Iterated Belief Revision: A Computational Approach

Yi Jin and Michael Thielscher

Department of Computer Science
Dresden University of Technology, Germany
{yijin,mit}@inf.tu-dresden.de

Abstract. The capability of revising its beliefs upon new information in a rational and efficient way is crucial for an intelligent agent. The classic AGM theory studies mathematically idealized models of belief revision in two aspects: the properties (i.e., the AGM postulates) a rational revision operator should satisfy; and how to construct concrete revision operators. In scenarios where new information arrives in sequence, rational revision operators should also respect postulates for iterated revision (e.g., the DP postulates). When applications are concerned, the idealization of the AGM theory has to be lifted, in particular, beliefs of an agent should be represented by a finite belief base. In this paper, we present a computational base revision operator, which satisfies the AGM postulates and postulates for iterated revision. We will show that our base revision operator is almost optimal in terms of computational complexity. Furthermore, the base revision operator's degrees of syntax relevance and minimal change are also formally analyzed.

Topic Area: Iterated Base Revision, Computational Complexity, Syntax Irrelevance, Minimal Change

Acknowledgments The first author is funded by the *Deutsche Forschungsgemeinschaft* under grant no. Gr 334/3.

Dagstuhl Seminar Proceedings 05321
Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics
http://drops.dagstuhl.de/opus/volltexte/2005/359

1 Introduction

The capability of revising its beliefs upon new information in a rational and efficient way is crucial for an intelligent agent. Technically, belief revision is the process of changing the beliefs of an agent to accommodate new, more precise, or more reliable information that is possibly inconsistent with the existing beliefs. In situations where the new information is consistent with the existing beliefs, the two can just be merged; we call this *mild revision*. More interesting and complicated are situations where the information conflicts with the prior beliefs, in which case the agent need to remove some of its currently held beliefs in order to consistently accommodate the new information. This kind of revision is referred to as *severe revision* [Lehmann, 1995].

In the literature, the classical AGM theory [Alchourrón *et al.*, 1985; Gärdenfors and Makinson, 1988; Gärdenfors, 1988] is a formal account of mathematically idealized models of belief revision in single agent environments. In AGM theory, beliefs of the agent are represented by a set of sentences in an underlying logical language. The agent is supposed to be a *Besserwisser*, i.e., it is aware of and responsible for all logical consequences of it beliefs. A belief revision operator is mathematically a mapping from the old beliefs and new information to the new beliefs. Any rational belief revision operator is supposed to satisfy a set of properties (AGM postulates). The guiding principle of AGM postulates is that of *economy of information* or *minimal change* of belief sets, which means not to give up currently held beliefs and not to generate new beliefs unless necessary.

For the incremental adaptation of beliefs, these postulates proved to be too weak [Darwiche and Pearl, 1997]. This has led to the development of additional postulates for iterated belief revision by Darwiche and Pearl (DP), among others. The underlying principle of DP postulates is that of minimal change of *preference information* which is exploited by revision operators. Still, however, the two sets of postulates are too permissive in that they support belief revision operators which assume arbitrary *implicit dependencies* among the pieces of information which an agent acquires along its way [Nayak *et al.*, 1996; Nayak *et al.*, 2003]. Based on a formal analysis, we proposed an *Independence* postulate as a solution to rule out revision operators which may abuse implicit dependence [Jin and Thielscher, 2005].

When applications are concerned, the mathematical idealization of AGM theory has to be lifted. Inevitably, the beliefs of the agent should be represented finitely by a belief base [Wobcke, 1992; Nebel, 1992; Rott, 1991]. Computational complexity of a belief revision operator becomes a very important criterion [Eiter and Gottlob, 1992; Nebel, 1998; Liberatore, 1997], since it determines how efficiently the agent can revise its beliefs . In this paper, we present a general base revision operator, which satisfies all AGM and DP postulates. The base revision operator does not abuse the implicit dependence in the sense it respects the postulate of Independence. We will show that our base revision operator is almost optimal in terms of computational complexity. Furthermore, we will formally analyze the base revision operator's degrees of syntax relevance and minimal change.

The rest of the paper is organized as follows. In the next section, we first introduce some notation which will be used throughout the paper. Then we recall the classical AGM theory in a propositional setting (including the AGM postulates and two characterizations of them), followed by a single-step base revision. In Section 3, we present a general base revision operator after recapitulating the basic ideas of iterated belief revision. The whole Section 4 contributes to the formal assessment of our iterated base revision operator, including its logical set-theoretical properties, computational complexity, and degrees of syntax relevance and minimal change. We conclude in Section 5 with a detailed comparison to related work.

2 Background

2.1 Preliminaries

We will work in a propositional language $\mathcal{L}$. The language is that of classical propositional logic with an associated consequence operation Cn in the sense that $Cn(X) = \{A : X \vdash A\}$, where X is a set of sentences. A set K of sentences is *logically closed* or called a *belief set* when $K = Cn(K)$. We denote the set of all belief sets in $\mathcal{L}$ by $\mathcal{K}$. If X, Y are two sets of sentences, $X + Y$ denotes $Cn(X \cup Y)$. $K + \varphi$ is a shorthand of $K + \{\varphi\}$, which is called *expansion* of K by φ. The set of all propositional interpretations (worlds) of $\mathcal{L}$ is denoted by $\mathcal{W}$. A set X of sentences is true in a world $w \in \mathcal{W}$ (denoted by $w \models X$), iff every element in X is true in w. For any set X of sentences, $[X]$ (called models of X) denotes the set of worlds in which every sentence in X is true. Conversely, given any set $S \subseteq \mathcal{W}$, $Th(S)$ denotes the set of sentences true in every world in S.

Given a total pre-order $\leq_U$ [1], we denote by $<_U$ its strict part, i.e., $\varphi <_U \psi$ iff $\varphi \leq_U \psi$ and $\psi \not\leq_U \varphi$, and $\alpha =_U \beta$ is just a shorthand of $\alpha \leq_U \beta$ and $\beta \leq_U \alpha$. Given a set U and a total pre-order $\leq_U$, the partition induced by $\leq_U$ on U is the one, s.t., two elements s, t are in the same class iff $s =_U t$. If the rest of the paper, whenever we talk about classes, they refer to those of the induced partition, unless explicitly specified. A class $C_1 \subseteq U$ is higher than another class $C_2 \subseteq U$, if there are $s \in C_1$ and $t \in C_2$, s.t., $t \leq_U s$. Given any set U and total pre-order $\leq_U$, we denote by $\min(U, \leq_U)$ the lowest class of the partition.

2.2 AGM Postulates

As a mathematical model, AGM theory assumes the beliefs of the agent are represented by a belief set, and the new information is represented by a sentence. Formally, for any belief set K and sentence φ, $K * \varphi$ stands for the result of belief revision when K is revised by φ. To provide general design criteria for

[1] A pre-order is a reflexive, transitive binary relation. A binary relation $\leq$ is total if $\alpha \leq \beta$ or $\beta \leq \alpha$, for arbitrary pair of elements α, β.

belief revision operators, Alchourrón, Gärdenfors, and Makinson (AGM) have developed a set of postulates: [2]

$(*1)$	$K * \varphi = Cn(K * \varphi)$	(Closure)
$(*2)$	$\varphi \in K * \varphi$	(Success)
$(*3)$	$K * \varphi \subseteq K + \varphi$	(Inclusion)
$(*4)$	If $\varphi \cup K$ is consistent, $K + \varphi \subseteq K * \varphi$	(Vacuity)
$(*5)$	$K * \varphi$ is consistent if φ is consistent	(Consistency)
$(*6)$	If $\varphi_1 \equiv \varphi_2$, $K * \varphi_1 = K * \varphi_2$	(Extensionality)
$(*7)$	$K * (\varphi_1 \wedge \varphi_2) \subseteq (K * \varphi_1) + \varphi_2$	(Superexpansion)
$(*8)$	If $\varphi_2 \cup (K * \varphi_1)$ is consistent, then $(K * \varphi_1) + \varphi_2 \subseteq K * (\varphi_1 \wedge \varphi_2)$	(Subexpansion)

One interesting point of AGM postulates is that they do not constrain operations wrt. varying belief sets. In other words. we can consider a revision operator a mapping (with an underlying belief set K) from $\mathcal{L}$ to $\mathcal{K}$. A revision operator of such kind is also referred to as a *local revision* [Hansson, 1998].

2.3 Characterizations of AGM Postulates

It is well known, a belief set K usually does not contain enough information to uniquely determine a revision operator [Gärdenfors, 1988; Gärdenfors and Makinson, 1988]. Therefore, we should exploit some kind of extra-logical *preference information*.[3] In particular, an *epistemic entrenchment* (EE for short) wrt. a belief set K is a binary relation over $\mathcal{L}$ which satisfies the following conditions:

(EE1) If $\alpha \leq_K \beta$ and $\beta \leq_K \gamma$, then $\alpha \leq_K \gamma$
(EE2) If $\alpha \models \beta$, then $\alpha \leq_K \beta$
(EE3) $\alpha \leq_K \alpha \wedge \beta$ or $\beta \leq_K \alpha \wedge \beta$, for any α and β
(EE4) If K is consistent, then $\alpha \notin K$ iff $\alpha \leq_K \beta$ for all β
(EE5) If $\beta \leq_K \alpha$ for all β , then $\models \alpha$

It follows from (EE1)-(EE3), that $\leq_K$ is a total pre-order. Conditions (EE2) and (EE5) says the highest class comprises all (and only) tautologous sentences. Condition (EE5) identifies $\mathcal{L} \backslash K$ to be the lowest class, if $K \neq \mathcal{L}$.

The intuitive meaning of $\alpha \leq_K \beta$ is sentence β is at least as entrenched as α , in the sense, when one of them has to be canceled, removing of α will be preferred. Given an EE $\leq_K$ wrt. the belief set K , we can uniquely determine a revision operator. The cut-set of a sentence φ wrt. $\leq_K$ [Rott, 1991] (denoted by $cut_{<_K}(\varphi)$) is defined as

$$cut_{<_K}(\varphi) = \{\psi \in \mathcal{L} \mid \varphi <_K \psi\} \tag{1}$$

[2] Readers are referred to [Alchourrón *et al.*, 1985] for detailed explanation and justification.

[3] Cf., Section 5 a discussion on revision operators which does not exploit extra-logical preference information.

The revision operator $*_{\leq_K}$ based on $\leq_K$ maps a sentence φ to the expansion of $cut_{<_K}(\neg\varphi)$ and φ:

$$K *_{\leq_K} \varphi \stackrel{\text{def}}{=} cut_{<_K}(\neg\varphi) + \varphi \tag{2}$$

Since the revision operator defined above amounts to cut away all sentences which is less entrenched than the negation of the sentence to be added, hence it is also called *cut revision* [Nebel, 1994].

It is not difficult to see, for a belief set there could be multiple corresponding EEs. As we have seen that an EE uniquely determines the result of the revision operator, which implies that the same belief set when revised with the same sentence may lead to different new belief sets due to different EEs. Such revision operators, whose results do not only depend on the logical contents of the original belief sets, are also referred to as *syntax based* [Nebel, 1991].

The following result shows a revision operator based on an EE satisfies all AGM postulate. Conversely, any revision operator which satisfies all AGM postulates can be generated by $*_{\leq_K}$ based on some EE $\leq_K$.

Theorem 1. [Gärdenfors and Makinson, 1988; Rott, 1991] *Given a belief set K, a revision operator $*$ satisfies all AGM postulates iff there exists an EE $\leq_K$ wrt K, s.t., for any sentence φ*

$$K * \varphi = K *_{\leq_K} \varphi$$

Revision operators based on EEs give a nice characterization of AGM postulates. Note, an EE is a preference relation over the set of sentences. There is another elegant characterization of the AGM postulates which is however more from semantics point of view, where the preference information is a relation over the set of possible worlds.

Given a belief set K, a total pre-order $\leq_{f(K)}$ on $\mathcal{W}$ is called a *faithful ranking* wrt. K [Katsuno and Mendelzon, 1991] iff it satisfies the following conditions:

- If $w_1, w_2 \models K$, then $w_1 =_{f(K)} w_2$.
- If $w_1 \models K$ and $w_2 \not\models K$, then $w_1 <_{f(K)} w_2$.

Intuitively, $w_1 \leq_{f(K)} w_2$ means w_1 is at least as plausible as w_2. Given a faithful ranking $\leq_{f(K)}$ wrt. the belief set K, the revision operator based it takes exactly the minimal models of the new sentence φ wrt. $\leq_{f(k)}$, as the models of the revised belief set.

Theorem 2. [Katsuno and Mendelzon, 1991] *Given a belief set K, a revision operator $*$ satisfies all AGM postulates iff there exists a faithful ranking $\leq_{f(K)}$ wrt. K s.t., for any sentence φ*

$$K * \varphi = Th(\min([\varphi], \leq_{f(K)}))$$

The vigilant reader may have observed that, in both characterization of AGM postulates, there is a problem of so-called categorial mis-matching [Hansson, 2003], in the sense the revision operator takes as input the belief set together with some preference information, whereas the output is merely a belief set. A direct consequence is that after one step of revision, we cannot iteratively apply the revision operator when there comes successively new information. This has been considered a difficult problem and extensively studied as the topic of iterated belief revision [Boutilier, 1993; Darwiche and Pearl, 1997; Lehmann, 1995; Jin and Thielscher, 2005]. We will come back to this point in Section 3, before we present the iterated base revision.

2.4 Base Revision

AGM theory gives a formal characterization of the space of rational revision operators and shows how to construct revision operators based on some preference information. However, there are two major problems if we want to apply AGM theory in a computer science or artificial intelligence application [Nebel, 1998; Jin and Thielscher, 2004]. First of all, it is assumed beliefs of the intelligent agent is logically closed, which seems representationally infeasible, since logically closed set of sentences are in general infinite. Similarly, extra-logic preference information is usually a relation over the set of all sentences [Gärdenfors and Makinson, 1988], or a relation over all possible worlds [Katsuno and Mendelzon, 1991], or a relation over all subsets of the belief set [Alchourrón *et al.*, 1985], which is huge even the belief set is finite modulo logical equivalence. Second problem is revision operators cannot be iterated [Darwiche and Pearl, 1997], when the new information comes in sequence.

In this section, we will address the first problem. Solutions to the second the problem are delegated to next section. In a computational framework, we assume that the beliefs of an agent are represented by *finite* sets of sentences, which are called *belief bases* [Wobcke, 1992]. From the computational point of view, the size of extra-logical preference information should be bounded polynomially in the size of the belief base. A *prioritized base* $\langle B, \leq_B \rangle$ consists of a belief base B and a total pre-order $\leq_B$ (called *epistemic relevance ordering* [Nebel, 1991], ERO for short) over the sentences in B .

Similar to the cut-set wrt. an EE as defined by (1), one can define cut-set wrt. an ERO $\leq_B$ (denoted by $cut_{\leq_B}(\varphi)$) for any sentence φ as

$$cut_{<_B}(\varphi) = \{\psi \in B \mid \{\chi \in B \mid \psi \leq_B \chi\} \not\models \varphi\} \tag{3}$$

Put in words, the cut-set of φ wrt. $\leq_B$ consists of all sentences in all high class, s.t., adding next lower class leads to the implication of φ .

Given an ERO $\leq_B$, the cut base revision $*_{\leq B}$ base on $\leq_B$ is defined for any sentence φ as

$$B *_{\leq_B} \varphi = cut_{<_B}(\neg\varphi) \cup \{\varphi\} \tag{4}$$

It is interesting to see how many of the AGM postulates are satisfied by the cut base revision. For this reason, we take $Cn(B)$ as the original belief set, and

$Cn(B *_{\leq_B} \varphi)$ as the revised belief set. First to observe is that although an ERO $\leq_B$ is a total pre-order over B, it can be generalized to an EE $\leq_{Cn(B)}$ wrt. $Cn(B)$ by letting

$$\varphi \leq_{Cn(B)} \psi \text{ iff } cut_{<_B}(\psi) \subseteq cut_{<_B}(\varphi) \tag{5}$$

Theorem 3. [Nebel, 1994] *Let B be a belief base with an ERO $\leq_B$. Then the relation $\leq_{Cn(B)}$ generated by (3) and (5) is an EE wrt. $Cn(B)$.*

Based on the generalized EE $\leq_{Cn(B)}$, we can define a revision operator $*_{\leq_{Cn(B)}}$ as defined by (2). The cut base revision $*_{\leq_B}$ based on $\leq_B$ has a very nice property that in terms of change of the logical contents it is equivalent to the revision operator $*_{\leq_{Cn(B)}}$. Therefore, we also say that the revision operator $*_{\leq_{Cn(B)}}$ is generated by $*_{\leq_B}$.

Theorem 4. [Nebel, 1994; Williams, 1994] *Let B be a belief base with an ERO $\leq_B$. Let $\leq_{Cn(B)}$ be the EE generated by (3) and (5). Then*

$$Cn(B *_{\leq_B} \varphi) = Cn(B) *_{\leq_{Cn(B)}} \varphi$$

It follows directly from Theorem 1, that cut base revisions based on EROs satisfy all AGM postulates. Furthermore, cut base revisions based on EROs are also flexible enough to generate all AGM revision operators.

Theorem 5. [Nebel, 1994] *The class of revision operations generated by cut base revision defined by (4) coincides with the class of revision operations satisfying all AGM postulates.*

3 Iterated Base Revision

Cut base revisions seems to be a perfect solution to the first problem which is mentioned in the last section. It is not difficult to observe that the second problem remains, since the result of cut base revision is merely a belief base.

3.1 Iterated Revision

Iterated belief revision is a very important topic from both theoretical and practical point of view, and has been studied by many researchers [Boutilier, 1993; Darwiche and Pearl, 1997; Lehmann, 1995; Jin and Thielscher, 2005]. Let's recall some basic ideas of the iterated revision: An iterated revision should not only produce a revised belief set (base), but also new preference information. A commonly accepted way to generate the new preference information is to apply the *minimal change* to the old preference information. Darwiche and Pearl have shown that AGM postulates are so overly permissive that they allow many unreasonable behaviors [Darwiche and Pearl, 1997] . The underlying reason is that AGM postulates put almost no constraints on the change of preference information. Therefore, Darwiche and Pearl proposed a set of (DP) postulates to complement the AGM postulates:[4]

[4] Readers are referred to [Darwiche and Pearl, 1997] for detailed explanation and justification.

(C1) If $\beta \models \varphi$, then $(K * \varphi) * \beta = K * \beta$.
(C2) If $\beta \models \neg\varphi$, then $(K * \varphi) * \beta = K * \beta$.
(C3) If $\varphi \in K * \beta$, then $\varphi \in (K * \varphi) * \beta$.
(C4) If $\neg\varphi \notin K * \beta$, then $\neg\varphi \notin (K * \varphi) * \beta$.

The following representation theorem shows the constraints DP postulates put on change of preference information. In particular, Condition (CR1) (Condition (CR2)) says the plausibilities of the worlds which satisfy (violate) the new information should changed in a uniformed way, s.t., their relative positions do not change.

Theorem 6. [Darwiche and Pearl, 1997] *Suppose that a revision operator $*$ satisfies all AGM postulates. The operator satisfies (C1)–(C4) iff its faithful rankings $f(K)$ and $f(K * \varphi)$ satisfy:*[5]

(CR1) If $w_1, w_2 \models \varphi$, then $w_1 \leq_{f(K)} w_2$ iff $w_1 \leq_{f(K*\varphi)} w_2$.
(CR2) If $w_1, w_2 \not\models \varphi$, then $w_1 \leq_{f(K)} w_2$ iff $w_1 \leq_{f(K*\varphi)} w_2$.
(CR3) If $w_1 \models \varphi$ and $w_2 \not\models \varphi$, then $w_1 <_{f(K)} w_2$ implies $w_1 <_{f(K*\varphi)} w_2$.
(CR4) If $w_1 \models \varphi$ and $w_2 \not\models \varphi$, then $w_1 \leq_{f(K)} w_2$ implies $w_1 \leq_{f(K*\varphi)} w_2$.

We have pointed out DP postulates are still too permissive [Jin and Thielscher, 2005] . In particular, A revision operator satisfying all DP postulates can still suffer from the problem of so-called *implicit dependence*, which can be informally explained as follows. According to the revision operator defined by (2),when a belief set K revised by $\neg\varphi$, any sentence $\psi \in K$ will be removed if $\varphi \geq_K \varphi \vee \psi$. It seems that φ is the only reason for ψ to exist in K (even they are not logically related), i.e., ψ implicitly depends on φ . Implicit dependence seems inevitable with AGM postulates, since any revision operator satisfying all AGM postulates can be generated by a revision operator based on some EE (Theorem 1). Therefore, the only thing we can do is to require a revision operator does not abuse implicit dependence, i.e., does not introduce undesired implicit dependence. The over-permissiveness of DP postulates can be evidenced by the *natural revision* [Boutilier, 1993], an iterated revision operator which satisfies all DP postulates. In natural revision, the new sentence is always least entrenched in the revised belief set, hence it implicitly depends on all other beliefs. As a consequence, a severe revision could cancel out all previous learned information, which is of course too radical in general.

To rule out revision operators which abuse implicit dependence, we proposed a postulate of *Independence*: [6]

(Ind) If $\neg\mu \notin K{*}\beta$ then $\mu \in (K{*}\mu){*}\beta$

The following representation theorem shows, postulate (Ind) impose a very natural constraint on the change of preference information.

[5] Recall: $f(K)$ is the original faithful ranking and $f(K * \varphi)$ is the faithful ranking after revision

[6] Reader are once again referred to [Jin and Thielscher, 2005] for a formal analysis of implicit dependence and detailed explanation and justification of (Ind)

Theorem 7. [Jin and Thielscher, 2005] *Suppose that a revision operator $*$ satisfies all AGM postulates. The operator satisfies Postulate (Ind) iff its faithful rankings $f(K)$ and $f(K * \varphi)$ satisfy:*

(IndR) *If $w_1 \models \mu$ and $w_2 \models \neg\mu$, then*
$w_1 \leq_{f(K)} w_2$ *implies* $w_1 <_{f(K*\mu)} w_2$.

Basically, Condition (IndR) says if a world w_1 satisfying the new information and another world w_2 violating the new information, then their relative distance in terms of plausibilities increases. It follows directly from Theorem 6 and Theorem 7, that with the presence of AGM postulates, (Ind) implies both (C3) and (C4).

3.2 A (General) Iterated Base Revision

In this subsection, we will show how to obtain an iterated base revision by adding some small ingredients to the recipe of cut base revision defined by (4). We first introduce a compact and more general belief representation: An *epistemic entrenchment base* (EE base for short) is a pair $\langle B, f\rangle$ consists of a belief base B and a mapping f from B to $\mathbb{N}^+$ (natural numbers greater than 0). For any sentence $\beta \in B$, we call $f(\beta)$ its *evidence degree*. In an EE base the size of extra-logic preference information is equal to the cardinality of its belief base. We denote by B^m the set of sentences in B which have at least evidence degree m, i.e., $B^m = \{\beta \in B \mid f(\beta) \geq m\}$. The maximal evidence degree of non-tautologous sentences in B is denoted by $max(B)$, that is, $max(B) = max\{f(\beta) \mid \beta \in B \text{ and } \not\models \beta\}$.

For any sentence φ, its belief degree (also called *rank*) wrt. a given EE base $\langle B, f\rangle$ is defined as:

$$Rank_f(B, \varphi) = \begin{cases} 0 & \text{If } B \not\models \varphi \\ max(B) + 1 & \text{Else if } \models \varphi \\ max(\{m \mid B^m \models \varphi\}) & \text{Otherwise} \end{cases}$$

Note, it is possible that a sentence β in B has a higher belief degree $Rank_f(B, \beta)$ than its evidence degree $f(\beta)$. So that $f(\beta)$ should only be considered as the lower bound of β's belief degree.

Clearly, EE bases are more general forms of prioritized bases, therefore what applies to prioritized bases can also apply to EE bases. Given an EE base $\langle B, f\rangle$, the cut-set of a sentence φ (denoted by $cut_f(\varphi)$) is defined as

$$cut_f(\varphi) = \{\beta \in B \mid Rank_f(B, \varphi) < f(\beta)\} \tag{6}$$

Similarly, an EE base $\langle B, f\rangle$ can induce an EE wrt. $\leq_{Cn(B)}$ by letting

$$\alpha \leq_{Cn(B)} \beta \text{ iff } Rank_f(B, \alpha) \leq Rank_f(B, \beta) \tag{7}$$

Theorem 8. [Wobcke, 1992] *Given an EE base $\langle B, f\rangle$, the induced binary relation $\leq_{Cn(B)}$ defined by (7) is an EE relation wrt. $Cn(B)$.*

An iterated EE base revision operator should map an EE base and the new information to a new EE base. A natural question is of course what should be the evidence (belief) degree of the new sentence φ in the revised EE base? Obviously, if the new information is purely a sentence φ, the revision operator has to assign φ a belief degree via a fixed scheme. It is unlikely that there exists such a fixed scheme suitable for all different kinds of applications. Therefore, based on the same considerations of [Spohn, 1988], we consider a more general revision schema where the new information $\langle\varphi, e\rangle$ consists of a sentence φ and an *evidence degree* $e \in \mathbb{N}^+$. The evidence degree e is supposed to provide additional information for the revision operator to determine the belief degree of φ in the revised EE base.

Given an EE base $\langle B, f\rangle$, the EE base revision is defined as follows for any new information pair $\langle\varphi, e\rangle$:

$$\begin{aligned}\langle B', f'\rangle = \langle B, f\rangle * \langle\varphi, e\rangle \text{ iff } B' &= \{\beta \mid \beta \in cut_f(\neg\varphi)\} \cup \{\varphi \vee \beta \mid \beta \in B\} \cup \{\varphi\} \\ \text{where} f'(\beta) = f(\beta), f'(\beta \vee \varphi) &= max(f(\beta)+1, e) \text{ and } f'(\varphi) = e\end{aligned} \tag{8}$$

According to (8), the evidence degree of a sentence $\beta \in cut_f(\neg\varphi)$ does not change. Unlike in (4), for each sentence $\beta \in B$, a disjunction $\beta\vee\varphi$ with evidence degree $max(f(\beta)+1, e)$ is added to the new EE base. These are the ingredients for avoiding (undesired) implicit dependence. Note, intuitively if $\beta \vee \varphi$ is more entrenched than φ, it will disqualify β to be implicitly dependent on φ.

The EE base revision defined by (8) can be straightforwardly implemented by the following algorithm, where an EE base is represented as: $\{\langle\beta_1, f(\beta_1)\rangle, \ldots, \langle\beta_n, f(\beta_n)\rangle\}$.

Input : $\langle B, f\rangle = \{\langle\beta_1, f(\beta_1)\rangle, \ldots, \langle\beta_n, f(\beta_n)\rangle\}, \varphi, e$
Output: $\langle B', f'\rangle$ such that $\langle B', f'\rangle = \langle B, f\rangle * \langle\varphi, e\rangle$
begin
  $\langle B', f'\rangle = \{\ \}$;
  $\overline{r} = Rank_f(B, \neg\varphi)$;
  for $i = 1 \ldots n$ **do**
    if $f(\beta_i) > \overline{r}$ **then**
      $\langle B', f'\rangle = \langle B', f'\rangle \cup \{\langle\beta_i, f(\beta_i)\rangle\}$;
    end
    $\langle B', f'\rangle = \langle B', f'\rangle \cup \{\langle\beta_i \vee \varphi, max(f(\beta_i)+1, e)\rangle\}$;
  end
  $\langle B', f'\rangle = \langle B', f'\rangle \cup \{\langle\varphi, e\rangle\}$
end

Algorithm 1: Algorithm of the EE base revision

4 Formal Assessment

We present in this section a formal assessment of the EE base revision, including its satisfiability of logical set-theoretical postulates, computational complexity, degrees of syntax relevance and minimal change.

4.1 Logical Set-Theoretical Properties

We first show that the EE base revision satisfies all AGM, DP postulates and additionally (Ind). For this reason, we first present another revision operator, which is equivalent to the EE base revision and satisfies all above the mentioned postulates. The revision operator is a variant of Spohn's proposal of revising *ordinal conditional functions* [Spohn, 1988].

An ordinal conditional function (OCF) is a mapping k from $\mathcal{W}$ to $\mathbb{N}$ (natural numbers including 0). Given a world $w \in \mathcal{W}$, $k(w)$ is called the *rank* of w. Clearly, OCFs are more general forms of faithful rankings. Intuitively, the rank of a world represents its degree of plausibility. The lower its rank, the more plausible is a world. The belief set $Bel(k)$ is the set of sentences which hold in all worlds of rank 0 :

$$Bel(k) = Th(\{w|k(w) = 0\}) \tag{9}$$

From now on, we use an ordinal conditional function and its belief set interchangeably; e.g., $\mu \in k$ means $\mu \in Bel(k)$, and $k_1 \equiv k_2$ denotes $Bel(k_1) = Bel(k_2)$.

Given an OCF k, we denote by $max(k)$ its maximal rank, i.e., $max(k) = max\{k(w) \mid w \in \mathcal{W}\}$. An OCF be extended to a ranking of sentences as follows:

$$k(\mu) = \begin{cases} max(k) + 1 & \text{If } \models \mu \\ min\{k(w) \mid w \models \neg\mu)\} & \text{Otherwise} \end{cases} \tag{10}$$

Put in words, the rank of a sentence is the lowest rank of a world in which the sentence does not hold. Hence, the higher the rank of a sentence, the firmer the belief in it.

As for the EE base revision, the new information consists of a sentence and an evidence degree. An OCF k is revised according to new evidence μ with evidence degree $e \in \mathbb{N}^+$ as follows:

$$k^*_{\varphi,e}(w) = \begin{cases} max(k(w) + 1, e) & \text{If } w \models \neg\varphi \\ k(w) & \text{Else if } k(w) > k(\neg\varphi) \\ 0 & \text{Otherwise} \end{cases} \tag{11}$$

First of all, we show that varying evidence degrees of the new information do not affect the logic contents of revised OCF.

Theorem 9. *Given any OCF k and sentence φ, let $k_1 = K^*_{\varphi,e_1}$ and $k_2 = K^*_{\varphi,e_2}$ where $e_1 \neq e_2$, we have*

$$Bel(k_1) = Bel(k_2)$$

The following results show that the OCF revision defined by (11) satisfies all AGM, DP postulates and (Ind), regardless the evidence degree of the new sentence.

Theorem 10. *For any* $e \in \mathbb{N}^+$ *, the OCF revision operator defined by (11) satisfies all AGM postulates, where* K *and* $K * \varphi$ *are, respectively, identified with* $Bel(k)$ *and* $Bel(k^*_{\varphi,e})$ *.*

Theorem 11. *For arbitrary* $e_1, e_2 \in \mathbb{N}^+$ *, the OCF revision operator defined by (11) satisfies the following conditions:*

(EC1) If $\alpha \models \mu$ *, then* $(k^*_{\mu,e_1})^*_{\alpha,e_2} \equiv k^*_{\alpha,e_2}$ *.*
(EC2) If $\alpha \models \neg\mu$ *, then* $(k^*_{\mu,e_1})^*_{\alpha,e_2} \equiv k^*_{\alpha,e_2}$ *.*
(EInd) If there exists e *such that* $\neg\mu \notin k^*_{\neg\beta,e}$ *, then* $\mu \in (k^*_{\mu,e_1})^*_{\neg\beta,e_2}$

As shown by the following result, the belief degree of any sentence in the revise OCF is uniquely determine by the contents of the original OCF, despite its syntax form (A point we will return to in Section 4.3).

Lemma 1. *Let* k *be an OCF and* $\langle \varphi, e \rangle$ *an arbitrary pair of new information, then the following condition holds for any non-tautologous sentence* β *:*

$$k^*_{\varphi,e}(\beta) = \begin{cases} max(k(\beta)+1, e) & \textit{If } \models \neg\varphi \vee \beta \\ 0 & \textit{Else if } k(\neg\varphi) \geq k(\varphi \rightarrow \beta) \\ k(\beta) & \textit{Else if } k(\varphi \rightarrow \beta) \leq k(\beta) \\ k(\beta)+1 & \textit{Else if } e \leq k(\beta) \\ min(k(\varphi \rightarrow \beta), e) & \textit{Otherwise} \end{cases}$$

Lemma 1 gives a full characterization of the OCF revision. The following result shows the EE base revision satisfies exactly the same conditions of Lemma 1. This essentially means the EE base revision is equivalent to the OCF revision.

Lemma 2. *Given an EE base* B *and a new information pair* $\langle B, \varphi, e \rangle$ *, let* $\langle B_1, f_1 \rangle = \langle B, f \rangle * \langle \varphi, e \rangle$ *, then the following condition holds for any non-tautologous sentence* β *:*

$$Rank_{f_1}(B_1, \beta) = \begin{cases} max(t+1, e) & \textit{If } \models \neg\varphi \vee \beta \\ 0 & \textit{Else If } \overline{r} \geq t' \\ t & \textit{Else if } t' \leq t \\ t+1 & \textit{Else if } e \leq t \\ min(t', e) & \textit{Otherwise} \end{cases}$$

where $\overline{r} = Rank_f(B, \neg\varphi)$ *,* $t = Rank_f(B, \beta)$ *and* $t' = Rank_f(B, \varphi \rightarrow \beta)$ *.*

Theorem 12. *Let* $\langle B, f \rangle$ *be an EE base and* k *an OCF, s.t.,* $Rank_f(B, \beta) = k(\beta)$ *for any sentence* β *. Let* $\langle \varphi, e \rangle$ *be any new information pair. We have for any sentence* α *the following condition holds:*

$$Rank_{f'}(B', \alpha) = k^*_{\varphi,e}(\alpha)$$

where $\langle B', f' \rangle = \langle B, f \rangle * \langle \varphi, e \rangle$

Theorem 13. *The EE base revision operator defined by (8) satisfies all AGM, DP postulates and (Ind).*

Since revision operators defined on the EE bases are more general than belief revision operators defined on belief sets (bases), it is no surprise that we can prove more properties for them. For the EE base revision defined by (8), we can prove that the evidence degree e of the new sentence φ is the lower bound of belief degree of φ in revised EE base.[7] This is a very nice property from pragmatic point of view. Image a scenario where information sources provide new information with evidence degrees to its best knowledge. It is of course proper to take the evidence degree as the lower bound of the information's belief degree

Theorem 14. *For any non-tautologous sentence φ, the following condition holds:*

$$Rank_{f'}(B', \varphi) = max(Rank_f(B, \varphi) + 1, e)$$

where $\langle B', f' \rangle = \langle B, f \rangle * \langle \varphi, e \rangle$.

4.2 Computational Complexity

In this subsection, we will show that the EE base revision is almost optimal in terms of computational complexity. [8] As in [Eiter and Gottlob, 1992], we consider the problem of COUNTERFACTUAL (CF for short), which is to decide whether $B * \varphi \models \beta$ for any belief base B and any sentences φ, β. It is not difficult to see, when a revision operator satisfies (*4) and (*5) of AGM postulates, both SAT and VALID can be polynomially many-to-one reduced to CF. Hence,in general CF is both NP and coNP hard.

Theorem 15. [Nebel, 1992] *For any revision operation, which satisfies (*4) and (*5) ,the problem CF is* NP*-hard and* coNP*-hard.*

A direct consequence of the above theorem is, in general CF is unlikely to be a member of NP $\cup$ coNP, or otherwise it implies NP=coNP. Moreover, Nebel has identified the complexity class of CF wrt. the cut base revision.

Theorem 16. [Nebel, 1994] *For the cut base revision operator defined by (4), the problem CF is $\Delta_2^{\mathrm{p}}[\mathrm{O}(\log \mathrm{n})]$ -complete.*

Since EE base revision is based on cut base revision without introducing additional computational overload, it is not difficult to see, the CF wrt. the EE base revision is in the same complexity class.

Theorem 17. *For EE base revision operator defined by (8), the problem CF is $\Delta_2^{\mathrm{p}}[\mathrm{O}(\log \mathrm{n})]$ -complete.*

[7] There are other possibilities, e.g., the EE base revision defined by [Jin and Thielscher, 2004] adds e to φ's old belief degree.

[8] We assume the reader has basic knowledge on complexity theory, or otherwise can be found in [Papadimitriou, 1994]

Together with NP-hardness and coNP-hardness of CF, the above result suggests that the complexity of the EE base revision is almost optimal and there is no space for fundamental improvement.

Since the problem of computing the revised EE base is obviously not hard than CF, it is in the same complexity class.

Theorem 18. *The problem of computing the revised EE base as defined by (8) is* $\Delta_2^{\mathrm{p}}[\mathrm{O}(\log \mathrm{n})]$ *-complete.*

4.3 Degree of Syntax Relevance

Technically, EE base revision violates Dalal's principle of *Irrelevance of Syntax* [Dalal, 1988b], in the sense the new EE base is not determined merely by the logical contents of the original EE base. However, we will show its degree of syntax relevance is so low, that it does not even deserve to be called syntax relevant.

Two EE bases $\langle B_1, f_1 \rangle$ and $\langle B_2, f_2 \rangle$ are called *epistemically equivalent* iff their induced EEs (as defined by (7)) are equivalent, i.e., for any sentences α, β :

$$\alpha \leq_{Cn(B_1)} \beta \text{ iff } \alpha \leq_{Cn(B_2)} \beta$$

Two EE bases $\langle B_1, f_1 \rangle$ and $\langle B_2, f_2 \rangle$ are called *equivalent* iff the following condition holds for any sentence β :

$$Rank_{f_1}(B_1, \beta) = Rank_{f_2}(B_2, \beta)$$

The following result shows in the EE base revision the logical contents of revised EE base is determined by the induced EE of the original EE base.

Theorem 19. *Let* $\langle B_1, f_1 \rangle$, $\langle B_2, f_2 \rangle$ *be two epistemically equivalent EE bases, then for any sentence* φ *with evidence degrees* e_1, e_2 *:*

$$Cn(B_1') = Cn(B_2')$$

where $\langle B_1', f_1' \rangle = \langle B_1, f_1 \rangle * \langle \varphi, e_1 \rangle$ *and* $\langle B_2', f_2' \rangle = \langle B_2, f_2 \rangle * \langle \varphi, e_2 \rangle$

Finally, it is not difficult to see that in the EE base revision equivalent original EE bases (despite their syntax forms) lead to equivalent new EE bases.

Theorem 20. *Let* $\langle B_1, f_1 \rangle$, $\langle B_2, f_2 \rangle$ *be two equivalent EE bases, then for any sentence* φ *with evidence degree* e *:* $\langle B_1', f_1' \rangle = \langle B_1, f_1 \rangle * \langle \varphi, e \rangle$ *is equivalent to* $\langle B_2', f_2' \rangle = \langle B_2, f_2 \rangle * \langle \varphi, e \rangle$

4.4 Degree of Minimal Change

Minimal change is both underlying principle of AGM theory (wrt. logical contents) and iterated revision (wrt. preference information). The EE base revision's satisfiability of AGM postulates is an evidence of its minimal change in logical contents. Obviously, the change of preference information in EE base revision is not absolutely minimal. This is not at all a drawback, since as shown in [Darwiche and Pearl, 1997; Jin and Thielscher, 2005], the absolute minimal change of preference information implies radical behaviors and is not proper in general.

In the sequel, we will show that the change of preference information in the EE base revision is all necessary in order to satisfy AGM, DP postulates and (Ind). Since the EE base revision is shown to be equivalent to the OCF revision, we will base our analysis on a special case of the OCF revision, where the evidence degree e of the new sentence is fixed to be 1. This OCF revision can be visualized by the Fig 4.4, where the dots on the left (right) side of the vertical dotted line represent worlds in which φ is true (false); the red arrowed lines denote the way how ranks of worlds change.

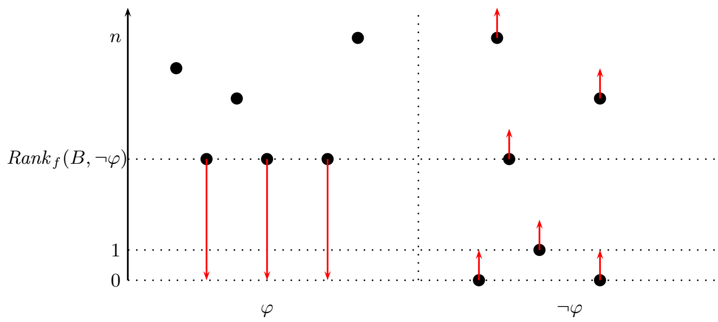

Fig. 1. Visualization of an OCF revised by $\langle \varphi, 1 \rangle$

According to Theorem 2, if a revision operator satisfies AGM postulates, the new belief set (base) should take the most plausible worlds satisfying φ as it most plausible worlds. Hence, the downward movement of worlds on the left side, and upward movement of worlds with rank 0 on the right side is compulsory. According to Theorem 6, if a revision operator satisfies DP postulates, the worlds the left (right) side should have somehow uniformed way of change their ranks, such that their relative positions do not change. As a direct consequence, the upward movement of worlds with rank 0 on the right side should squeeze the rest worlds on the right side upward (image domino effect). Altogether, the OCF revision does not do any change which is unnecessary. If fact, as a side effect of the way OCF revision changes ranks of worlds on the right side, the Condition (IndR) of Theorem 7 is automatically satisfied.

The general OCF (EE base) revision allows the evidence degree of the new sentence to vary, this may introduce more changes to the preference information, but it also has its merit in practice as discussed above. In this sense, we agree on on those authors [Hansson, 2003], who argued that we should not take minimal change as the only criterion.

5 Related Work and Conclusions

The EE base revision proposed is technically syntax based, as argued it is however from pragmatic point of view more fine-grained. There are revisions, such as model-based revision [Katsuno and Mendelzon, 1991; Dalal, 1988a], do no need extra-logical preference information. Since such revision operators are always uniquely determined by logical contents of the belief set (also called *syntax irrelevant* [Dalal, 1988b]), they are quite *inflexible* [Nebel, 1998], in the sense the class of generated revision operations is quite limited. On the other hand, almost all syntax irrelevant revision operators violate both DP postulates and (Ind).

Syntax irrelevant revision operators are in principle iterated operators. As already shown in [Eiter and Gottlob, 1992; Nebel, 1998], CF wrt. most of them are Π_2^p-complete, except Dalal's operator and full meet revision. Moreover, Liberatore [1997] has shown that CF wrt. most well-know syntax based iterated revisions are Δ_2^p-complete. Due to the limited space, we only show in Table 5 a comparison between revision operators, which satisfy all AGM postulates [9].

Revision operator	DP	Ind	Complexity of CF
EE base	Yes	Yes	$\Delta_2^p[\log n]$-complete
Natural	Yes	No	Δ_2^p-complete
Prioritized full meet	No	Yes	Δ_2^p-complete
Adjustment	No	No	Δ_2^p-complete
Conditionalization	No	No	Δ_2^p-complete
Dalal	No	No	$\Delta_2^p[\log n]$-complete
Full meet	No	No	coNP(3)-complete

Table 1. Comparison of well-known iterated revision operators

The first three columns show that the EE base revision and prioritized full meet revision are the only operators which satisfy all DP postulates and (Ind). It it worth to mention, that the prioritized full meet revision [Liberatore, 1997] is essentially same as the lexicographic revision [Abhaya, 1994] with "naked evidence"and the basic memory operator [Konieczny and Pérez, 2000]. Hence its satisfies the following postulate of *Recalcitrance* which is criticized as too radical in [Jin and Thielscher, 2005]:

(Rec) If $\not\models \beta \rightarrow \neg\mu$, then $(\Psi * \mu) * \beta \models \mu$.

The last column shows the complexity results, most of which can be found in [Eiter and Gottlob, 1992; Nebel, 1998; Liberatore, 1997]. As shown, although the EE base revision's computational complexity is at same level as other syntax

[9] We only consider special cases of Adjustment and Conditionalization where the new evidence degree $e > 0$.

based revisions, it requires much less calls (specially when size of EE base is large) of NP oracles. Since in practice, NP oracles cost exponentially time (provided $P \neq NP$), less invocations of NP oracles is already a big advantage. For general propositional logic, Dalal's operator is in the same complexity class as the EE base revision. A drawback of Dalal's operator is even the underlying language is constrained to *Horn sentences*, it remains in the same complexity class, while CF wrt. the EE base revision becomes *tractable*. The only revision operator whose complexity class is lower than the EE base revision is the full meet revision, which is too radical, in the sense a severe revision with full meet revision only takes the new information as the revised belief base.

To summarize, we have proposed a general iterated base revision operator, which satisfies all AGM, DP postulates and (Ind), in addition to other nice properties from the practical point of view. The proposed iterated is almost syntax irrelevant and does only changes of preference information, which is necessary. A comparison to others shows it is also well-performed in computation.

Appendix: proofs

Theorem 9. *Given any OCF* k *and sentence* φ, *let* $k_1 = K^*_{\varphi,e_1}$ *and* $k_2 = K^*_{\varphi,e_2}$ *where* $e_1 \neq e_2$, *we have*

$$Bel(k_1) = Bel(k_2)$$

Proof. According to (11), the set of most plausible worlds (with rank 0) in the revised OCF is not influenced by the value of evidence degree of the new information. Furthermore, $Bel(k)$ of an OCF k as defined by (9) only depends on the set of most plausible world.

Theorem 10. *For any* $e \in \mathbb{N}^+$, *the OCF revision operator defined by (11) satisfies all AGM postulates, where* K *and* $K * \varphi$ *are, respectively, identified with* $Bel(k)$ *and* $Bel(k^*_{\varphi,e})$.

Proof. Let $w_1 \leq_k w_2$ iff $k(w_1) \leq k(w_2)$. Clearly, $\leq_k$ is a faithful ranking wrt. $Bel(k)$. According to (11), the set of most plausible worlds (with rank 0) in the revised OCF is exactly $\min([\varphi], \leq_k)$. It follows from Theorem 2, that EE base revision satisfies all AGM postulates.

Theorem 11. *For arbitrary* $e_1, e_2 \in \mathbb{N}^+$, *the OCF revision operator defined by (11) satisfies the following conditions:*

(EC1) If $\alpha \models \mu$, *then* $(k^*_{\mu,e_1})^*_{\alpha,e_2} \equiv k^*_{\alpha,e_2}$.
(EC2) If $\alpha \models \neg\mu$, *then* $(k^*_{\mu,e_1})^*_{\alpha,e_2} \equiv k^*_{\alpha,e_2}$.
(EInd) If there exists e *such that* $\neg\mu \notin k^*_{\neg\beta,e}$, *then* $\mu \in (k^*_{\mu,e_1})^*_{\neg\beta,e_2}$

Proof. It follows directly from Theorem 6, Theorem 7, that Conditions (EC1), (EC2) and (EInd) hold for OCF revision with fixed evidence degree for the new information. From Theorem 9, then it follows that those conditions also holds for varying evidence degrees.

Lemma 1. *Let* k *be an OCF and* $\langle\varphi, e\rangle$ *an arbitrary pair of new information, then the following condition holds for any non-tautologous sentence* β:

$$k^*_{\varphi,e}(\beta) = \begin{cases} max(k(\beta)+1, e) & \textit{If } \models \neg\varphi \vee \beta \\ 0 & \textit{Else if } k(\neg\varphi) \geq k(\varphi \to \beta) \\ k(\beta) & \textit{Else if } k(\varphi \to \beta) \leq k(\beta) \\ k(\beta)+1 & \textit{Else if } e \leq k(\beta) \\ min(k(\varphi \to \beta), e) & \textit{Otherwise} \end{cases}$$

Proof. Assume $\models \neg\varphi \vee \beta$. Since β is non-tautologous, if follows that there exist $w \models \neg\beta$, s.t., $k(w) = k(\beta)$. Follows from $\models \neg\varphi \vee \beta$, for any w' if $w' \models \neg\beta$ then $w' \models \neg\varphi$. According to (11), $k^*_{\varphi,e}(w) = max(k(w)+1, e)$. Let w' be an arbitrary world, s.t., $w' \models \neg\beta$. It follows from $k(w) = k(\beta)$, $k(w') \geq k(w)$. Hence, we have $k^*_{\varphi,e}(w') \geq k^*_{\varphi,e}(w)$. It follows from $k^*_{\varphi,e}(w) = max(k(w)+1, e)$, we have $k^*_{\varphi,e}(\beta) = max(k(\beta)+1, e)$. In the rest of the proof, we only consider the case $\not\models \neg\varphi$ and $\not\models \beta$.

Assume $k(\neg\varphi) \geq k(\varphi \to \beta)$. It follows from (EE2), $k(\neg\varphi) = k(\varphi \to \beta)$. There exists $w \models \varphi \wedge \neg\beta$, s.t., $k(w) = k(\neg\varphi) = k(\varphi \to \beta)$. According to (11), $k^*_{\varphi,e}(w) = 0$. Hence, we have $k^*_{\varphi,e}(\beta) = 0$.

Assume $k(\neg\varphi) < k(\varphi \to \beta)$ and $k(\varphi \to \beta) \leq k(\beta)$. It follows from (EE2), $k(\varphi \to \beta) = k(\beta)$. There exists $w \models \varphi \wedge \neg\beta$, s.t., $k(w) = k(\varphi \to \beta) = k(\beta)$. According to (11), $k^*_{\varphi,e}(w) = k(w)$, since $k(w) > k(\neg\varphi)$. To prove $k^*_{\varphi,e}(\beta) = k(w)$, we need to show, if $w' \models \neg\beta$ then $k^*_{\varphi,e}(w') \geq k(w)$ for any w' . Let w' be an arbitrary world, s.t., $w' \models \neg\beta$. It follows from $k(w) = k(\beta)$, $k(w') \geq k(w)$. We consider two cases: **1)** If $w' \models \neg\varphi$, then according to (11) $k^*_{\varphi,e}(w') = max(k(w') + 1, e) \geq k(w)$. **2)** If $w' \models \varphi$, then according to (11) $k^*_{\varphi,e}(w') = k(w') \geq k(w)$.

Assume $k(\neg\varphi) < k(\varphi \to \beta)$, $k(\varphi \to \beta) > k(\beta)$ and $e \leq k(\beta)$. There exists $w \models \neg\beta$, s.t., $k(w) = k(\beta)$. It follows from $k(\varphi \to \beta) > k(\beta)$, $w \models \neg\varphi$. According to (11), $k^*_{\varphi,e}(w) = max(k(w) + 1, e)$. Since $e \leq k(w)$, we have $k^*_{\varphi,e}(w) = k(w) + 1$. To prove $k^*_{\varphi,e}(\beta) = k(w) + 1$, we need to show, if $w' \models \neg\beta$ then $k^*_{\varphi,e}(w') \geq k(w) + 1$ for any w' . Let w' be an arbitrary world, s.t., $w' \models \neg\beta$. It follows from $k(w) = k(\beta)$, $k(w') \geq k(w)$. We consider again two cases: **1)** Assume $k(w') \geq k(\varphi \to \beta)$. Since $k(\varphi \to \beta) > k(\neg\varphi)$, according to (11) $k^*_{\varphi,e}(w')$ is either $max(k(w')+1, e)$ or $k(w')$. Clearly, $k^*_{\varphi,e}(w') \geq k(w)+1$, since $k(w') \geq k(\varphi \to \beta) > k(w)$. **2)** If $k(w') < k(\varphi \to \beta)$, then $w' \models \varphi \to \beta$. Since $w' \models \neg\beta$, we have $w' \models \neg\varphi$. According to (11) , $k^*_{\varphi,e}(w') = max(k(w')+1, e)$. Since $k(w') \geq k(w)$, $k^*_{\varphi,e}(w') \geq k(w) + 1$.

Assume $k(\neg\varphi) < k(\varphi \to \beta)$, $k(\varphi \to \beta) > k(\beta)$ and $e > k(\beta)$. we distinguish two cases: **1)** Assume $e \leq k(\varphi \to \beta)$. There exists $w \models \neg\beta$, s.t., $k(w) = k(\beta)$. It easy to see $w \models \neg\varphi$, since $k(\varphi \to \beta) > k(\beta)$. According to (11), $k^*_{\varphi,e}(w) = max(k(w) + 1, e)$. It follows from $e > k(\beta)$, $k^*_{\varphi,e}(w) = e$. We will show, that if $w' \models \neg\beta$ then $k^*_{\varphi,e}(w') \geq e$ for any w' . Let w' be an arbitrary world, s.t., $w' \models \neg\beta$. It follows from $k(w) = k(\beta)$, $k(w') \geq k(w)$. There are two sub-cases: a) If $k(w') < k(\varphi \to \beta)$, then $w' \models \varphi \to \beta$. It follows from $w' \models \neg\beta$, $w' \models \neg\varphi$. According to (11) $k^*_{\varphi,e}(w') = max(k(w') + 1, e) \geq e$. b) If $k(w') \geq k(\varphi \to \beta)$, then $k^*_{\varphi,e}(w')$ is either $k(w')$ or $max(k(w')+1, e)$. Since $k(w') \geq k(\varphi \to \beta) \geq e$, we have $k^*_{\varphi,e}(w') \geq e$. **2)** Assume $e > k(\varphi \to \beta)$. There exists $w \models \varphi \wedge \neg\beta$, s.t., $k(w) = k(\varphi \to \beta)$. According to (11) $k^*_{\varphi,e}(w) = k(w)$, since $k(w) > k(\neg\varphi)$. We will show, that if $w' \models \neg\beta$ then $k^*_{\varphi,e}(w') \geq k(\varphi \to \beta)$ for any w' . Let w' be an arbitrary world, s.t., $w' \models \neg\beta$. It follows from $k(w) = k(\beta)$, $k(w') \geq k(w)$. There are two sub-cases: a) If $k(w') \geq k(\varphi \to \beta)$, then $k^*_{\varphi,e}(w')$ is either $k(w')$ or $max(k(w') + 1, e)$, since $k(\varphi \to \beta) > k(\neg\varphi)$. Hence $k^*_{\varphi,e}(w') \geq k(\varphi \to \beta)$. b) If $k(w') < k(\varphi \to \beta)$, then $w' \models \varphi \to \beta$. It follows from $w' \models \neg\beta$, that $w' \models \neg\varphi$. According to (11) $k^*_{\varphi,e}(w') = max(k(w') + 1, e) \geq e > k(\varphi \to \beta)$.

Lemma 2. *Given an EE base B and a new information pair $\langle B, \varphi, e\rangle$, let $\langle B_1, f_1\rangle = \langle B, f\rangle * \langle\varphi, e\rangle$, then the following condition holds for any non-tautologous* *sentence* β *:*

$$Rank_{f_1}(B_1,\beta) = \begin{cases} max(t+1,e) & \textit{If } \models \neg\varphi \vee \beta \\ 0 & \textit{Else If } \overline{r} \geq t' \\ t & \textit{Else if } t' \leq t \\ t+1 & \textit{Else if } e \leq t \\ min(t',e) & \textit{Otherwise} \end{cases}$$

where $\overline{r} = Rank_f(B,\neg\varphi)$ *,* $t = Rank_f(B,\beta)$ *and* $t' = Rank_f(B,\varphi\rightarrow\beta)$ *.*

Proof. Assume $\models \neg\varphi\vee\beta$. Since β is non-tautologous, it follows that $\models \neg\varphi$. From $\models\neg\varphi$, it follows that $cut_f(\neg\varphi)=\emptyset$. We consider two cases: **1)** Assume $e > t+1$. According to (8), $\varphi \in B_1^e$. It follows from that $\models\neg\varphi$, we have $B_1^e \models \beta$. It is easy too see, according to (8), $B_1^{e+1} = \{\varphi\vee\psi \,|\, \psi\in B^e\}$ which is logically equivalent to B^e, since $\models\neg\varphi$. It follows from $e>t+1$, $B^e \not\models \beta$. Thus, we have $Rank_{f_1}(B_1,\beta) = e$. **2)** Assume $e \leq t+1$. According to (8), $\varphi\vee\beta \in B_1^{t+1}$. It follows from that $\models\neg\varphi$, we have $B_1^{t+1}\models\beta$. According to (8), $B_1^{t+2} = \{\varphi\vee\psi\,|\,\psi\in B^{t+1}\}$, which is logically equivalent to B^{t+1}, since $\models\neg\varphi$. Since $B^{t+1}\not\models\beta$, it follows that $Rank_{f_1}(B_1,\beta)=t+1$.

For the rest of the proof, we assume $\not\models\neg\varphi$ and $\not\models\beta$.

Assume $\overline{r}\geq t'$. According to (8), $B_1^0 = B^{\overline{r}+1}\cup\{\varphi\}\cup\{\varphi\vee\psi\,|\,\psi\in B\}$. It follows from $t' = Rank_f(B,\varphi\rightarrow\beta)$, that $B^{\overline{r}+1}\not\models\varphi\rightarrow\beta$. According to the Deduction Theorem, we have $B_1^0\not\models\beta$. Hence $Rank_{f_1}(B_1,\beta)=0$.

Assume $\overline{r}<t'$ and $t'\leq t$. It follows from (EE2), we have $t=t'$. According to (8), $B^t\subseteq B_1^t$, since $\overline{r}<t'$. From $B^t\models\beta$, it follows that $B_1^t\models\beta$. It is easy to see, $B_1^{t+1}\subseteq B^{t+1}\cup\{\varphi\}\cup\{\varphi\vee\psi\,|\,\psi\in B\}$. It follows from $B^{t+1}\not\models\varphi\rightarrow\beta$, we have $B_1^{t+1}\not\models\beta$. Hence $Rank_{f_1}(B_1,\beta)=t$.

Assume $\overline{r}<t'$ and $t'>t$ and $e\leq t$. We consider two cases: **1)** Assume $t\geq\overline{r}$. According to (8), $B_1^{t+1} = B^{t+1}\cup\{\varphi\vee\psi\,|\,\psi\in B^t\}$, since $e\leq t$. It follows from $B^t\models\beta$, we have $B^{t+1}\models\bigwedge(B^t\backslash B^{t+1})\rightarrow\beta$. It is easy to see that $\{\varphi\vee\psi\,|\,\psi\in B^t\}\models\bigwedge(B^t\backslash B^{t+1})\vee\varphi$. Since $t'>t$, we have $B^{t+1}\models\varphi\rightarrow\beta$. Hence, we have $B_1^{t+1}\models\beta$. According to (8), $B_1^{t+2} = B^{t+2}\cup\{\varphi\vee\psi\,|\,\psi\in B^{t+1}\}$. Assume $B^{t+2}\cup\{\varphi\vee\psi\,|\,\psi\in B^{t+1}\}\models\beta$. It is easy to see that $B^{t+2}\cup\{\varphi\vee\psi\,|\,\psi\in B^{t+1}\}\models\beta$ is logically equivalent to $B^{t+2}\cup\{\varphi\vee\bigwedge(B^{t+1}\backslash B^{t+2})\}$. Thus $B^{t+2}\models(\varphi\vee\bigwedge(B^{t+1}\backslash B^{t+2}))\rightarrow\beta$. Since $B^{t+2}\not\models\beta$, we have $B^{t+2}\models\neg\varphi\wedge\neg\bigwedge(B^{t+1}\backslash B^{t+2})\}$, which contradicts with $t>\overline{r}$. **2)** Assume $t<\overline{r}$. According to (8) $B_1^{t+1} = B^{\overline{r}+1}\cup\{\varphi\vee\psi\,|\,\psi\in B^t\}$, since $e\leq t$. Since $B^t\models\beta$, we have $\{\varphi\vee\psi\,|\,\psi\in B^t\}\models\varphi\vee\beta$. It follows from $\overline{r}<t'$, that $B^{\overline{r}+1}\models\varphi\rightarrow\beta$. Thus, we obtain $B_1^{t+1}\models\beta$. $B_1^{t+2} = B^{\overline{r}+2}\cup\{\varphi\vee\psi\,|\,\psi\in B^{t+1}\}$ .It is easy to see that $B^{t+1}\models\{\varphi\vee\psi\,|\,\psi\in B^{t+1}\}$. Since $t<\overline{r}$, we have $B_1^{\overline{r}+2}\subseteq B^{t+1}$. It follows from $B^{t+1}\not\models\beta$, we have $B_1^{t+2}\not\models\beta$.

Assume $\overline{r}<t'$ and $t'>t$ and $e>t$. We consider two cases: **1)** Assume $t'\geq e$. According to (8), $B^{t'}\cup\{\varphi\}\subseteq B_1^e$, since $\overline{r}<t'$. It follows from that $B^{t'}\models\varphi\rightarrow\beta$, we have $B_1^e\models\beta$. According to (8), $B_1^{e+1}\subseteq B^{e+1}\cup\{\varphi\vee\psi\,|\,\psi\in B^e\}$. It follows from $e>t$, that $B^e\not\models\beta$. Hence, we have $B_1^{e+1}\not\models\beta$. **2)** Assume $t'<e$. According to (8), $B^{t'}\cup\{\varphi\}\subseteq B_1^{t'}$, since $\overline{r}<t'$. It follows from that $B^{t'}\models\varphi\rightarrow\beta$, we have $B_1^{t'}\models\beta$. According to (8), $B_1^{t'+1}\subseteq B^{t'+1}\cup\{\varphi\vee\psi\,|\,\psi\in$

$B^e\}$. Since $e > t'$, we have $B^{t'+1} \cup \{\varphi \vee \psi \mid \psi \in B^e\}$ is logically equivalent to $B^{t'+1}$. It follows from $B^{t'+1} \not\models \varphi \rightarrow \beta$, $B^{t'+1} \not\models \beta$ Thus, we obtain $B_1^{t'+1} \not\models \beta$.

Theorem 12. *Let $\langle B, f\rangle$ be a EE base and k a OCF, s.t., $Rank_f(B, \beta) = k(\beta)$ for any sentence β. Let $\langle \varphi, e\rangle$ be any new information pair. We have for any sentence α the following condition holds:*

$$Rank_{f'}(B', \alpha) = k^*_{\varphi,e}(\alpha)$$

where $\langle B', f'\rangle = \langle B, f\rangle * \langle \varphi, e\rangle$

Proof. As a direct consequence of Lemma 1 and 2, the condition holds for non-tautologous sentences. Since the ranks of tautologous sentences is the maximal rank of non-tautologous sentences plus one, the condition also holds for tautologous sentences.

Theorem 13. *The EE base revision operator defined by (8) satisfies all AGM, DP postulates and (Ind).*

Proof. If follows directly from Lemma 1, Lemma 2, Theorem 11 and Theorem 10.

Theorem 14. *For any non-tautologous sentence φ, the following condition holds:*

$$Rank_{f'}(B', \varphi) = max(Rank_f(B, \varphi) + 1, e)$$

where $\langle B', f'\rangle = \langle B, f\rangle * \langle \varphi, e\rangle$.

Proof. If follows directly from Lemma Lemma 2.

Theorem 17. *For EE base revision operator defined by (8), the problem CF is $\Delta_2^p[O(\log n)]$-complete.*

Proof. To prove that CF is in $\Delta_2^p[O(\log n)]$, it is enough to see that the main computational affect of Algorithm 1 is to compute $Rank(B, \neg\varphi)$. Clearly, computing $Rank(B, \neg\varphi)$ requires in worst-case to call $log\, n$ (based on the ideas of the Binary Search) times a NP oracle (which solves the implication (IMPL) problem). After computed the revised EE base B', we just can call one more time IMPL oracle to decide whether $B' \models \beta$.

What remains to show is CF's $\Delta_2^p[O(\log n)]$-hardness. We prove this by showing a polynomially many-to-one reduction from CF wrt. cut base revision to CF wrt. EE base revision. Given any prioritized base $\langle B, \leq_B\rangle$, we can construct a EE base $\langle B, f\rangle$ as follows. The all sentences in the lowest class get as evidence degrees 1, the sentences in the next higher class get evidence degree 2, and so on unit all sentences in B gets an evidence degree. Obviously, EEs induced by $\langle B, \leq_N\rangle$ and $\langle B, f\rangle$ are equivalent. As a consequence of Theorem 19, for any sentences φ, ψ and we have

$$B *_{\leq_B} \varphi \models \psi \text{ iff } B_1 \models \psi$$

where $\langle B_1, f_1\rangle = \langle B, f\rangle * \langle \varphi, 1\rangle$

Theorem 18. *The problem of computing the revised EE base as defined by (8) is* $\Delta_2^{\mathrm{p}}[\mathrm{O}(\log \mathrm{n})]$ *-complete.*

Proof. Base on same ideas of the proof of Theorem 17, and observation that CF can be solved by first computing the revised EE base, then calling a NP oracle.

Theorem 19. Let $\langle B_1, f_1 \rangle$, $\langle B_2, f_2 \rangle$ be two epistemically equivalent EE bases, then for any sentence φ with evidence degrees e_1, e_2 :

$$Cn(B_1') = Cn(B_2')$$

where $\langle B_1', f_1' \rangle = \langle B_1, f_1 \rangle * \langle \varphi, e_1 \rangle$ and $\langle B_2', f_2' \rangle = \langle B_2, f_2 \rangle * \langle \varphi, e_2 \rangle$

Proof. According to Lemma 2, when a EE base $\langle B, f \rangle$ revised with φ with some arbitrary evidence degree e , we have a sentence β 's rank is e in the revised EE base (which means it is not in the revised EE base) iff $\varphi \to \beta \leq_{Cn(B)} \neg\varphi$.

Theorem 20. *Let* $\langle B_1, f_1 \rangle$ *,* $\langle B_2, f_2 \rangle$ *be two equivalent EE bases, then for any sentence* φ *with evidence degree* e *:* $\langle B_1', f_1' \rangle = \langle B_1, f_1 \rangle * \langle \varphi, e \rangle$ *is equivalent to* $\langle B_2', f_2' \rangle = \langle B_2, f_2 \rangle * \langle \varphi, e \rangle$

Proof. It follows from the Lemma 2, all non-tautologous sentences will have the same rank in both revised EE bases. Since ranks of tautologous sentences is the maximal rank of the non-tautologous sentences plus one, they are also equal.

References

[Abhaya, 1994] N. Abhaya. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 4:353–390, 1994.

[Alchourrón *et al.*, 1985] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic 50(2)*, pages 510–530, 1985.

[Boutilier, 1993] C. Boutilier. Revision sequences and nested conditionals. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 519–525, 1993.

[Dalal, 1988a] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In Paul Rosenbloom and Peter Szolovits, editors, *Proceedings of the Seventh National Conference on Artificial Intelligence*, volume 2, pages 475–479, Menlo Park, California, 1988. AAAI Press.

[Dalal, 1988b] M. Dalal. Updates in propositional databases. Technical Report DCS-TR 222, Rutgers University, Dept. of Computer Science, Feb 1988.

[Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence 89*, 1–29, 1997.

[Eiter and Gottlob, 1992] Thomas Eiter and Georg Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. pages 261–273, 1992.

[Gärdenfors and Makinson, 1988] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of TARK*, pages 83–95, Asilomar, CA, 1988.

[Gärdenfors, 1988] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.

[Hansson, 1998] S. O. Hansson. Revision of belief sets and belief bases. In Didier Dubois and Henri Prade, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 3: Belief Change*, pages 17–75. Kluwer Academic Publishers, Dordrecht, 1998.

[Hansson, 2003] S. O. Hansson. Ten philosophical problems in belief revision. *Journal of Logic Computation*, 13(1):37–49, 2003.

[Jin and Thielscher, 2004] Y. Jin and M. Thielscher. Representing beliefs in the fluent calculus. In R. L. de Mántras and L. Saitta, editors, *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pages 823–827, Valencia, Spain, August 2004. IOS Press.

[Jin and Thielscher, 2005] Y. Jin and M. Thielscher. Iterated belief revision, revised. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 478–483, Edinburgh, Scotland, August 2005.

[Katsuno and Mendelzon, 1991] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

[Konieczny and Pérez, 2000] S. Konieczny and R. Pérez. A framework for iterated revision. *Journal of Applied Non-Classical Logics*, 10(3-4), 2000.

[Lehmann, 1995] D. J. Lehmann. Belief revision, revised. In C. S. Mellish, editor, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1534–1540, 1995.

[Liberatore, 1997] Paolo Liberatore. The complexity of iterated belief revision. In Foto N. Afrati and Phokion G. Kolaitis, editors, *Database Theory - ICDT '97, 6th International Conference, Delphi, Greece, January 8-10, 1997, Proceedings*, volume 1186 of *Lecture Notes in Computer Science*, pages 276–290. Springer, 1997.

[Nayak *et al.*, 1996] A. C. Nayak, N. Y. Foo, M. Pagnucco, and A. Sattar. Changing conditional belief unconditionally. In *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, pages 119–135. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.

[Nayak *et al.*, 2003] A. C. Nayak, M. Pagnucco, and P. Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003.

[Nebel, 1991] B. Nebel. Belief revision and default reasoning: Syntax-based approaches. In J. A. Allen, R. Fikes, and E. Sandewall, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 2nd International Conference*, pages 417–428, San Mateo, 1991. Morgan Kaufmann.

[Nebel, 1992] Bernhard Nebel. Syntax-based approaches to belief revision. In P. Gärdenfors, editor, *Belief Revision*, volume 29, pages 52–88. Cambridge University Press, Cambridge, UK, 1992.

[Nebel, 1994] B. Nebel. Base revision operations and schemes: Semantics, representation and complexity. In *ECAI*, pages 341–345, 1994.

[Nebel, 1998] B. Nebel. How hard is it to revise a belief base? In Didier Dubois and Henri Prade, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 3: Belief Change*, pages 77–145. Kluwer Academic Publishers, Dordrecht, 1998.

[Papadimitriou, 1994] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[Rott, 1991] H. Rott. Two methods of constructing contractions and revisions of knowledge systems, 1991.

[Spohn, 1988] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic state. In W. L. Harper and B. Skyrms, editors, *Causation in Decision: Belief, Change and Statistics: Proceedings of the Irvine Conference on Probability and Causation*, volume II, pages 105–134, Dordrecht, 1988. Kluwer Academic Publisher.

[Williams, 1994] M-A. Williams. On the logic of theory base change. In C. MacNish, D. Pearce, and L.M. Pereira, editors, *Logics in Artificial Intelligence*, volume LNCS No 835, pages 86–105. Springer Verlag, 1994.

[Wobcke, 1992] W. R. Wobcke. On the use of epistemic entrenchment in nonmonotonic reasoning. In *Proceedings of the Tenth European Conference on Artificial Intelligence*, pages 324–328, 1992.