Base Revision Operations and Schemes: Semantics, Representation, and Complexity

Bernhard Nebel[1]

Abstract. The theory of belief revision developed by Gärdenfors and his colleagues characterizes the classes of reasonable belief revision operations. However, some of the assumptions made in the theory of belief revision are unrealistic from a computational point of view. We address this problem by considering revision operations that are based on a priority ordering over a set of sentences representing a belief state instead of using preference relations over all sentences that are accepted in a belief state. In addition to providing a semantic justification for such operations, we investigate also the computational complexity. We show how to generate an epistemic entrenchment ordering for a belief state from an arbitrary priority ordering over a set of sentences representing the belief state and show that the resulting revision is very efficient. Finally, we show that some schemes for generating revision operations from bases can encode the preference relations more concisely than others.

1 INTRODUCTION

The problem of changing a belief state in the face of new information arises in a number of areas in computer science and artificial intelligence, e.g., in updating logical database, in hypothetical reasoning, etc. Most of the research in this area is influenced by work in philosophical logic, in particular by Gärdenfors and his colleagues [1, 10, 9], who developed the *theory of belief revision* (see also [13].)

The main topic of the work by Gärdenfors and his colleagues is the characterization of *classes* of reasonable *revision operations*, where belief states are modeled by *deductively closed sets of propositional sentences* (so-called *belief sets*) [10] or, equivalently, as a set of models [17]. Starting with a number of *rationality postulates*, the set of possible change operations on a belief state is constrained to those that fulfill the postulates. Based on that, it is possible to define and analyze specific *revision schemes* that generate revision operations by employing some *preference* information. One such revision scheme is, for instance, the *partial meet revision* scheme [1]. Another scheme uses so-called *epistemic entrenchment orderings* in order to generate revision operations [12].

If one wants to apply this theory in a computer science or artificial intelligence application, there are two severe problems. First of all, the assumption that belief states are modelled by deductively closed sets of sentences seems to be computationally infeasible. However, we can, of course, represent belief sets by a finite *belief base*, a finite set of sentences that is logically equivalent with the belief set (provided the belief set is finite modulo logical equivalence). Secondly, there is the problem that the extra information required by the above mentioned revision schemes is usually a relation over the set of all sentences in a deductively closed theory, which is representationally infeasible.

One way to address the two problems is to consider revision operations on belief bases (so-called *base revisions*), i.e., operations that modify a belief base instead of a belief set [7, 8, 15, 18, 20, 24]. Such an approach often also matches certain characteristics of an application setting very well. If, for instance, a code of norms or a scientific or naive theory of the world is represented by a set of explicitly stated sentences, one may want to express preferences between these sentences.

However, the base revision approach violates the principle of "irrelevance of syntax" [3, 17], i.e., base revisions may lead to different results for belief bases that are syntactically different but logically equivalent. We address this criticism by proposing a different view on base revisions. Instead of analyzing them on the base level, we consider equivalent belief revision operations that are *generated* by the preferences on the base.

We focus on revision schemes that are as efficient as possible, namely, those that generate revision operations that are computable using a polynomial number of NP-oracle calls and which allow a concise representation of the revised base.

2 PRELIMINARIES

Throughout this paper, a finitary propositional language $\mathcal{L}$ with the usual logical connectives ($\neg$, $\vee$, $\wedge$, $\rightarrow$ and $\leftrightarrow$) is assumed. The finite alphabet of propositional variables $p, q, r \ldots$ is denoted by Σ, propositional sentences by $\tau, \omega, \phi, \psi, \chi, \ldots$, constant truth by $\top$, its negation by $\bot$, and 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., $Cn(K) \stackrel{\text{def}}{=} \{\phi \in \mathcal{L} | \; K \vdash \phi\}$. Instead of $Cn(\{\phi\})$, we will also write $Cn(\phi)$. Deductively closed sets of propositional sentences, i.e., $K = Cn(K)$, are denoted by the capital letters $K, L, M \ldots$ and are called **belief sets**. The set of all belief sets over $\mathcal{L}$ is denoted by $Th_{\mathcal{L}}$. The monotonic addition of a propositional sentence ϕ to a belief set K, i.e., $Cn(K \cup \{\phi\})$, is denoted by $K + \phi$ and called **expansion** of K by ϕ. Arbitrary sets of sentences are called **belief bases** and are denoted by capital letters from the beginning

[1] DFKI, Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany and Universität Ulm, Fakultät für Informatik, James-Franck-Ring, 89081 Ulm, Germany, e-mail: nebel@informatik.uni-ulm.de.

© 1994 B. Nebel
ECAI 94. *11th European Conference on Artificial Intelligence* Edited by A. Cohn
Published in 1994 by John Wiley & Sons, Ltd.

of the alphabet. 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 $\bigvee \emptyset = \bot$.

A **belief revision operation** is a function [1, 10]:

$$\dot{+}\colon Th_{\mathcal{L}} \times \mathcal{L} \quad \rightarrow \quad Th_{\mathcal{L}}, \tag{1}$$

where the result of the operation is denoted by $K \dot{+} \phi$, which is intended to be a consistent belief set that differs *minimally* from K and contains ϕ. While these conditions do not lead to a unique operation, it is possible to constrain the space of reasonable belief revision operations. Gärdenfors proposed the following set of **rationality postulates for revision operations**:

($\dot{+}$1) $K \dot{+} \phi$ is a belief set;
($\dot{+}$2) $\phi \in K \dot{+} \phi$;
($\dot{+}$3) $K \dot{+} \phi \subseteq K + \phi$;
($\dot{+}$4) If $\neg\phi \notin K$, then $K + \phi \subseteq K \dot{+} \phi$;
($\dot{+}$5) $K \dot{+} \phi = Cn(\bot)$ only if $\vdash \neg\phi$;
($\dot{+}$6) If $\vdash \phi \leftrightarrow \psi$ then $K \dot{+} \phi = K \dot{+} \psi$;
($\dot{+}$7) $K \dot{+} (\phi \wedge \psi) \subseteq (K \dot{+} \phi) + \psi$;
($\dot{+}$8) If $\neg\psi \notin K \dot{+} \phi$, then $(K \dot{+} \phi) + \psi \subseteq K \dot{+} (\phi \wedge \psi)$.

These postulates intend to capture the intuitive meaning of minimal change—from a logical point of view [1, 10, 11, 13]. The first six postulates, which are straightforward, are called **basic postulates**. The two last postulates, which are called **supplementary postulates**, are less obvious. They capture the idea that a revision of K by a conjunction $(\phi \wedge \psi)$ should be achieved through a revision by ϕ and an expansion by ψ, if this is possible at all, i.e., if ψ is consistent with $K \dot{+} \phi$.

One interesting point to note is that the postulates do not constrain the revision operation with respect to varying belief sets. In other words, we may restrict ourselves to a given belief set and consider the mapping from $\mathcal{L}$ (the new information) to $Th_{\mathcal{L}}$ (the revised belief set).

Based on the above framework, one can consider different *schemes* to generate revision operations. Formally, a **revision scheme** maps a belief set and some extra information, which encodes preferences, to a belief revision operation on the given belief set. In our setting (of finitary propositional logic), such a scheme can be considered as an *algorithm* that computes the revision.

In [1], so-called partial meet revisions are investigated. This notion is based on systems of maximal (w.r.t. to set-inclusion) subsets of a given belief set K that do not allow the derivation of ϕ, called the **removal** of ϕ and written $K \downarrow \phi$:

$$K \downarrow \phi \quad \stackrel{\text{def}}{=} \quad \{L \subseteq K \mid L \nvdash \phi, \forall M\colon L \subset M \subseteq K \Rightarrow M \vdash \phi\}. \tag{2}$$

A **partial meet revision** (on K for all ϕ) is defined by a *selection function* γ that selects a nonempty subset of $K \downarrow \neg\phi$ (provided $K \downarrow \neg\phi$ is nonempty, $\emptyset$ otherwise) in the following way:

$$K \dot{+} \phi \quad \stackrel{\text{def}}{=} \quad \left(\bigcap \gamma(K \downarrow \neg\phi)\right) + \phi. \tag{3}$$

Such partial meet revisions satisfy unconditionally the first six postulates. Furthermore, it is possible to show that all revision operations satisfying the basic postulates are partial meet revisions [1, Observation 2.5].

Instead of providing the preference information by a selection function, one may also think of preference relations over sentences. **Epistemic entrenchment orderings**, written as $\phi \preceq \psi$, are defined over the entire set of sentences $\mathcal{L}$ and have to satisfy the following properties:

($\preceq$1) If $\phi \preceq \psi$ and $\psi \preceq \chi$, then $\phi \preceq \chi$.
($\preceq$2) If $\phi \vdash \psi$, then $\phi \preceq \psi$.
($\preceq$3) For any ϕ, ψ, $\phi \preceq (\phi \wedge \psi)$ or $\psi \preceq (\phi \wedge \psi)$.
($\preceq$4) When $K \neq Cn(\bot)$, then $\phi \notin K$ iff $\phi \preceq \psi$ for all $\psi \in \mathcal{L}$.
($\preceq$5) If $\psi \preceq \phi$ for all $\psi \in \mathcal{L}$, then $\vdash \phi$.

From ($\preceq$1) and ($\preceq$3), it follows that an epistemic entrenchment ordering is a complete preorder over $\mathcal{L}$. The strict part of this preorder will be denoted by $\prec$ in the following.

Using such a relation, one can define a revision scheme, which we will call **cut revision**:

$$K \dot{+} \phi \quad \stackrel{\text{def}}{=} \quad \{\psi \in K \mid \neg\phi \prec \psi\} + \phi. \tag{4}$$

From results by Gärdenfors and Makinson [12] and Rott [23], it follows that class of belief revision operations generated by this scheme coincides with the class of revision operations satisfying all rationality postulates.

3 PRIORITIZED MEET BASE-REVISION

Although the theory sketched above provides us with a good picture of the ways a belief set can be revised, it does seem not to be possible to use this theory for implementing belief revision on a computer system. First of all, the assumption that a belief state is modeled as a deductively closed set of sentences sounds unrealistic. Secondly, the amount of preference information needed for the partial meet revision and the epistemic entrenchment scheme seems to be prohibitive, e.g. as has been shown by Gärdenfors and Makinson, one needs an ordering over $2^{|\Sigma|}$ sentences to specify an entrenchment ordering.

Addressing these problems, we consider revision schemes that use preference information that has a size polynomial in the size of a belief base. In particular, we focus on schemes where the preference information is encoded by a *complete preorder* $\trianglelefteq$ over the set of sentences in a belief base, also called **epistemic relevance ordering** [19, 20], with the intuitive meaning that if $\phi \trianglelefteq \psi$, then ψ is at least as relevant, important, or reliable than ϕ (see also, e.g., [2, 7, 14]). Equivalently, we can view a base A as partitioned into n **priority classes** or *ranks* $A_1, \ldots, A_n$ as follows (using $\triangleleft$ to denote the strict part of $\trianglelefteq$):

$$\phi \in A_1 \quad \text{iff} \quad \forall \psi \in A\colon \phi \trianglelefteq \psi,$$
$$\phi \in A_{i+1} \quad \text{iff} \quad \exists \psi \in A_i\colon \Big(\psi \triangleleft \phi \wedge \neg(\exists \chi \in A\colon \psi \triangleleft \chi \triangleleft \phi)\Big).$$

The union of the highest priority classes down to the jth class, i.e., $\bigcup_{i=j}^{n} A_i$, will also be written as $\overline{A_j}$.

Based on a such an epistemic relevance ordering, one can define a removal operation , called **prioritized base-removal** and written $A \Downarrow \phi$, that keeps as many sentences with high priority as possible:

$$A \Downarrow \phi \quad \stackrel{\text{def}}{=} \quad \{B \subseteq A \mid B \nvdash \phi, \forall C, j\colon B \cap \overline{A_j} \subset C \cap \overline{A_j} \Rightarrow C \cap \overline{A_j} \vdash \phi\}. \tag{5}$$

Similar to partial meet revision, we define **prioritized meet base-revision**, written $A \oplus \phi$, on a prioritized base A as follows:

$$A \oplus \phi \quad \stackrel{\text{def}}{=} \quad \Big(\bigcap_{B \in (A \Downarrow \neg\phi)} Cn(B) \Big) + \phi. \tag{6}$$

As mentioned above, we will view $\oplus$ as a base-revision *scheme*, i.e., as defining a belief-revision operation $\dot{+}$ on $Cn(A)$ using A and $\trianglelefteq$ as "parameters:"

$$Cn(A) \dot{+} \phi \quad \stackrel{\text{def}}{=} \quad A \oplus \phi. \tag{7}$$

Although the construction looks quite plausible, it leads to a number of problems. First of all, revision operations generated by the prioritized meet base-revision scheme do not satisfy all rationality postulates [20]. Secondly, there is the problem of constructing a representation of the result of the revision operation. In our case of belief sets that are finite modulo logical equivalence, such a base can be easily specified. If A is a prioritized belief base then

$$A \oplus \phi \quad = \quad Cn\Big((\bigvee (A \Downarrow \neg\phi)) \wedge \phi \Big). \tag{8}$$

However, the resulting base looks quite unintuitive. Furthermore, the cardinality of $(A \Downarrow \phi)$ cannot be polynomially bounded as the following example demonstrates. Let

$$A \quad = \quad \{p_1, \ldots, p_m, q_1, \ldots, q_m\} \tag{9}$$

$$\phi \quad = \quad \bigwedge_{i=1}^{m} (p_i \leftrightarrow \neg q_i), \tag{10}$$

and assume that there is just one priority class. Clearly, $(A \Downarrow \neg\phi)$ has exponentially many elements.[2]

Thirdly, prioritized meet base-revision is computational very difficult. Determining whether a proposition follows from a revised prioritized base is Π_2^p-complete[3][20, 6].

The representational problem mentioned above could be avoided, if we defined the revision by considering the intersection over all maximal consistent sub-bases instead of using the intersection over the *consequences* of the maximal consistent sub-bases—a method that has been called *when in doubt, throw it out* (WIDTIO) [25]. However, the computational problems would remain. Deciding whether a proposition follows from a WIDTIO-revision is also Π_2^p-hard [6].

4 CUT BASE-REVISION

As mentioned above, prioritized meet base-revision suffers from a number of problems. Accounting for these problems, we will consider a base-revision scheme that resembles the *cut revision scheme* on belief sets (see Eq. (4)).

As before, we assume that the base is partitioned into priority classes $A_1, \ldots, A_n$. The **cut base-revision**, written $A \otimes \phi$, is then defined as follows:

$$A \otimes \phi \quad \stackrel{\text{def}}{=} \quad Cn(\{\psi \in A | \, \psi \in A_j, \overline{A_j} \not\vdash \neg\phi\}) + \phi. \tag{11}$$

[2] In fact, from a recent result by Cadoli, Donini, and Schaerf (personal communication), it follows that in the general case it is impossible to find a dense (i.e., polynomial) representation of a revised base unless the polynomial hierarchy collapses.

[3] It is assumed that the reader is familiar with basic notions of complexity theory (see, e.g., [16]).

Viewing $\otimes$ as a revision scheme, one easily verifies that all basic revision postulates are satisfied.

Proposition 1 *Belief revision operations generated by the cut base-revision scheme satisfy the basic rationality postulates.*

Moreover, contrary to the prioritized meet base-revision scheme, the result of a revision can be straightforwardly represented as a belief base, namely, as $\overline{A_j} \cup \{\phi\}$ with j being the smallest index such that $\overline{A_j}$ is consistent with $\{\phi\}$. Besides the representational advantages, cut base-revisions are also better behaved from a computational point of view.

Proposition 2 *Deciding $A \otimes \phi \vdash \psi$ is in $\mathrm{P}^{\mathrm{NP}[O(\log n)]}$.*

Since belief revision in general is a problem that is NP-hard and co-NP-hard [20], the above result is close to the optimum. Further, if we reduce the complexity of propositional reasoning, for instance, by restricting the expressiveness, we can obtain a polynomial-time revision scheme.

While all the above sounds as if cut base-revisions are much more well-behaved than prioritized meet base-revisions, there are also apparent disadvantages. Firstly, cut base-revision is much more radical in giving up beliefs than prioritized meet base-revision (or even WIDTIO-revisions). It cuts away all priority classes that have a priority equal or lower to the class that is "responsible" for an inconsistency. For this reason, one might argue that this kind of revision violates the principle of *minimal modification.* Secondly, it is not clear whether cut base-revision satisfies the supplementary postulates as well. In order to address these problems, we will investigate the relationship between epistemic relevance and epistemic entrenchment.

5 RELEVANCE VS. ENTRENCHMENT IN CUT BASE-REVISION

Epistemic relevance differs from epistemic entrenchment in two aspects. Firstly, the former is a complete preorder over a belief base while the second is a complete preorder over $\mathcal{L}$. Secondly, epistemic entrenchment respects the logical contents of the sentences while epistemic relevance is an arbitrary preorder. For instance, we may well have the case that $\phi \vdash \psi$ but $\psi \lhd \phi$, contradicting ($\preceq$2).

In the example above, it does not seem to make much sense to give ϕ higher priority than ψ since ϕ has to be retracted in any case if ψ is forced to be deleted. More generally, if the sub-base $C \subseteq A$ implies $\phi \in A$, then it does not make much sense that ϕ has a priority that is lower than the minimum of the priorities of the sentences in C. For this reason, let us assume that the epistemic relevance ordering satisfies the following **priority consistency condition** (PCC) (see also [21, 22]):

> For all $\phi \in A$, if C is a nonempty subset of A such that $C \vdash \phi$, then there exists $\chi \in C$ such that $\chi \trianglelefteq \phi$.

As has been shown by Rott [22], this condition is necessary and sufficient for the **extendibility** of $\trianglelefteq$ on A to an epistemic entrenchment ordering, i.e., an epistemic entrenchment ordering $\preceq$ satisfying for all $\phi, \psi \in A$ that $\phi \trianglelefteq \psi$ iff $\phi \preceq \psi$.

Rott [22] gives the following construction to show sufficiency (adapted to finite bases). Let A be a prioritized base that satisfies (PCC) and assume w.l.g. that $A_n = \emptyset$. Then define $\preceq$ on $\mathcal{L}$ by

$$\phi \preceq \psi \quad \text{iff} \quad A \not\vdash \phi \text{ or } \max_j(\overline{A_j} \vdash \phi) \leq \max_i(\overline{A_i} \vdash \psi) \quad (12)$$

This construction does not only show that a priority consistent epistemic relevance ordering $\trianglelefteq$ can be extended to an entrenchment ordering $\preceq$, but also provides us with an entrenchment ordering that **minimally extends** the relevance ordering, i.e., if $\phi \prec \psi$, then all other entrenchments $\preceq'$ extending $\trianglelefteq$ also contain $\phi \prec' \psi$.

Theorem 3 *Let A be a base with an associated epistemic relevance ordering $\trianglelefteq$ satisfying (PCC). Then $\preceq$ determined by Eq. (12) is the unique minimal extension of $\trianglelefteq$ to an epistemic entrenchment ordering over $Cn(A)$.*

Even more interestingly, the cut base-revision operation on a priority consistent base coincides with the cut revision using the epistemic entrenchment ordering generated by Eq. (12).

Theorem 4 *Let A be a belief base and $\trianglelefteq$ be an associated epistemic relevance ordering satisfying (PCC). Let $\preceq$ be the epistemic entrenchment order generated from $A, \trianglelefteq$ according to Eq. (12) and $\dot{+}$ be the cut revision on $Cn(A)$ using $\preceq$. Then*

$$Cn(A) \dot{+} \phi = Cn(A \otimes \phi). \quad (13)$$

In other words, priority-consistent epistemic relevance orderings can be viewed as a dense representation of an epistemic entrenchment ordering. Furthermore, assuming that belief sets are finite modulo logical equivalence, all belief revision operations generated by the cut revision (i.e., all fully rational revision operations) are generated by some priority-consistent epistemic relevance ordering. Hence, the cut base-revision scheme can generate all fully rational revision operations.

This leaves us with the question of what kind of belief revision operation is generated if the epistemic relevance ordering is not priority consistent. This question is important because it seems not very realistic to put the burden of guaranteeing this condition, which involves deciding propositional decidability, on the shoulders of a user.

As it turns out, we can interpret the specified priorities as an *approximation* specifying *lower bounds* for the intended priorities (see also [4, 5]). Using (12) on some arbitrary prioritized base A, the resulting relation $\preceq$ is again an epistemic entrenchment ordering for the generated belief set $Cn(A)$. Furthermore, $\preceq$ satisfies a number of conditions that show that $\preceq$ can be indeed regarded as the epistemic entrenchment relation intended by the priorization of A.

Theorem 5 *Let A be an arbitrary prioritized base. Then the ordering $\preceq$ generated by Eq. (12) is an epistemic entrenchment ordering for $Cn(A)$, and the restriction of $\preceq$ to A is a priority-consistent epistemic relevance ordering that generates $\preceq$ according to Eq. (12).*

Furthermore, the cut base-revision coincides with the cut revision using the generated epistemic entrenchment.

Theorem 6 *Let A be an arbitrary prioritized base and $\dot{+}$ be a cut revision operation based on the epistemic entrenchment ordering generated by (12). Then*

$$Cn(A) \dot{+} \phi = Cn(A \otimes \phi). \quad (14)$$

From that the following corollary follows straightforwardly.

Corollary 7 *The cut base-revision scheme coincides with the set of belief revision operations that satisfy all rationality postulates.*

6 LINEAR BASE-REVISION

While prioritized meet base-revision has a lot of conceptual, representational, and computational problems, cut base-revision seems to be too radical in deleting beliefs. In trying to find a compromise, one may consider the method of deleting an entire priority class if one sentence in it leads to a contradiction that cannot be blamed on sentences in lower priority classes, but keeping as many of the other priority classes as possible.[4] We could view this scheme as if we had prioritized bases where each priority class contains only one element. Formally, given a prioritized base A with n priority classes, we form a new prioritized base with n classes where $B_i = \{\bigwedge A_i\}$. The resulting epistemic relevance ordering on B is then a linear order.

Interestingly, the prioritized meet base-revision scheme on linearly ordered bases behaves in the way as spelled out above, i.e., it deletes a priority class if it is to be blamed for a contradiction and no lower class can be blamed for it. Since prioritized meet base-revision on linearly ordered bases seems to be an important special case, we will consider it as a base-revision scheme on its own, as the **linear base-revision scheme**. The **linear base-revision operation** will be denoted by $\odot$.

As can be easily shown, the linear base-revision operation picks just one subset of the base as the result of the revision [20, Proposition 7] (see also [18]). Furthermore, the linear base-revisions scheme satisfies all rationality postulates [20, Theorem 8]. Finally, also the computational properties of this scheme are very attractive.

Theorem 8 *The problem of deciding $A \odot \phi \vdash \psi$ is $\mathrm{P}^{\mathrm{NP}[O(n)]}$-complete.*

One interesting question is how linear base-revision relates to cut base-revision. As can be easily shown, cut base-revisions can be polynomially transformed to linear base-revisions.

Proposition 9 *Let A be a prioritized base. Then there exists a linearly prioritized base B that can be computed in polynomial time such that*

$$Cn(A \otimes \phi) \quad = \quad Cn(B \odot \phi). \quad (15)$$

Using this result and [20, Theorem 8], the following corollary is immediate.

Corollary 10 *The linear base-revision scheme coincides with the set of belief revision operations satisfying all rationality postulates.*

[4] This is also the intention of the revision of possibilistic knowledge bases described in [5, p. 167].

Although this result shows that everything that can be expressed as a cut base-revision can be expressed as a linear base-revision (and *vice versa*), it does not imply that the representation of the bases are equally concise. In fact, the complexity results show that this can only be the case if $\mathrm{P}^{\mathrm{NP}[O(\log^c n)]} = \mathrm{P}^{\mathrm{NP}[O(n)]}$, which is considered to be rather unlikely. In fact, the most natural transformation of linear base-revision to cut base-revision blows up the belief base exponentially.

Assume a linearly ordered prioritized base A with n priority classes. Then we define a logically equivalent belief base $B = \pi(A)$ with $2^n - 1$ priority classes. The priority classes of B are again singletons, and the elements of these classes are disjunctions of the classes in A. In particular, the sentence $A_{j_1} \vee \ldots \vee A_{j_k}$ is in the priority class B_l, where $l = \sum_{l=1}^{k} 2^{(j_l - 1)}$.

Theorem 11 *Let A be a prioritized base with a linear epistemic relevance ordering. Then it holds that*

$$Cn(A \odot \phi) \quad = \quad Cn(\pi(A) \otimes \phi). \tag{16}$$

As mentioned above, it seems unlikely that we can find a transformation that is "cheaper," i.e., uses less than exponential time. Hence, it appears to be the case that the linear base-revision scheme uses a more concise coding of preference information than the cut base-revision scheme.

7 CONCLUDING REMARKS

Revision schemes that generate a belief revision operation from a given belief base and some additional preference information that is not of infeasible size seem to be of practical interest in the area of belief revision. The most straightforward such scheme, *prioritized meet base-revision*, has a number of conceptual and computational problems, though. Adopting the notion of *epistemic entrenchment*, we were able to show that priorities of sentences in belief bases can be interpreted naturally as lower bounds of epistemic entrenchment, and based on this view, it is possible to define an elegant and efficient revision scheme, called *cut base-revision* scheme. Relating this scheme to earlier results concerning prioritized meet base-revision on linearly ordered prioritized bases, we noted that the latter is expressively equivalent to the former. Furthermore, we noted that it appears to be the case that prioritized meet base revision on linearly ordered bases permits to state the preference information in a way that is more concise than in the case of cut base revisions.

ACKNOWLEDGEMENTS

I would like to thank Didier Dubois, Dov Gabbay, and Henri Prade for discussion on the subject of this paper, Georg Gottlob and Werner Nutt for some comments on computational complexity issues, and Hans Rott and the anonymous referees for comments on an earlier version of this paper.

This work was supported by the German Ministry for Research and Technology (BMFT) under grant ITW 8901 8 as part of the WIP project and by the European Commission as part of DRUMS-II, the ESPRIT Basic Research Project P6156.

REFERENCES

[1] C. E. Alchourrón, Peter Gärdenfors, and David Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *J. Symbolic Logic*, **50**(2), 510–530, (1985).

[2] G. Brewka, 'Preferred subtheories: An extended logical framework for default reasoning', in *Proc. IJCAI-89*, pp. 1043–1048, Detroit, MI, (1989).

[3] M. Dalal, 'Investigations into a theory of knowledge base revision: Preliminary report', In *Proc. AAAI-88*, pp. 475–479.

[4] D. Dubois and H. Prade, 'Epistemic entrenchment and possibilistic logic', *Artificial Intelligence*, **50**, 223–239, (1991).

[5] D. Dubois and H. Prade, 'Belief change and possibility theory', In Gärdenfors [9], 142–182.

[6] T. Eiter and G. Gottlob, 'On the complexity of propositional knowledge base revision, updates, and counterfactuals', *Artificial Intelligence*, **57**, 227–270, (1992).

[7] R. Fagin, J. D. Ullman, and M. Y. Vardi, 'On the semantics of updates in databases', in *Proc. PODS-83*, pp. 352–365, Atlanta, Ga., (1983).

[8] A. Fuhrmann, 'Theory contraction through base contraction', *J. Philosophical Logic*, **20**, 175–203, (1991).

[9] *Belief Revision*, ed., P. Gärdenfors, Cambridge University Press, Cambridge, UK, 1992.

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

[11] P. Gärdenfors, 'Belief revision: An introduction', In Gärdenfors [9], 1–28.

[12] P. Gärdenfors and D. Makinson, 'Revision of knowledge systems using epistemic entrenchment', in *Proc. TARK-88*, Asilomar, CA, (1988).

[13] P. Gärdenfors and H. Rott, 'Belief revision', Lund University Cognitive Studies 11, Cognitive Science, Department of Philosophy, University of Lund, Lund, Sweden, (July 1992). To appear in: D. Gabbay (ed.), *Handbook of Logic in AI and Logic Programming*.

[14] M. L. Ginsberg, 'Counterfactuals', *Artificial Intelligence*, **30**(1), 35–79, (October 1986).

[15] S. O. Hannson, *Belief base dynamics*, Doctoral Dissertation, Uppsala University, Sweden, 1991.

[16] D. S. Johnson, 'A catalog of complexity classes', in *Handbook of Theoretical Computer Science, Vol. A*, ed., J. van Leeuwen, 67–161, MIT Press, (1990).

[17] H. Katsuno and A. O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**, 263–294, (1991).

[18] B. Nebel, 'A knowledge level analysis of belief revision', in *Proc. KR-89*, eds., R. Brachman, H. J. Levesque, and R. Reiter, pp. 301–311, Toronto, ON, (1989). Morgan Kaufmann.

[19] B. Nebel, *Reasoning and Revision in Hybrid Representation Systems*, Springer-Verlag, Berlin, 1990.

[20] B. Nebel, 'Belief revision and default reasoning: Syntax-based approaches', in *Proc. KR-91*, eds., J. A. Allen, R. Fikes, and E. Sandewall, pp. 417–428, Cambridge, MA, (1991). Morgan Kaufmann.

[21] N. Rescher, *The Coherence Theory of Truth*, Oxford University Press, Oxford, UK, 1973.

[22] H. Rott, 'A nonmonotonic conditional logic for belief revision I', in *The Logic of Theory Change*, eds., A. Fuhrmann and M. Morreau, 135–183, Springer-Verlag, Berlin, (1991).

[23] H. Rott, 'Two methods of constructing contractions and revisions of knowledge systems', *J. Philosophical Logic*, **20**, 149–173, (1991).

[24] H. Rott, 'Belief contraction in the context of the general theory of rational choice', *J. of Symbolic Logic*, **58**, (1993).

[25] M. S. Winslett, 'Reasoning about action using a possible models approach', In *Proc. AAAI-88*, pp. 89–93.