemic relevance ordering*. It induces an equivalence relation, written $\phi \simeq \psi$, as follows:

$$\phi \simeq \psi \quad \text{iff} \quad (\phi \preceq \psi \text{ and } \psi \preceq \phi). \tag{3}$$

The corresponding equivalence classes are denoted by $\overline{\chi}$ and are called *degrees of epistemic relevance* of C. The set of equivalence classes $C/{\simeq}$ is denoted by $\overline{C}$. Since the preorder is complete, $\preceq$ is a linear order on $\overline{C}$. Further, there exists a maximal such degree because the preorder contains maximal elements.

A belief base together with an epistemic relevance ordering will be called *prioritized base*. If the belief base is finite, we will also use the notation $C_1, \ldots, C_n$ to denote the n degrees of epistemic relevance of C with the convention that C_1 has highest relevance.

Employing an epistemic relevance ordering, the *prioritized removal* of ϕ from C, written $C \Downarrow \phi$, will be defined as a system S of subsets of C. Each element $B \in S$ in turn is the union over a family consisting of subsets of all degrees of epistemic relevance, i.e.,

$$B \quad = \quad \bigcup \{B_{\overline{\chi}}\}_{\overline{\chi} \in \overline{C}} \text{ where } B_{\overline{\chi}} \subseteq \overline{\chi}. \tag{4}$$

Formally, $B \in (C \Downarrow \phi)$ if, and only if,

1. $B = \bigcup_{\overline{\chi} \in \overline{C}} B_{\overline{\chi}}$,
2. for all $\overline{\chi} \in \overline{C}$, $B_{\overline{\chi}} \subseteq \overline{\chi}$, and

[2] In particular, (Ginsberg, Smith 1987) makes clear, however, that usually more than one level of protected sentences is needed. For instance, the rule that an object can only occupy one place is, of course, an undeniable truth in our commonsense view of the world, while the rule that a room becomes stuffy when the ventilation is blocked may well be violated by an open window.

3. for all $\overline{\chi} \in \overline{C}$, $B_{\overline{\chi}}$ is set-inclusion maximal among the subsets of $\overline{\chi}$ such that $\bigcup_{\overline{\psi} \succeq \overline{\chi}} B_{\overline{\psi}} \not\vdash \phi$.

Intuitively, the elements of $C \Downarrow \phi$ are constructed by selecting a maximal subset not implying ϕ from the greatest degree of epistemic relevance, then a maximal subset of the next important degree is added such that ϕ is not implied, and so on. Note, however, that this intuition about *constructing* the elements of $C \Downarrow \phi$ may fail in the general case. Since we did not place restrictions on the relevance ordering, it can happen that there are infinitely ascending chains of degrees of epistemic relevance. Nevertheless, also in this case the existence of elements of B's satisfying the above conditions is guaranteed by Zorn's lemma.

A prioritized removal operation selects by definition a subset of the maximal subsets of a base not implying a given proposition.

Proposition 1 *Given a base C and a relevance ordering $\preceq$, for all ϕ:*

$$(C \Downarrow \phi) \quad \subseteq \quad (C \perp \phi). \tag{5}$$

Thus, it makes sense to use $\Downarrow$ instead of $\perp$ in the definition (2). The resulting operation is called *prioritized base revision*, denoted by $\hat{\oplus}$. This operation is identical to simple base revision in case that there is only one degree of epistemic relevance.

In the interesting special case when we are dealing with finite belief bases—which corresponds to prioritized logical databases investigated in (Fagin, Ullman, Vardi 1983)—the result of a prioritized base revision can be finitely represented.

Proposition 2 *If C is a finite belief base then*

$$C \hat{\oplus} \phi \quad = \quad Cn\Big(\big(\bigvee(C \Downarrow \neg\phi)\big) \wedge \phi\Big), \tag{6}$$

for every prioritized base revision $\hat{\oplus}$ on C.

Proof. Since C is finite, there can be only a finite number of finite degrees of epistemic relevance, hence, $C \Downarrow \neg\phi$ is a finite set of finite belief bases. In this case, the following equivalences hold

$$Cn(\bigwedge_{i=1}^{n} \phi_i) \quad = \quad Cn(\{\phi_1, \ldots, \phi_n\}) \tag{7}$$

$$Cn(\bigvee_{i=1}^{n} \phi_i) \quad = \quad \bigcap_{i=1}^{n} Cn(\phi_i), \tag{8}$$

and the proposition follows immediately. ■

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 C:

$$\begin{aligned} C_1 &= \{((b \wedge s) \rightarrow t)\}, \\ C_2 &= \{s\}, \\ C_3 &= \{b\}, \\ C &= C_1 \cup C_2 \cup C_3. \end{aligned}$$

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

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

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 A and B that have the same meaning, i.e., $Cn(A) = Cn(B)$, base revision can lead to different results, i.e., $Cn(A \hat{\oplus} \phi) \neq Cn(B \hat{\oplus} \phi)$. Base revision has a "morbid sensitivity to the syntax of the description of the world" (Winslett 1988), which is considered as an undesirable property. Dalal (1988) formulated the *principle of irrelevance of syntax* which states that a revision operation 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, Mendelzon 1989)), i.e., revision operations shall operate on the *knowledge level* (Newell 1982). In the theory of epistemic change this is accomplished by the requirements that the objects to be revised are belief sets and that the result of a revision does not depend on the syntactical form of the sentence to be added (AGM postulate (K $\dot{+}$6)).

Obviously, base revision does not satisfy the principle of irrelevance of syntax—and is not a belief 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 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, 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 complete preorders 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. This is actually one of the basic messages of the theory of epistemic change. We have to make up our minds about the importance of propositions or sets of propositions in order to select among the alternatives which are logically possible. If we consider all of them as equally important and combine them (by using full meet revision), we end up with nothing. Similarly, in case of a model-theoretic perspective, we cannot consider all models as equally possible candidates for a revision, since this would lead to a similar result.

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*.

4 BELIEF REVISION GENERATED BY EPISTEMIC RELEVANCE

The intention behind base revision is that all the sentences in a belief base A are considered as *relevant*—some perhaps more so than others. For this reason we want to give up as few sentences from A as possible, while with sentences that are only derivable we are more liberal. Formalizing this idea we employ as in the case of belief bases an *epistemic relevance ordering*, i.e. a complete pre-order with maximal elements on the entire *belief set*, with the intention of assigning the least degree of relevance to sentences that are only derivable. Based on these orderings, selection functions are constructed that select subsets that are maximally preferred with respect to epistemic relevance orderings.

We start by defining a *strict partial ordering* expressing preferences on subsets $A, B \in 2^K$, written as $A \ll B$, by

$$A \ll B \quad \text{iff} \quad \exists \overline{\tau}\colon \Big((A \cap \overline{\tau} \subset B \cap \overline{\tau}) \text{ and } \forall \overline{\omega} \succ \overline{\tau}\colon (A \cap \overline{\omega} = B \cap \overline{\omega}) \Big), \tag{9}$$

which in turn can be used to define a function $\gamma_{\preceq}$ that selects all *maximally preferred* elements of $K \perp \phi$:

$$\gamma_{\preceq}(K \perp \phi) \quad \stackrel{\text{def}}{=} \quad \{ L \in (K \perp \phi) \,|\, \forall M \in (K \perp \phi)\colon L \not\ll M \}. \tag{10}$$

Note that such maximally preferred sets always exist as can be easily inferred from the following lemma that relates maximally preferred sets to the elements of a prioritized removal.

Lemma 3 *Let K be a belief set with an epistemic relevance ordering $\preceq$. Then for any sentence ϕ:*

$$L \text{ is maximally preferred in } (K \perp \phi) \quad \text{iff} \quad L \in (K \Downarrow \phi). \tag{11}$$

Proof. Note that Proposition 1 applies also to belief sets because any belief set is also a belief base by definition. Hence, for all K and epistemic relevance orderings on K, for all ϕ:

$$(K \Downarrow \phi) \quad \subseteq \quad (K \perp \phi). \tag{12}$$

Assume that $B \in (K \Downarrow \phi)$. Assume for contradiction that there is a set $M \in (K \perp \phi)$ such that $B \ll M$. This means there exists a degree $\overline{\tau}$ such that $B_{\overline{\tau}} \subset M \cap \overline{\tau}$ while for all $\overline{\omega} \succ \overline{\tau}$ we have $B_{\overline{\omega}} = M \cap \overline{\omega}$. However, the set $B_{\overline{\tau}}$ is by definition of $\Downarrow$ a set-inclusion maximal subset of $\overline{\tau}$ such that $(\bigcup_{\overline{\omega} \succ \overline{\tau}} B_{\overline{\omega}}) \cup B_{\overline{\tau}}$ does not imply ϕ, hence, $M \cap \overline{\tau}$ cannot be a proper superset of $B_{\overline{\tau}}$.

For the other direction, assume L is maximal w.r.t. $\ll$ in $(K \perp \phi)$. Set $B = L$ and $B_{\overline{\chi}} = L \cap \overline{\chi}$. Obviously, the following conditions are satisfied:

1. $L = B = \bigcup_{\overline{\chi} \in \overline{K}} B_{\overline{\chi}}$,
2. $L \cap \overline{\chi} = B_{\overline{\chi}} \subseteq \overline{\chi}$, and

3. $L \cap \overline{\chi} = B_{\overline{\chi}}$ is set-inclusion maximal among the subsets of $\overline{\chi}$ such that $\bigcup_{\overline{\psi} \succeq \overline{\chi}} (L \cap \overline{\psi}) = \bigcup_{\overline{\psi} \succeq \overline{\chi}} B_{\overline{\psi}} \not\vdash \phi$.

Hence, $L \in (K \Downarrow \phi)$. ∎

This means that $\gamma_{\preceq}$ selects a nonempty subset of $(K \bot \phi)$ provided $(K \bot \phi)$ is nonempty, i.e., $\gamma_{\preceq}$ is a *selection function* (see (Gärdenfors, this book)) that may be used to define a partial meet revision operation. Revisions defined in this way will be called *revisions generated by epistemic relevance.* Analyzing the properties of such revisions, we note that they satisfy most of the AGM postulates.

Theorem 4 *Revisions generated by epistemic relevance satisfy* (K $\dot{+}$*1*)–(K $\dot{+}$*7*)*.*

Proof. Since $\gamma_{\preceq}$ is a selection function, revisions generated by epistemic relevance satisfy (K $\dot{+}$1)–(K $\dot{+}$6) by (Alchourrón, Gärdenfors, Makinson 1985, Observation 2.3).

Further, we have by definition of the selection function that there exists a relation $\sqsubseteq$, defined by putting

$$M \sqsubseteq L \quad \text{iff} \quad L \not\ll M, \tag{13}$$

such that for all ϕ,

$$\gamma(K \bot \phi) \quad = \quad \{L \in (K \bot \phi) |\, \forall M \in (K \bot \phi)\colon M \sqsubseteq L\} \tag{14}$$

Thus, $\dot{+}$ is a *relational* partial meet revision, which by (Alchourrón, Gärdenfors, Makinson 1985, Observations 3.1, 4.2 and 4.3) satisfies (K $\dot{+}$7). ∎

Note that the relation $\not\ll$ is not transitive and therefore revisions generated by epistemic relevance do not satisfy (K $\dot{+}$8) in general.[3] The interesting point about such revisions is that they coincide with prioritized base revision as defined in Section 3. That any revision generated by epistemic relevance can be conceived as a prioritized base revision follows already from Lemma 3. In order to show the other direction of the correspondence, the following Lemma (adapted from (Nebel 1989)) is helpful.

Lemma 5 *Let K be a belief set and ϕ be sentence such that $\neg\phi \in K$. Let C be any subset of K such that $C \not\vdash \neg\phi$. Then*

$$\left(\bigcap\{M \in K \bot \neg\phi |\, C \subseteq M\}\right) + \phi \quad = \quad Cn(C) + \phi. \tag{15}$$

[3]For a counter-example consult Section 6.

Proof. "$\supseteq$": Since by the assumption of the lemma $C \subseteq K$ and $C \not\vdash \neg\phi$, $(\bigcap\{M \in K\bot\neg\phi | \; C \subseteq M\})$ contains C as a subset by definition. Further, since all elements of $(K\bot\phi)$ are belief sets and the intersection of belief sets is a belief set again, $(\bigcap\{M \in K\bot\neg\phi | \; C \subseteq M\})$ contains $Cn(C)$, hence, the right hand side is a subset of the left hand side.

"$\subseteq$": Assume the contrary, i.e., there is a sentence ψ that is an element of the left hand side of equation (15), but $\psi \notin Cn(C \cup \{\phi\})$. By the latter assumption $Cn(C \cup \{\neg\psi\} \cup \{\phi\})$ is consistent and $\neg\phi \notin Cn(C \cup \{\neg\psi\}) \supseteq Cn(C \cup \{\neg\psi \vee \neg\phi\})$. By the assumption of the lemma that $\neg\phi \in K$, we have $(\neg\phi \vee \neg\psi) \in K$. Since also $C \subseteq K$, there is at least one element in $(K\bot\neg\phi)$ that contains $C \cup \{\neg\psi \vee \neg\phi\}$. Call this set L.

From the first assumption that $\psi \in Cn((\bigcap\{M \in K\bot\neg\phi | \; C \subseteq M\}) \cup \{\phi\})$, we conclude $(\phi \rightarrow \psi) \in Cn(\bigcap\{M \in K\bot\neg\phi | \; C \subseteq M\})$. However, the set $L \in (K\bot\neg\phi)$ that contains C and $(\neg\psi \vee \neg\phi)$ cannot contain $(\phi \rightarrow \psi)$ because otherwise $L \vdash \neg\phi$.

By the fact that the intersection over a system of belief sets is already a belief set, we have $Cn(\bigcap\{M \in K\bot\neg\phi | \; C \subseteq M\}) = \bigcap\{M \in K\bot\neg\phi | \; C \subseteq M\}$.

Finally, because $L \in \{M \in K\bot\neg\phi | \; C \subseteq M\}$, it cannot be the case that $(\phi \rightarrow \psi) \in (\bigcap\{M \in K\bot\neg\phi | C \subseteq M\})$. Thus, we have a contradiction of our assumption. Hence, the left hand side must be a subset of the right hand side. ■

It should be noted that if in the above lemma the set C is empty, the lemma describes the behavior of full meet revision. Another way to look at this lemma is that if the selection function selects elements of $K\bot\neg\phi$ by *focusing* on a particular set C, then the result of the revision is the set of consequences of the union of C and the new sentence. This result can be easily generalized to systems of *focusing sets*.

Lemma 6 *Let K be a belief set, and let ϕ be a sentence such that $\neg\phi \in K$. Let S be a system of subsets of K, where $C \not\vdash \neg\phi$ for all $C \in S$. Then*

$$\Big(\bigcap\{M \in K\bot\neg\phi | \; \exists C \in S \colon C \subseteq M\}\Big) + \phi \quad = \quad \Big(\bigcap_{C \in S} Cn(C)\Big) + \phi. \tag{16}$$

Proof.

$$\begin{aligned} &\Big(\bigcap\{M \in K\bot\neg\phi | \; \exists C \in S \colon C \subseteq M\}\Big) + \phi = \\ &= \; Cn\Big((\bigcap\{M \in K\bot\neg\phi | \; \exists C \in S \colon C \subseteq M\}) \cup \{\phi\}\Big) \end{aligned} \tag{17}$$

$$= Cn\Big(\big(\bigcap(\bigcup_{C\in S}\{M \in K\bot\neg\phi |\, C \subseteq M\})\big) \cup \{\phi\}\Big) \quad (18)$$

$$= Cn\Big(\big(\bigcap_{C\in S}(\bigcap\{M \in K\bot\neg\phi |\, C \subseteq M\})\big) \cup \{\phi\}\Big) \quad (19)$$

$$= Cn\Big(\bigcap_{C\in S}\big((\bigcap\{M \in K\bot\neg\phi |\, C \subseteq M\}) \cup \{\phi\}\big)\Big) \quad (20)$$

$$= Cn\Big(\bigcap_{C\in S} Cn\big((\bigcap\{M \in K\bot\neg\phi |\, C \subseteq M\}) \cup \{\phi\}\big)\Big) \quad (21)$$

$$= Cn\Big(\bigcap_{C\in S} Cn\big(C \cup \{\phi\}\big)\Big) \quad (22)$$

$$= Cn\Big(\bigcap_{C\in S} Cn\big(Cn(C) \cup \{\phi\}\big)\Big) \quad (23)$$

$$= Cn\Big(\big(\bigcap_{C\in S} Cn(C)\big) \cup \{\phi\}\Big) \quad (24)$$

$$= \big(\bigcap_{C\in S} Cn(C)\big) + \phi. \quad (25)$$

Equation 17 is the application of the definition of the expansion of a belief set. (18)–(20) follow by set theory. (21) follows because for any system of belief sets S the following equation holds:

$$Cn(\bigcap_{K\in S}(K \cup \{\phi\})) \;=\; Cn(\bigcap_{K\in S} Cn(K \cup \{\phi\})) \quad (26)$$

The "$\subseteq$" direction is obvious. For the other direction assume a sentence ψ that is an element of the right hand side, i.e., such that for all $K \in S$ we have $\psi \in Cn(K \cup \{\phi\})$. By the deduction theorem, $(\phi \to \psi) \in Cn(K)$ for all $K \in S$. Since $K = Cn(K)$, it holds that $(\phi \to \psi) \in (\bigcap_{K\in S} K)$, hence, $\psi \in Cn((\bigcap_{K\in S} K) \cup \{\phi\}) = Cn(\bigcap_{K\in S}(K \cap \{\phi\}))$.

(22) is an application of Lemma 5. (23) follows from properties of Cn, (24) is another application of equation (26), and, finally, (25) is another application of the definition of the expansion of a belief set. ■

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

Theorem 7 *For any revision operation $\dot{+}$ on a belief set K generated by epistemic relevance, there exists a corresponding prioritized base revision $\hat{\oplus}$ on some base C of*

K, and vice versa, such that for all ϕ:

$$K \dotplus \phi \;=\; C \hat{\oplus}\, \phi. \tag{27}$$

Proof. Assume a belief set K and an epistemic relevance ordering $\preceq$ on K. By definition, any belief set is also a belief base. Applying Lemma 3, it follows that

$$\gamma_{\preceq}(K \perp \neg\phi) = (K \Downarrow \neg\phi). \tag{28}$$

Hence, for a given revision on K generated by epistemic relevance, there is a prioritized base revision on some base C of K (namely, the base $C = K$)[4] such that for all ϕ:

$$K \dotplus \phi = C \hat{\oplus}\, \phi. \tag{29}$$

For the other direction, assume a prioritized belief base C with degrees of epistemic relevance $\overline{C}$. Set $K = Cn(C)$ and set $\overline{K} = \overline{C} \cup \{\overline{0}\}$, where $\overline{0} = K - C$ and $\overline{0} \prec \overline{\chi}$ for all $\overline{\chi} \in \overline{C}$. Now we will show that

$$\gamma_{\preceq}(K \perp \neg\phi) = \{M \in (K \perp \neg\phi) \,|\, \exists A \in (C \Downarrow \neg\phi)\colon\ A \subseteq M\}. \tag{30}$$

"$\supseteq$": Let $A \in (C \Downarrow \neg\phi)$ and let $L \in (K \perp \neg\phi)$ such that $A \subseteq L$. Such a set L exists because $A \subseteq C \subseteq K$ and $A \not\vdash \neg\phi$. Then L must be maximal w.r.t. $\ll$ in $(K \perp \neg\phi)$. Assuming otherwise would mean that there is a degree $\overline{\chi}$ and the selected subset $A_{\overline{\chi}} \subseteq \overline{\chi}$ was not maximal w.r.t. to the conditions in the definition of the elements of a prioritized removal, or there is another set $N \in (K \perp \neg\phi)$ that is identical to L for all priority $\overline{\chi} \in \overline{C}$ but contains a larger subset of $\overline{0}$, which is impossible, however, because L is already a maximal subset of K.

"$\subseteq$": Assume that L is an element of the left hand side of equation (30), i.e., L is a maximal element w.r.t. $\ll$. Consider the set $A = L \cap C$. Assume for contradiction that $A \notin (C \Downarrow \neg\phi)$. Since $A \not\vdash \neg\phi$, this means that there is set $B \in (C \Downarrow \neg\phi)$ such that $\bigcup_{\overline{\psi} \succeq \overline{\chi}} (A \cap \overline{\psi}) \subset \bigcup_{\overline{\psi} \succeq \overline{\chi}} B_{\overline{\psi}}$ for some degree $\overline{\chi} \in \overline{C}$. Now, since $B \not\vdash \neg\phi$ and $B \subseteq K$, there must be a set $M \in (K \perp \neg\phi)$ that contains B. By definition of $\ll$, we would then have $L \ll M$. Hence, L cannot be a maximal element w.r.t. $\ll$ and we have a contradiction. Thus, the left hand side is a subset of the right hand side.

Applying Lemma 6 to equation (30) we get

$$\Big(\bigcap \gamma_{\preceq}(K \perp \neg\phi)\Big) + \phi \;=\; \Big(\bigcap_{B \in (C \Downarrow \neg\phi)} Cn(B)\Big) + \phi, \tag{31}$$

[4]Note, however, that a smaller base would be sufficient as can be seen from the proof of Proposition 8.

i.e., for any prioritized base revision on C there exists an equivalent revision on $Cn(C)$ generated by epistemic relevance. ■

This means that prioritized base revision coincides with revision generated by epistemic relevance in the sense that the class of prioritized base revisions is identical with the class of revisions generated by 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 generated by 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 correspondence between belief revision generated by epistemic relevance and prioritized base revision can be exploited computationally cannot be answered positively in the general case. Although Theorem 7 states that it is possible to compute a revision on a belief set K generated by epistemic entrenchment by performing a prioritized base revision on some base of K, this does not help very much because in the proof we used K itself as the base. For the case of belief sets that are finite modulo logical equivalence, however, a revision operation generated by epistemic relevance can be performed by a prioritized base revision on a finite base.

Proposition 8 *Let K be a belief set finite modulo logical equivalence. If $\dot{+}$ is a revision on K generated by epistemic relevance, then there exists a finite prioritized base C, such that for all ϕ:*

$$K \dot{+} \phi \quad = \quad C \hat{\oplus} \phi. \tag{32}$$

Proof. Define C such that it contains one representative χ for each class of logically equivalent sentences $[\chi] = \{\phi \in K | \vdash \chi \leftrightarrow \phi\}$. These representatives are chosen to be maximal elements w.r.t. $\preceq$ in $[\chi]$. The relevance ordering on C is defined as the restriction of the epistemic relevance ordering on K.

Since K is finite modulo logical equivalence, C is finite. In order to show that (32) holds, it obviously suffices to prove the following condition:

$$A \in (C \Downarrow \neg\phi) \quad \text{iff} \quad Cn(A) \in \gamma_{\preceq}(K \bot \neg\phi). \tag{33}$$

"$\Rightarrow$": Assume $A \in (C \Downarrow \neg\phi)$. First, we verify that $Cn(A) \in (K \perp \neg\phi)$. By definition of $\Downarrow$ A does not imply $\neg\phi$. Furthermore, $Cn(A)$ is a maximal subset of K. Assuming otherwise, i.e., $Cn(A) \cup \psi \not\vdash \neg\phi$ for some $\psi \in K$, would mean that there is a sentence $\chi \in [\psi]$ such that $A \cup \chi \not\vdash \neg\phi$, which is impossible by the construction of C and the definition of $\Downarrow$.

Second, $Cn(A)$ must be maximal w.r.t. $\ll$ in $(K \perp \neg\phi)$. Let us assume the contrary, i.e., there is a set $L \in \gamma_{\preceq}(K \perp \neg\phi)$ and $Cn(A) \ll L$. This means for some degree $\overline{\tau}$: $Cn(A) \cap \overline{\tau} \subset L \cap \overline{\tau}$ while for all larger degrees the sets are identical. Chose a proposition $\psi \in (L \cap \overline{\tau}) - (Cn(A) \cap \overline{\tau})$. Let $\chi \in [\psi]$ be maximal w.r.t. $\preceq$. Note that $\chi \notin A$ and that $\tau \simeq \psi \preceq \chi$. By this we conclude that $L \supseteq (\bigcup_{\overline{\tau} \succeq \overline{\chi}} A_{\overline{\tau}}) \cup \{\chi\} \vdash \neg\phi$. This means however, that there cannot be a set L that is larger than $Cn(A)$ w.r.t. $\ll$.

"$\Leftarrow$": Assume a set $L \in (K \perp \neg\phi)$ that is maximally preferred. Set $A = L \cap C$. Because of the construction of C, we have $Cn(A) = L$. Assume for contradiction that $A \notin (C \Downarrow \neg\phi)$. This means for some degree $\overline{\tau}$ there is a sentence $\chi \in \overline{\tau}$ such that $\chi \notin A$ but $(\bigcup_{\overline{\omega} \succeq \overline{\tau}} A_{\overline{\omega}}) \cup \{\chi\} \not\vdash \neg\phi$. However, in this case there is also a set $M \in (K \perp \neg\phi)$ that contains $(\bigcup_{\overline{\omega} \succeq \overline{\tau}} A_{\overline{\omega}}) \cup \{\chi\}$ and which is therefore more preferred than L. ■

5 EPISTEMIC RELEVANCE AND EPISTEMIC ENTRENCHMENT

Although revisions generated by epistemic relevance do not satisfy all AGM postulates, there are special cases that do so. A trivial special case is a revision generated by only one degree of epistemic relevance, which is equivalent to full meet revision. There are more interesting cases, however.

Gärdenfors and Makinson claim that the notion of *epistemic entrenchment* introduced in (Gärdenfors, Makinson 1988) is closely related to the notion of *database priorities* as proposed in (Fagin, Ullman, Vardi 1983). Since the notion of database priorities is the finite special case of *epistemic relevance orderings* on belief bases as introduced in Section 3, which can in turn be used to generate belief revision operations, one would expect that epistemic entrenchment is closely related to epistemic relevance. Although the intuitions are clearly similar, the question is whether the different formalizations lead indeed to identical results.

Epistemic entrenchment orderings, written as $\phi \preceq_\epsilon \psi$, are defined over the entire set of sentences $\mathcal{L}$ and have to satisfy the following properties (see also (Gärdenfors, this book)):

(EE1) If $\phi \preceq_\epsilon \psi$ and $\psi \preceq_\epsilon \chi$, then $\phi \preceq_\epsilon \chi$.

(EE2) If $\phi \vdash \psi$, then $\phi \preceq_\epsilon \psi$.

(EE3) For any ϕ, ψ, $\phi \preceq_\epsilon (\phi \wedge \psi)$ or $\psi \preceq_\epsilon (\phi \wedge \psi)$.

(EE4) When $K \neq Cn(\bot)$, then $\phi \notin K$ iff $\phi \preceq_\epsilon \psi$ for all $\psi \in \mathcal{L}$.

(EE5) If $\psi \preceq_\epsilon \phi$ for all $\psi \in \mathcal{L}$, then $\vdash \phi$.

Using such a relation, Gärdenfors and Makinson define *belief contraction generated by epistemic entrenchment*, written $K \stackrel{\epsilon}{-} \phi$, by

$$\psi \in K \stackrel{\epsilon}{-} \phi \text{ iff } \psi \in K \text{ and } ((\phi \vee \psi) \not\preceq_\epsilon \phi \text{ or } \vdash \phi) \tag{34}$$

and show that such a belief contraction operation satisfies all AGM postulates for contraction as well as the following condition (Gärdenfors, Makinson 1988, Theorem 4):

$$\phi \preceq_\epsilon \psi \quad \text{iff} \quad \phi \notin K \stackrel{\epsilon}{-} (\phi \wedge \psi) \text{ or } \vdash (\phi \wedge \psi). \tag{35}$$

Further, they show that any belief contraction operation satisfying all of the AGM postulates is generated by some epistemic entrenchment ordering (Gärdenfors, Makinson 1988, Theorem 5).

The question is now how to interpret these results in the framework of epistemic relevance orderings on belief sets. First of all, from (EE2), reflexivity follows. Second, from (EE2) and (EE3), it follows that either $\phi \preceq_\epsilon (\phi \wedge \psi) \preceq_\epsilon \psi$ or $\psi \preceq_\epsilon (\phi \wedge \psi) \preceq_\epsilon \phi$. This means, $\preceq_\epsilon$ is a complete preorder on $\mathcal{L}$. For the strict part of this ordering we will use the symbol $\prec_\epsilon$. Further, from (EE2) it follows that there are maximal elements, namely, all sentences logically equivalent to $\top$ (and perhaps some other sentences as well). Ignoring the minimal elements (the sentences that are not elements of the belief set (EE4)), 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. In this case, using interdefinability of revision and contraction, definition (34) coincides with a contraction operation that is defined by using the *Harper identity* and a revision operation generated by epistemic relevance.

Theorem 9 *Suppose a belief set K, an epistemic entrenchment ordering $\preceq_\epsilon$, and the contraction operation $\stackrel{\epsilon}{-}$ generated by $\preceq_\epsilon$. Let $\preceq$ be the epistemic relevance ordering that is the restriction of $\preceq_\epsilon$ to K, and let $\dot{+}$ be the revision generated by the epistemic relevance ordering $\preceq$. Then*

$$K \stackrel{\epsilon}{-} \phi \quad = \quad (K \dot{+} \neg\phi) \cap K. \tag{36}$$

Proof. For the limiting case $\vdash \phi$, we have $(K \dot{+} \neg\phi) = Cn(\bot)$, hence the right hand side equals K. By (34) we also get for the left hand side K.

For the case $\phi \notin K$, again $(K \cup \{\neg\phi\}) \cap K = K$. That the left hand side has the same value follows from (34) and the observation that by (EE4) $\phi \notin K$ and $\psi \in K$ implies $\phi \prec_\epsilon \psi$, which in turn implies by $\psi \vdash (\phi \vee \psi)$ and (EE2): $\phi \prec_\epsilon \psi \preceq_\epsilon (\phi \vee \psi)$.

For the principal case, $\phi \in K$ and $\not\vdash \phi$, we will show that

$$\psi \in \bigcap \gamma_{\preceq}(K \bot \phi) \quad \text{iff} \quad \psi \in K \text{ and } \phi \prec_\epsilon (\phi \vee \psi). \tag{37}$$

If this condition is satisfied, then equation (36) holds obviously for the principal case as well.

"$\Leftarrow$": Suppose $\psi \in K$ and $\phi \prec_\epsilon (\phi \vee \psi)$. Note that because of $((\phi \vee \psi) \wedge (\phi \vee \neg\psi)) \vdash \phi$ and (EE2) we have $((\phi \vee \psi) \wedge (\phi \vee \neg\psi)) \preceq_\epsilon \phi$, which leads by our assumption and (EE1) to $((\phi \vee \psi) \wedge (\phi \vee \neg\psi)) \prec_\epsilon (\phi \vee \psi)$. Because of (EE3), either $(\phi \vee \psi)$ or $(\phi \vee \neg\psi)$ is less entrenched than the conjunction of them. It cannot be the former since that is strictly more entrenched, hence

$$(\phi \vee \neg\psi) \preceq_\epsilon ((\phi \vee \psi) \wedge (\phi \vee \neg\psi)) \prec_\epsilon (\phi \vee \psi). \tag{38}$$

Consider an arbitrary set $L \in \gamma_{\preceq}(K \bot \phi)$. Assume that $(\phi \vee \psi) \in L$. Then $(\phi \vee \neg\psi) \notin L$, or equivalently $L \cup \{\psi\} \not\vdash \phi$. Since L is a maximal subset of K not implying ϕ, we have $\psi \in L$. Thus, assume $(\phi \vee \psi) \notin L$. Consider $M = L \cap \{\chi \in K | \, (\phi \vee \psi) \preceq_\epsilon \chi\}$. Because L is a maximally preferred subset in $(K \bot \phi)$, we must have $M \cup \{(\phi \vee \psi)\} \vdash \phi$, or, using the deduction theorem $M \vdash ((\phi \vee \psi) \rightarrow \phi)$, hence $M \vdash (\neg\psi \vee \phi)$. By the compactness of propositional logic, there is a finite subset $N \subseteq M$ such that $\bigwedge N \vdash (\neg\psi \vee \phi)$, hence, by (EE2) $\bigwedge N \preceq_\epsilon (\neg\psi \vee \phi)$, which by (EE3) implies that there is a sentence $\tau \in N$ such that $\tau \preceq_\epsilon (\neg\psi \vee \phi)$. By (38) we get $\tau \prec_\epsilon (\psi \vee \phi)$ which is in contradiction to the construction of M, however. Thus, ψ is a member of every maximally preferred set in $(K \bot \phi)$.

"$\Rightarrow$": Assume $\psi \in \bigcap \gamma_{\preceq}(K \bot \phi)$. Assume for contradiction that we nevertheless have $(\phi \vee \psi) \preceq_\epsilon \phi$. By the fact that $\phi \vdash (\phi \vee \neg\psi)$, we conclude

$$(\phi \vee \psi) \preceq_\epsilon \phi \preceq_\epsilon (\phi \vee \neg\psi). \tag{39}$$

Since $\psi \in \bigcap \gamma_{\preceq}(K \bot \phi)$, every set $L \in \gamma_{\preceq}(K \bot \phi)$ must contain ψ and, hence, $(\phi \vee \psi)$, i.e., $(\phi \vee \neg\psi) \notin L$. Consider the set $M = L \cap \{\chi \in K | (\phi \vee \neg\psi) \prec_\epsilon \chi\}$. Since no element of $\gamma_{\preceq}(K \bot \phi)$ contains $(\phi \vee \neg\psi)$, all such sets M must already contain propositions that together with $(\phi \vee \neg\psi)$ leads to the derivation of ϕ, i.e., $M \cup \{(\phi \vee \neg\psi)\} \vdash \phi$, or, by the deduction theorem $M \vdash ((\phi \vee \neg\psi) \rightarrow \phi)$, hence $M \vdash (\phi \vee \psi)$. By compactness, (EE2), and (EE3) we conclude that there exists a proposition $\tau \in M$ such that $\tau \preceq_\epsilon (\phi \vee \psi)$, and by the construction of M: $(\phi \vee \neg\psi) \prec_\epsilon (\phi \vee \psi)$, contradicting (39). ■

Thus, the notion of epistemic entrenchment can indeed be viewed as a special case of epistemic relevance orderings—and, in the finite case, as a special case of database priorities.

The next corollary makes explicit which of the conditions (EE1)–(EE5) are actually needed to lead to a fully rational revision operation generated by epistemic relevance.

Corollary 10 *Any revision generated by an epistemic relevance ordering $\preceq$ such that*

1. *if $\phi \vdash \psi$ then $\phi \preceq \psi$, and*
2. *for any ϕ, ψ: $\phi \preceq (\phi \wedge \psi)$ or $\psi \preceq (\phi \wedge \psi)$,*

satisfies all AGM postulates.

Proof. (EE1) is already entailed by the fact that $\preceq$ is a preorder. (EE4) concerns only elements that are not in the belief set, and are therefore not related by $\preceq$. Further, as can be seen from the proof of Theorem 9, (EE5) is not necessary at all. We can always add a maximal degree that contains all logically valid sentences and remove them from other degrees without changing the outcome of a revision. ■

Epistemic entrenchment orderings lead to "fully rational" contraction and revision, and, moreover all such belief change operations are generated by some epistemic entrenchment ordering. It is not obvious, however, how to arrive at such epistemic entrenchment orderings. While epistemic relevance can be easily derived from a given prioritized belief base, it is not clear whether there are natural ways to generate epistemic entrenchment orderings. In (Gärdenfors, Makinson 1988) it is proposed to start with a complete ordering over the maximal disjunctions derivable from a belief base. 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 Proposition 17 in Section 7), in order to change a belief set.

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

Proposition 11 *Let C be a prioritized belief base such that all degrees of epistemic relevance are singletons. Then $(C \Downarrow \phi)$ is a singleton iff $\not\vdash \phi$.*

Proof. Note that $(C \Downarrow \phi) \neq \emptyset$ if and only if $\not\vdash \phi$.

If $C \not\vdash \phi$ then trivially $(C \Downarrow \phi) = \{C\}$.

For the case $\not\vdash \phi$ and $C \vdash \phi$, assume for contradiction that $A, A' \in (C \Downarrow \phi)$ and $A \neq A'$. By the definition of $\Downarrow$ there must be some degree $\overline{\chi}$ such that $A_{\overline{\chi}} \neq A'_{\overline{\chi}}$. Let $\overline{\chi}$ be the greatest such class. Now, since the degrees of epistemic relevance are singletons, we either have $(\bigcup_{\overline{\tau} \succ \overline{\chi}} A_{\overline{\tau}} \cup \overline{\chi}) = (\bigcup_{\overline{\tau} \succ \overline{\chi}} A'_{\overline{\tau}} \cup \overline{\chi}) \vdash \phi$ or not. In both cases, A and A' would agree on whether they contain $\overline{\chi}$. Hence, they cannot be different. ■

Note that even when $(C \Downarrow \phi)$ is always a singleton (for $\not\vdash \phi$), the corresponding selection function $\gamma_{\preceq}$ does not necessarily select singletons from $(Cn(C) \bot \phi)$, i.e., the corresponding belief revision operation is not a *maxi-choice revision*.

Clearly, the epistemic relevance ordering on the belief set $Cn(C)$ cannot always be extended to an epistemic entrenchment ordering. Nevertheless, belief revisions corresponding to unambiguous prioritized base revisions satisfy all AGM postulates.

Theorem 12 *Let $\preceq$ be an unambiguous epistemic relevance ordering on a belief set K. Then the revision generated by this ordering satisfies all AGM postulates.*

Proof. By Theorem 4, the revision operation satisfies (K $\dot{+}$1)–(K $\dot{+}$7). Thus, we only have to verify (K $\dot{+}$8). By (Alchourrón, Gärdenfors, Makinson 1985, Corollary 4.5) it suffice to show that the revision operation is transitively relational, i.e., using definition (13), we have to show that $\not\ll$ is transitive.

Since $\preceq$ is an unambiguous epistemic relevance ordering on K, all degrees of epistemic relevance except for the least one are singletons. The least degree will be denoted by $\overline{0}$.

In order to show transitivity of $\not\ll$, we first show that incomparability of two sets $L, M \in (K \bot \neg\phi)$, written $L \| M$ and defined by

$$L \| M \quad \text{iff} \quad L \not\ll M \text{ and } M \not\ll L,$$

is an equivalence relation on $K \bot \neg\phi$. Symmetry and reflexivity of $\|$ are immediate consequences of the definition. For showing transitivity, suppose $L, M, N \in (K \bot \neg\phi)$ and $L \| M \| N$. If $L = M$ or $M = N$, then $L \| N$ follows immediately. Therefore assume $L \neq M \neq N$. If $L \| M$ and $L \neq M$, then there is a degree $\overline{\tau} \in \overline{K}$ such that

$$L \cap \overline{\tau} \not\subset \neq \not\supset M \cap \overline{\tau} \text{ and } \forall \overline{\omega} \succ \overline{\tau} \colon (L \cap \overline{\omega} = M \cap \overline{\omega}). \tag{40}$$

Since all degrees except $\overline{0}$ are singletons, it follows that $\overline{\tau} = \overline{0}$, i.e., $L \cap (K - \overline{0}) = M \cap (K - \overline{0})$. With the same argument, we conclude that $M \cap (K - \overline{0}) = N \cap (K - \overline{0})$, hence $L \cap (K - \overline{0}) = N \cap (K - \overline{0})$. Since L and N are maximal subsets of K, it cannot be the case that $L \cap \overline{0} \subset N \cap \overline{0}$ or $L \cap \overline{0} \supset N \cap \overline{0}$, hence $L \| N$.

From the fact that $\|$ is an equivalence relation, it follows straightforwardly that $L \| M$ and $M \ll N$ implies that $L \ll N$. For contradiction assume $L \not\ll N$. Then we must have $N \ll L$ because otherwise by transitivity of $\|$ we could conclude $M \| N$, which is a contradiction of the assumption. From $N \ll L$, the assumption that $M \ll N$ and the transitivity of $\ll$ it follows that $M \ll L$, which again contradicts the assumption. With the same argument, it follows that $L \ll M$ and $M \| N$ implies $L \ll N$.

Now assume $L \not\ll M \not\ll N$. By considering cases, transitivity of $\ll$ follows. (1) Assuming $L \| M$ and $M \| N$ leads to $L \| N$, hence, $L \not\ll N$. (2) Assuming $L \| M$ and $N \ll M$ leads to $N \ll L$, hence $L \not\ll N$. (3) Assuming $M \ll L$ and $M \| N$ leads to $N \ll L$, hence $L \not\ll N$. (4) Assuming $M \ll L$ and $N \ll M$ leads to $N \ll L$, hence $L \not\ll N$. Thus, belief revisions generated by unambiguous epistemic relevance are transitively relational and satisfy for this reason (K $\dot{+}$8). ■

Although an unambiguous relevance ordering is not necessarily an entrenchment ordering, it is possible to generate an epistemic entrenchment ordering using (35) that leads to an identical revision operation because unambiguous revisions are fully rational. Given an unambiguous prioritized base C, the epistemic entrenchment ordering can be derived as follows. For every pair of propositions $\phi, \psi \in Cn(C)$, determine $\{A\} = C \Downarrow \phi$ and $\{B\} = C \Downarrow \psi$, and set $\psi \prec_\epsilon \phi$ if and only if $A \ll B$. The verification that this is indeed the right epistemic entrenchment ordering is left as an exercise to the reader.

6 BELIEF REVISION AND DEFAULT REASONING

Doyle has 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.

Formally, the term *logical nonmonotonicity* refers to nonmonotonicity found in nonmonotonic logics, i.e., given a deductive closure operation $C(\cdot)$ of a nonmonotonic logic,

$$A \subseteq B \quad \not\Rightarrow \quad C(A) \subseteq C(B). \tag{41}$$

The notion of *temporal nonmonotonicity* refers to the development of a set of beliefs over time, where K_t will be used to refer to K at time point t:

$$t_1 \leq t_2 \quad \not\Rightarrow A_{t_1} \subseteq A_{t_2} \tag{42}$$

Although these two forms of nonmonotonicity should not be confused, sometimes they turn out to be intimately connected. In particular, the temporal nonmonotonicity induced by belief revision, i.e., the fact that in general we do not have $K \subseteq K \dot{+} \phi$, is related to logical nonmonotonicity induced by some forms of default reasoning. Further, there exists also a connection between a form of contraction and default reasoning, as we will see below.

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 $Cn(\phi) \subseteq Cn(\psi)$ does not imply $K \dot{+} \phi \subseteq K \dot{+} \psi$. Indeed, as is well known, requiring monotony in the second operand of a belief revision operation is impossible in the general case. Exploring the space of possible revision operations that imply monotony shows that the revision either violates one of the basic AGM postulates or it is a trivial revision on $Cn(\emptyset)$ or $Cn(\perp)$.

Proposition 13 *Let $\dot{+}$ be a belief revision operation defined on a belief set K. If for all ϕ, ψ*

$$K \dot{+} \phi \subseteq K \dot{+} \psi \quad \textit{if} \quad Cn(\phi) \subseteq Cn(\psi), \tag{43}$$

then

1. *The operation $\dot{+}$ violates one of the basic AGM postulates* (K $\dot{+}$*1*)*–*(K $\dot{+}$*6*)*, or*

2. $K = Cn(\emptyset)$ *and* $K \dot{+} \phi = Cn(\phi)$*, or*

3. $K = Cn(\bot)$ *and* $K \dot{+} \phi = Cn(\phi)$*.*

Proof. Assume $K \neq Cn(\bot)$ and a proposition ϕ with $\neg\phi \in K$ and $\not\vdash \neg\phi$. By (43) we would have $K \dot{+} \top \subseteq K \dot{+} \phi$. Because of (K $\dot{+}$3) and (K $\dot{+}$4), $K \dot{+} \top = K$. By assumption, we thus have $\neg\phi \in K \dot{+} \top$. Now, by (K $\dot{+}$2) $\phi \in K \dot{+} \phi$. Because of (K $\dot{+}$5) and the assumption $\not\vdash \neg\phi$, $\neg\phi \notin K \dot{+} \phi$. Thus, either the requirement $K \dot{+} \top \subseteq K \dot{+} \phi$ or one of the basic postulates is violated.

Let $\dot{+}$ a belief revision operation on $Cn(\emptyset)$ and assume that all basic postulates are satisfied. Then by (K $\dot{+}$1)–(K $\dot{+}$3) it follows that $Cn(\phi) \subseteq Cn(\emptyset) \dot{+} \phi \subseteq Cn(\emptyset \cup \{\phi\}) = Cn(\phi)$ and (43) is trivially satisfied.

Assume $K = Cn(\bot)$. If $\vdash \neg\phi$ then clearly $K \dot{+} \phi = Cn(\bot) = Cn(\phi)$ by (K $\dot{+}$1) and (K $\dot{+}$2). Thus, assume $\not\vdash \neg\phi$. By (K $\dot{+}$1) and (K $\dot{+}$2), we have $Cn(\phi) \subseteq K \dot{+} \phi$. Now assume there is a proposition $\chi \in K \dot{+} \phi$ such that $\chi \notin Cn(\phi)$. By (43) we would have $K \dot{+} \phi \subseteq K \dot{+} (\neg\chi \wedge \phi)$. However, this would violate (K $\dot{+}$2) or (K $\dot{+}$5). Thus, if the basic postulates and (43) are satisfied, $Cn(\bot) \dot{+} \phi = Cn(\phi)$. ■

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 a very general level. They compare various general conditions on nonmonotonic provability relations with the AGM postulates.

For the approaches to belief revision described in the previous section there is an even stronger connection to some models of nonmonotonic reasoning. Prioritized base revision, and hence partial meet revision generated by epistemic relevance, is expressively equivalent to *skeptical* provability[5] in Poole's (1988) *theory formation* approach and Brewka's (1989) *level default theories* (LDT)—in 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.

[5] A correspondence to *credulous* derivability could be achieved if a notion of *nondeterministic* revision as proposed in (Doyle 1990) is adopted.

An *extension* of Δ is a deductively closed set of propositions

$$E = Cn((\bigcup_{i=1}^{n} \mathcal{R}_i) \cup \mathcal{F}) \tag{44}$$

such that for all i with $1 \leq i \leq n$:

1. $\mathcal{R}_i \subseteq \mathcal{D}_i$,
2. $\mathcal{R}_i$ is set-inclusion maximal among the subsets of $\mathcal{D}_i$ such that $(\bigcup_{j=1}^{i} \mathcal{R}_j) \cup \mathcal{F}$ is consistent.[6]

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

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 $\perp$ while this is impossible in LDT's.

Theorem 14 *Let $\Delta = (\langle \mathcal{D}_1, \ldots, \mathcal{D}_n \rangle, \mathcal{F})$ be a* RDT*. Let $C = \bigcup_{i=1}^{n} \mathcal{D}_i$ be a prioritized base with degrees of epistemic relevance $\mathcal{D}_1, \ldots, \mathcal{D}_n$. Then for all ϕ:*

$$\Delta \mid\!\sim \phi \quad \textit{iff} \quad \phi \in (C \hat{\oplus} \mathcal{F}). \tag{45}$$

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

When $\mathcal{F}$ is consistent, then $(C \Downarrow \neg(\bigwedge \mathcal{F}))$ is by definition a system S of subsets $B \subseteq C$ such that

1. $B = \bigcup_{i=1}^{n} B_i$,
2. $B_i \subseteq C_i$, for all $1 \leq i \leq n$, and
3. for all $1 \leq i \leq n$, B_i is set-inclusion maximal among the subsets of C_i such that $\bigcup_{j=1}^{i} B_j \not\vdash \neg(\bigwedge \mathcal{F})$.

[6]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 $\perp$ can be derived iff $\mathcal{F}$ is inconsistent.

Since the second condition of 3. is equivalent with the condition that $(\bigcup_{j=1}^{i} B_j) \cup \mathcal{F})$ is consistent, it follows that by definition for every extension E of Δ there exists a set $B \in (C \Downarrow \neg(\bigwedge \mathcal{F}))$ such that $E = Cn(B \cup \mathcal{F})$ and vice versa, hence

$$\bigcap_{E \text{ is an extension of } \Delta} E = \bigcap_{B \in (C \Downarrow \neg(\bigwedge \mathcal{F}))} Cn(B) + (\bigwedge \mathcal{F}), \tag{46}$$

which completes the proof. ∎

This means that ranked default theories have the same expressive power as finitary prioritized base revision operations, which coincide with finitary belief revisions generated by epistemic relevance.

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 14 that may lead to $\perp \in Cn(C)$, 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 15 *Let Δ be a* RDT *as above. Then there exists a consistent prioritized base C and a proposition ψ such that for all ϕ*

$$\Delta \mid\!\sim \phi \quad \textit{iff} \quad \phi \in \left(C \hat{\oplus} (\psi \wedge \mathcal{F})\right). \tag{47}$$

Proof. Define C as in Theorem 14. Transform every sentence in C into negation normal form (i.e., into a formula that contains only $\wedge$, $\vee$ and $\neg$, and all negation signs appear only in front of propositional variables). Assuming without loss of generality that the alphabet of propositional variables Σ is finite, extend Σ to Σ' by adding for every propositional variable p a fresh variable p'. Now replace any negative literal $\neg p$ in all sentences of C by p', call the new belief base C' and define

$$y \stackrel{\text{def}}{=} \bigwedge_{p \in \Sigma} (\neg p \leftrightarrow p'). \tag{48}$$

Since no sentence in C' contains any negation sign, C' is consistent.

Let ϕ any proposition over Σ, we will show that for any two belief bases B and B', where B' is a transformed belief base according to the above rules, the following relation holds:

$$B \vdash \phi \quad \text{iff} \quad B' \cup \{\psi\} \vdash \phi. \tag{49}$$

Assume $B \vdash \phi$ but $B' \cup \{\psi\} \not\vdash \phi$. This means $B' \cup \{\psi\} \cup \neg\phi$ is satisfiable. Restricting the truth assignment of this belief set to Σ, we get one that must satisfy $B \cup \{\neg\phi\}$ by construction of B'. This is impossible, however. Conversely, assuming satisfiability of $B \cup \{\neg\phi\}$, a truth assignment can be extended to Σ' such that it satisfies $B \cup \{\psi\} \cup \neg\phi$, hence also $B' \cup \{\psi\} \cup \neg\phi$.

That means that for any maximal subset $B \subseteq C$ that is consistent with a given proposition ϕ there exists a corresponding set $B' \subseteq C'$, that is consistent with ψ and ϕ and maximal in C'. Further adding ψ to B' allows to derive the same propositions over Σ as can be derived from B. ∎

From the results above and the translation of (K $\dot{+}$8) to a condition on nonmonotonic derivability relations in (Makinson, Gärdenfors 1990), it follows that the derivability relation of RDT's w.r.t. the set of hard facts $\mathcal{F}$ does not satisfy *rational monotony* (see (Makinson, Gärdenfors 1990)).[7] This condition can be phrased as follows:

$$\text{If } \phi {\mid\!\sim} \psi \text{ and } \phi {\mid\!\not\sim} \neg\chi \text{ then } \phi \wedge \chi {\mid\!\sim} \psi \tag{50}$$

In plain words, if a proposition ϕ permits the plausible conclusion ψ, this conclusion continues to hold for the stronger premise $\phi \wedge \chi$ provided there is no plausible reason to deny χ given the assumption ϕ. Applying this condition to RDT's we consider the nonmonotonic derivability relation as parameterized by the defaults $\mathcal{D}$, written $\mathcal{F} {\mid\!\sim}_{\mathcal{D}} \phi$. For a counter-example to rational monotony, suppose a situation where two people of different sex meet the first time and try to get to know important facts about each other. Assume one person has the following background beliefs modeled as a set of defaults:

1. Being a parent implies being married ($p \rightarrow m$).
2. Living alone implies being a bachelor ($a \rightarrow b$).
3. Wearing a ring implies being a dandy or being married ($r \rightarrow (d \vee m)$).

[7]Note that this result depends on the exact correspondence between RDT's and belief revision generated by epistemic relevance. In (Makinson, 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.

All these defaults have the same priority. Further, suppose the postulate $(b \leftrightarrow \neg m)$ and the facts p, a, and r. One extension, which contains m, is the consequential closure of the facts and rules 1 and 3. The other possible extension, which contains d, is the closure of the facts and rule 2. This means that $(m \vee d)$ is a sceptical consequence:

$$(b \leftrightarrow \neg m) \wedge p \wedge a \wedge r \quad \mid\!\sim_{\mathcal{D}} \quad (m \vee d). \tag{51}$$

If $\neg d$ is added to the facts the expected conclusion m does not follow, however. In this case one extension, which contains m and $\neg d$, is generated by the facts and rule 1 and 3. The other possible extension is generated by the facts and rule 2 and contains $\neg m$ and $\neg d$. Hence,

$$\neg d \wedge ((b \leftrightarrow \neg m) \wedge p \wedge a \wedge r) \quad \not\mid\!\sim_{\mathcal{D}} \quad (m \vee d), \tag{52}$$

although

$$((b \leftrightarrow \neg m) \wedge p \wedge a \wedge r) \quad \not\mid\!\sim_{\mathcal{D}} \quad d. \tag{53}$$

Another interesting observation in this context is that the addition of *constraints* to RDT's is similar but not identical to a *belief contraction operation*. 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. $\mathcal{R}_i$ is set-inclusion maximal among the subsets of $\mathcal{D}_i$ such that $(\bigcup_{i=1}^{n} \mathcal{R}_i) \cup \mathcal{F} \cup \mathcal{C}$ is consistent.

It should be obvious that the addition of constraints is a generalization of the basic framework, i.e., for all $\mathcal{F}, \mathcal{D}, \phi$:

$$(\mathcal{D}, \mathcal{F}, \emptyset) \vdash \phi \quad \text{iff} \quad (\mathcal{D}, \mathcal{F}) \vdash \phi. \tag{54}$$

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)).

Theorem 16 *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* $C = \mathcal{F} \cup \bigcup_{i=1}^{n} \mathcal{D}_i$ *be a prioritized base with* $\mathcal{F}, \mathcal{D}_1, \ldots, \mathcal{D}_n$ *the degrees of relevance of* C. *Then*

$$\Delta \mid\!\sim \phi \quad \textit{iff} \quad \bigvee \Big(C \Downarrow \neg (\bigwedge \mathcal{C}) \Big) \vdash \phi \tag{55}$$

Proof. If $\mathcal{F} \cup \mathcal{C}$ is consistent, then every element $B \in (C \Downarrow \neg(\bigwedge \mathcal{C}))$ contains $\mathcal{F}$. Further, the subsets chosen from $\mathcal{D}_i$ are maximal subsets consistent with $\mathcal{F}$ and $\mathcal{C}$, hence, the extensions of Δ correspond to sets $B \in (C \Downarrow \neg(\bigwedge \mathcal{C}))$ and vice versa, such that $E = Cn(B)$. ■

This goes some way to answering the question whether there is a counter-part to contraction in nonmonotonic logics, raised in (Makinson, Gärdenfors 1990). Default reasoning with constraints in Poole's theory formation approach can be modeled by using base contraction.

Trying to lift this result to belief sets, however, is impossible in the general case. Usually, ranked default theories with constraints do not allow the derivation of $\bigwedge \mathcal{C}$, and this property is independent from consistency of the set of facts $\mathcal{F}$ with the set of defaults $\bigcup_i \mathcal{D}_i$. When contracting an inconsistent belief set, however, the contracted belief set contains the negation of the proposition used to contract the belief set. This property follows from the *Harper identity* when we set $A = Cn(\bot)$:

$$Cn(\bot) \dot{-} \chi = (Cn(\bot) \dot{+} \neg\chi) \cap Cn(\bot) = Cn(\bot) \dot{+} \neg\chi \supseteq Cn(\neg\chi) \qquad (56)$$

This means, provided we try to model derivability in such logics by belief contraction, in case when the defaults are inconsistent with the facts, a belief contraction would lead to the inclusion of the constraints—which may not be derivable in the corresponding default logic. Base contraction does not have this property because such operations remove more beliefs than belief contractions. In particular, while every contracted belief set $K \dot{-} \phi$ contains $Cn(K) \cap Cn(\neg\phi)$, a contracted base usually does not contain those beliefs (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 ψ in a belief set $K = Cn(C)$ revised by ϕ, i.e.,

$$\psi \in K \dot{+} \phi. \qquad (57)$$

As the input size we use the sum of the size $|C|$ of the belief base C that represents K and the sizes $|\phi|$ and $|\psi|$ of the sentences ϕ and ψ, respectively.

This assumption implies that the representation of the preference relation used to guide the revision process should be polynomially bounded by $|C|+|\phi|+|\psi|$. Although this sounds like a reasonable restriction, it is not met by all belief revision approaches. Belief revision generated by *epistemic entrenchment* orderings (Gärdenfors, 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, 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.

Proposition 17 *The set of maximal disjunctions implied by a belief base has a worst-case size that is exponential in the size of the belief base.*

A similar statement could be made about revisions generated by epistemic relevance. It is, of course, possible to have a belief base C that represents K and an epistemic relevance ordering over K that is not representable in a polynomial way w.r.t. $|C|$. However, if we consider only complete preorders over C with the understanding that the degree of least relevant sentences is $Cn(C) - C$, then the ordering is represented in a way that is polynomially bounded by $|C|$ and $\dot{+}$ can be computed by using the corresponding prioritized base revision. This means all "natural" epistemic relevance orderings are well-behaved.

Analyzing the computational complexity of the belief revision problems, the first thing one notes that deciding the trivial case $\psi \in Cn(\emptyset) \dot{+} \phi$ 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, 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

[8] We assume some familiarity with the basic notions of the theory of NP-completeness as presented in the first few chapters of (Garey, Johnson 1979). This means the terms *decision problem*, P, NP, co-NP, PSPACE, *polynomial transformation* (or *many-one reduction*), *polynomial Turing reduction*, *completeness* w.r.t. polynomial transformability or Turing reducibility should be familiar to the reader.

polynomial time on a deterministic 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]

$$
\begin{aligned}
\Delta_0^p &= \Sigma_0^p = \Pi_0^p = \mathsf{P}, && (58)\\
\Delta_{k+1}^p &= \mathsf{P}^{\Sigma_k^p}, && (59)\\
\Sigma_{k+1}^p &= \mathsf{NP}^{\Sigma_k^p}, && (60)\\
\Pi_{k+1}^p &= \mathsf{co}\text{-}\Sigma_{k+1}^p. && (61)
\end{aligned}
$$

Thus, $\Sigma_1^p = \mathsf{NP}$, $\Pi_1^p = \mathsf{co}\text{-}\mathsf{NP}$, 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:

$$\underbrace{\exists \vec{p}\, \forall \vec{q} \ldots}_{k \text{ alternating quantifiers starting with } \exists} F(\vec{p}, \vec{q}, \ldots). \tag{62}$$

The complementary problem, denoted by $\overline{k{-}\mathsf{QBF}}$, is complete for Π_k^p.

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 18 $\mathsf{FMR} \in \Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$ *provided* $\Sigma_1^p \neq \Pi_1^p$.

Proof. If $\dot{+}$ is a full meet revision, $\phi \in Cn(C) \dot{+} \psi$ can be solved by the following algorithm:

$$
\begin{array}{l}
\text{if } C \not\vdash \neg\phi \\
\quad \text{then } C \cup \{\phi\} \vdash \psi \\
\quad \text{else } \phi \vdash \psi
\end{array}
$$

From this, membership in Δ_2^p follows.

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

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

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 19 SBR *is* Π_2^p*-complete.*

Proof. We will prove that the complementary problem $C \oplus \phi \not\vdash \psi$, 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 $\vec{p} = p_1, \ldots, p_n$, let $\vec{q} = q_1, \ldots, q_m$, and let $\exists \vec{p} \forall \vec{q} F(\vec{p}, \vec{q})$ be an instance of 2-QBF. Set

$$C \quad = \quad \{p_1, \ldots, p_n, \neg p_1, \ldots, \neg p_n, \neg F(\vec{p}, \vec{q})\}. \tag{63}$$

Now we claim that

$$C \oplus \top \not\vdash \neg F(\vec{p}, \vec{q}) \quad \text{iff} \quad \exists \vec{p} \forall \vec{q} F(\vec{p}, \vec{q}) \ \text{ is true.} \tag{64}$$

$C \oplus \top \not\vdash \neg F(\vec{p}, \vec{q})$ if and only if there is an element $B \in (C \bot \top)$ such that $\neg F(\vec{p}, \vec{q}) \notin B$. Since every set of literals $\{l_1, \ldots l_n\}$ with $l_i = p_i$ or $l_i = \neg p_i$ is consistent, $\neg F(\vec{p}, \vec{q}) \notin B$ if and only if the set $\{l_1, \ldots, l_n\} \subseteq B$ is inconsistent with $\neg F(\vec{p}, \vec{q})$, i.e., $\{l_1, \ldots, l_n\} \vdash F(\vec{p}, \vec{q})$. This, in turn is equivalent with the fact that there is a truth assignment to $\vec{p}$ such that $F(\vec{p}, \vec{q})$ is true for all truth-assignments to $\vec{q}$. Thus, equivalence (64) holds.

Membership of $\overline{\mathsf{SBR}}$ in Σ_2^p follows from the following algorithm that needs nondeterministic polynomial time using an oracle for SAT:

1. Guess a set $B \subseteq C$.
2. Verify that there is no $\chi \in C - B$ such that $B \cup \{\chi\} \not\vdash \neg\phi$.
3. Verify $B \cup \{\phi\} \not\vdash \psi$.

∎

This means that SBR is, on one hand, not much more difficult than FMR, and, on the other hand, apparently easier than derivability in most modal logics (e.g., K, T, and

$S4$), which is a PSPACE-complete problem (Garey, 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. This means the worst-case behavior of any SBR algorithm is most probably not better or worse than the worst-case behavior of any propositional proof method. 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—for instance, by restricting the expressiveness—there would still be the problem of determining the maximal consistent 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 20 PBR *is* Π_2^p*-complete.*

Proof. Π_2^p-hardness is immediate by Theorem 19. Membership of PBR in Π_2^p also follows easily. The maximality test (step 2 in the algorithm used in the proof of Theorem 19) has to be performed as often as there are priority classes, which is polynomially bounded by $|Z|$.

■

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 19 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 21 UBR $\in \Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$*, provided* $\Sigma_1^p \neq \Pi_1^p$*.*

Proof. In order to show that UBR $\in \Delta_2^p$, we specify an algorithm to compute $C \hat{\oplus}\, \phi \vdash \psi$, for $\hat{\oplus}$ based on singleton degrees of relevance:

1. Initialize $A = \emptyset$ and $i = 1$.
2. Test $A \cup C_i \not\vdash \neg\phi$. If so, set $A = A \cup C_i$.
3. Increment i.
4. If there are only $i - 1$ degrees return with the result $(A \cup \{\phi\} \vdash \psi)$.
5. Otherwise continue with step 2.

Using an oracle for SAT, this algorithms runs in polynomial time. Thus, we have $\mathsf{UBR} \in \Delta_2^p$.

Using the same arguments as in the proof of Proposition 18 leads to $\mathsf{FMR} \notin \mathsf{NP} \cup \mathsf{co\text{-}NP}$ provided $\mathsf{NP} \neq \mathsf{co\text{-}NP}$. ■

From the proof, we can infer that if we can come up with a polynomial algorithm for satisfiability (by restricting the propositional language to Horn logic, 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. Recall that a model $\mathcal{I}$ of a belief base C is a truth assignment that satisfies all propositions in C. $mod(C)$ denotes the set of all models of C. $\delta(\mathcal{I}, \mathcal{J})$ denotes the number of propositional variables such that $\mathcal{I}$ and $\mathcal{J}$ map them to different truth-values. Assuming that $\mathcal{M}$ denotes a set of truth assignments, $g^m(\mathcal{M})$ 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 C is a finite belief base, then $G^m(C)$ is some belief base such that $mod(G^m(C)) = g^m(mod(C))$. Although G^m is not a deterministic function, all possible results are obviously logically equivalent.

[10] The RUP system 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].

Now, model-based revision, written $C \circ \phi$ is defined by:[11]

$$C \circ \phi \stackrel{\text{def}}{=} \begin{cases} G^m(C) \cup \{\phi\} & \text{for the least } m \text{ s.t.} \\ & G^m(C) \cup \{\phi\} \not\vdash \bot \\ \{\phi\} & \text{if } C \vdash \bot \text{ or } \phi \vdash \bot. \end{cases} \qquad (65)$$

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 22 $\mathsf{MBR} \in \Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$, *provided* $\Sigma_1^p \neq \Pi_1^p$.

Proof. Note that for any fixed i, $G^i(C) \not\vdash \phi$ 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. $\models_{\mathcal{I}} C$,
2. $\not\models_{\mathcal{J}} \phi$, and
3. $\delta(\mathcal{I}, \mathcal{J}) \leq i$.

Note further that solving $G^i(C) \cup \{\phi\} \not\vdash \psi$ can be reduced to solving $G^i(C) \not\vdash (\phi \rightarrow \psi)$.

Let n be the number of different propositional variables in C. Then it is obvious that $g^k(C) = g^{k+1}(C)$ for all $k \geq n$.

Membership of MBR in Δ_2^p follows from the following algorithm:

1. Determine the least i, where $0 \leq i \leq n$ such that $G^i(C) \not\vdash \neg\phi$.
   (a) If there is no such i, then return $\phi \vdash \psi$.
   (b) Otherwise, return $(G^i(C) \cup \{\phi\} \vdash \psi)$.

Since n is bounded polynomially by $|C|$, this algorithms runs in polynomial time using an oracle for the problem $G^i(C) \not\vdash \phi$.

[11]This definition is a slight extension of the definition given in (Dalal 1988) that takes also care of the limiting cases when C or ϕ is inconsistent.

$\mathsf{MBR} \notin \mathsf{NP} \cup \mathsf{co\text{-}NP}$ provided $\mathsf{NP} \neq \mathsf{co\text{-}NP}$ follows with the same argument as in the proof of Proposition 18. ∎

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 "fully rational" base revisions,[12] namely, FMR, UBR, and MBR (for the latter see (Dalal 1988)) turn out to be in this class, base revisions that are not "fully rational" cannot be shown to be in this class.

8 SUMMARY AND OUTLOOK

The class of *prioritized base revision* (a form of syntax-based approaches to belief revision) and the class of *belief revision operations generated by epistemic relevance* were shown to be identical, removing partially the restriction of the *theory of epistemic change* that states of beliefs have to be modeled as deductively closed sets of sentences.

Further, epistemic relevance orderings on belief sets were shown to be a generalization of *epistemic entrenchment orderings* confirming the intuition spelled out in (Gärdenfors, Makinson 1988) that epistemic entrenchment is related to the notion of *database priorities* as introduced in (Fagin, Ullman, Vardi 1983).

Complementing the results in (Makinson, Gärdenfors 1990), we showed that concrete models of nonmonotonic reasoning, namely, ranked default theories (RDT's)—a generalization of Poole's logic without constraints (Poole 1988) and Brewka's level default theories (Brewka 1989; Brewka 1990)—turn out to be expressively equivalent to prioritized base revision in the case of finitary propositional logic. In addition, some answer to the question raised in (Makinson, Gärdenfors 1990) whether contraction plays a role in nonmonotonic logics was given. The theory formation approach *with* constraints was shown to be equivalent—under some reasonable assumptions—to *base contraction*. It is not possible to lift this result to belief contraction, however.

Finally, the computational complexity of different base revision operations was investigated—where the results apply by the above mentioned correspondences to reasoning in default logics, as well.

The results confirm the intuition that unambiguous prioritized base revision is not harder but apparently less complex than general prioritized base revision (Doyle 1990,

[12]This means base revisions such that the corresponding belief revision operations satisfy all the AGM postulates.

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 the analyzed default logics 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 are natural postulates for belief revision 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 analyzing that class of revision operations in more detail in order to detect interesting and tractable special cases.

Acknowledgement

I would like to thank Gerd Brewka, Jon Doyle, Peter Gärdenfors, David Makinson, and David McAllester for discussions about the subject of this paper and Gerd Brewka and David Makinson for comments on an earlier draft. In particular, I am grateful to David Makinson for his extensive and helpful comments on the final draft. All remaining flaws are mine, of course.

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, C. E., Gärdenfors, P., Makinson, D. (1985): On the Logic of Theory Change: Partial Meet Contraction and Revision Functions, in: Journal of Symbolic Logic 50(2), 510–530.

Brachman, R. J. (1990): The Future of Knowledge Representation, in: Proceedings of the 8th National Conference of the American Association for Artificial Intelligence, Boston, Mass., 1082–1092.

Brewka, G. (1989): Preferred Subtheories: An Extended Logical Framework for Default Reasoning, in: Proceedings of the 11th International Joint Conference on Artificial Intelligence, Detroit, Mich., 1043–1048.

Brewka, G. (1990): Nonmonotonic Reasoning: Logical Foundations of Commonsense, Cambridge University Press, Cambridge, England. To appear.

Dalal, M. (1988): Investigations Into a Theory of Knowledge Base Revision: Preliminary Report, in: Proceedings of the 7th National Conference of the American Association for Artificial Intelligence, Saint Paul, Minn., 475–479.

Doyle, J. (1990): Rational Belief Revision. Presented at the Third International Workshop on Nonmonotonic Reasoning, Stanford Sierra Camp, Cal.

Fagin, R., Ullman, J. D., Vardi, M. Y. (1983): On the Semantics of Updates in Databases, in: 2nd ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, Atlanta, Ga., 352–365.

Gärdenfors, P., Makinson, D. (1988): Revision of Knowledge Systems Using Epistemic Entrenchment, in: Vardi, M. (ed.): Proceedings of the 2nd Workshop on Theoretical Aspects of Reasoning about Knowledge, Morgan Kaufmann, Los Altos, Cal.

Gärdenfors, P. (1988): Knowledge in Flux—Modeling the Dynamics of Epistemic States, MIT Press, Cambridge, Mass.

Gärdenfors, P. (1990): Belief Revision and Nonmonotonic Logic: Two Sides of the Same Coin?, in: Aiello, L. C. (ed.): Proceedings of the 9th European Conference on Artificial Intelligence, Stockholm, Sweden, 768–773.

Garey, M. R., Johnson, D. S. (1979): Computers and Intractability—A Guide to the Theory of NP-Completeness, Freeman, San Francisco, Cal.

Ginsberg, M. L., Smith, D. E. (1987): Reasoning About Action I: A Possible Worlds Approach, in: Brown, F. M. (ed.): The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop, Morgan Kaufmann, Los Altos, Cal., 233–258.

Ginsberg, M. L. (1986): Counterfactuals, in: Artificial Intelligence 30(1), 35–79.

Katsuno, H., Mendelzon, A. O. (1989): A Unified View of Propositional Knowledge Base Updates, in: Proceedings of the 11th International Joint Conference on Artificial Intelligence, Detroit, Mich., 1413–1419.

Katsuno, H., Mendelzon, A. O. (1990): 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.

Makinson, D., Gärdenfors, P. (1990): Relations between the Logic of Theory Change and Nonmonotonic Logic, in: Fuhrmann, A., Morreau, M. (eds.): Proceeding of the Konstanz Workshop on Belief Revision, Springer-Verlag, Berlin, Germany. To appear.

McAllester, D. A. (1982): Reasoning Utility Package User's Manual, AI Memo 667, AI Laboratory, Massachusetts Institute of Technology, Cambridge, Mass.

Myers, K. L., Smith, D. E. (1988): The Persistence of Derived Information, in: Proceedings of the 7th National Conference of the American Association for Artificial Intelligence, Saint Paul, Minn., 496–500.

Nebel, B. (1989): A Knowledge Level Analysis of Belief Revision, in: Brachman, R. J., Levesque, H. J., Reiter, R. (eds.): Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference, Toronto, Ont., 301–311.

Nebel, B. (1990): Reasoning and Revision in Hybrid Representation Systems, volume 422 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Germany.

Newell, A. (1982): The Knowledge Level, in: Artificial Intelligence 18(1), 87–127.

Poole, D. (1988): A Logical Framework for Default Reasoning, in: Artificial Intelligence 36, 27–47.

Rao, A. S., Foo, N. Y. (1989): Formal Theories of Belief Revision, in: Brachman, R. J., Levesque, H. J., Reiter, R. (eds.): Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference, Toronto, Ont., 369–380.

Winslett, M. S. (1988): Reasoning about Action Using a Possible Models Approach, in: Proceedings of the 7th National Conference of the American Association for Artificial Intelligence, Saint Paul, Minn., 89–93.

Winslett, M. S. (1989): Sometimes Updates are Circumscription, in: Proceedings of the 11th International Joint Conference on Artificial Intelligence, Detroit, Mich., 859–863.