Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence

A Constructive Approach to Independent and Evidence Retaining Belief Revision by General Information Sets

Gabriele Kern-Isberner and **Patrick Krümpelmann**
Technische Universität Dortmund
Dortmund, Germany
{gabriele.kern-isberner, patrick.kruempelmann}@tu-dortmund.de

Abstract

Recent years have seen a lot of work towards extending the established AGM belief revision theory with respect to iterating revision, preserving conditional beliefs, and handling sets of propositions as new information. In particular, novel postulates like *independence* and *evidence retainment* have been brought forth as new standards for revising epistemic states by (sets of) propositional information. In this paper, we propose a constructive approach for revising epistemic states by sets of (propositional and conditional) beliefs that combines ideas from nonmonotonic reasoning with conditional belief revision. We also propose a novel principle called *enforcement* that covers both *independence* and *evidence retainment*, and we show our revision operator to comply with major postulates from the literature. Moreover, we point out the relevance of our approach for default reasoning.

1 Introduction

Belief revision theory provides methods and axiomatic postulates for changing the minds of intelligent agents in a rational way, in particular, when the new information is in conflict with previously held beliefs. The seminal AGM theory [Alchourrón *et al.*, 1985] studied belief revision in a purely propositional scenario. Since then, various approaches have been proposed to extend this framework, going for iterated revision of more complex epistemic states and studying revision strategies given by conditionals (cf., e.g., [Darwiche and Pearl, 1997]). Recently, in particular the demand for independence of syntax which is responsible for considering information sets like $\{A, B\}$ and $\{A \wedge B\}$ as equivalent has been questioned since in the context of iterated revision, unwanted effects may occur (cf., e.g., [Delgrande and Jin, 2008]). This has raised interest in investigating multiple [Zhang, 2004], parallel [Delgrande and Jin, 2008] or simultaneous [Kern-Isberner, 2008] revision in which the new information is given by sets of propositions resp. conditionals. However, no consensus has been reached yet, and each new approach seems to set up a new framework.

This paper aims at providing some common ground for the theory of revising epistemic states by general information sets, i.e., sets of (conditional or propositional) sentences. So, given some prior epistemic state and some general information set, we will propose postulates that are mainly based on a principle of conditional preservation as stated in [Kern-Isberner, 2004], and on enforcing conditional beliefs in the posterior epistemic state. As a consequence of the principle of conditional preservation, each element of the information set is dealt with separately, thus distinguishing easily revisions by, e.g., $\{A, B\}$ from those by $\{A \wedge B\}$. The postulate of enforcing conditional beliefs exploits ideas from [Kern-Isberner, 2002] originally proposed for default reasoning and generalizes postulates like *independence* [Jin and Thielscher, 2007] and *evidence retainment* [Delgrande and Jin, 2008] for the task of revising by conditionals. Both principles make use of the theory of conditional structures that has been proposed as a fundamental algebraic approach to guide belief revision and default reasoning [Kern-Isberner, 2004; 2001]. Despite its complex theoretical underpinning, the simple rationale behind this theory is to make interactions between different conditionals (and propositions) visible on possible worlds, and manageable for information processing. Moreover, conditional structures provide schemata for revisions from which a constructive method for revising epistemic states by general information sets is obtained which satisfies the above mentioned principle of conditional preservation, the novel enforcement postulate, and success. We will show that this is enough to cover all major postulates from related works [Darwiche and Pearl, 1997; Jin and Thielscher, 2007; Delgrande and Jin, 2008] thus achieving a concise but, at the same time, broad and general picture.

Taking also conditionals as inputs to the revision procedure into regard allows for a more general and realistic view on belief revision. Propositions can be taken as a special case of conditionals, so methods for revising by conditionals can also be used for revising by plausible propositions. Moreover, agents often change contexts in their environment which may change their conditional beliefs. For instance, when passing the border of a country we will usually adopt new traffic rules without forgetting how to drive a car. So, it must be emphasized here that the approach to be presented in this paper is also applicable to propositional, or parallel [Delgrande and Jin, 2008] belief revision tasks, as propositions are considered as conditionals with tautological antecedent here. We will also establish the connection between our approach to belief revision and nonmonotonic reasoning.

This paper is structured as follows: Section 2 recalls basic facts about the logical setting of our approach. The constructive approach is developed in section 3; this section also presents the novel *enforcement postulate* and its semantical characterization. Section 4 elaborates on general links to related work and goes into more detail for the connection to parallel belief revision. We conclude in section 5.

2 Preliminaries on logic and conditionals

Let $\mathcal{L}$ be a finitely generated propositional language, with atoms $a, b, c, \ldots$, and with formulas $A, B, C, \ldots$. For conciseness of notation, we will omit the logical *and*-connector, writing AB instead of $A \wedge B$, and overlining formulas will indicate negation, i.e. $\overline{A}$ means $\neg A$. Let Ω denote the set of possible worlds over $\mathcal{L}$; Ω will be taken here simply as the set of all propositional interpretations over $\mathcal{L}$. $\omega \models A$ means that the propositional formula $A \in \mathcal{L}$ holds in the possible world $\omega \in \Omega$. By slight abuse of notation, we will use ω both for the model and the corresponding conjunction of all positive or negated atoms.

By introducing a new binary operator $|$, we obtain the set $(\mathcal{L} \mid \mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}$ of conditionals over $\mathcal{L}$. $(B|A)$ formalizes "*if* A *then* B" and establishes a plausible connection between the *antecedent* A and the *consequent* B. Here, conditionals are supposed not to be nested, that is, antecedent and consequent of a conditional will be propositional formulas from $\mathcal{L}$. Conditionals with tautological antecedents are taken as plausible statements about the world. Following De Finetti [DeFinetti, 1974], a conditional $(B|A)$ can be *verified* (*falsified*) by a possible world ω iff $\omega \models AB$ ($\omega \models A\overline{B}$). If $\omega \not\models A$, then the conditional is *not applicable* to ω. A conditional $(B|A)$ is said to be *non-trivial* if neither AB nor $A\overline{B}$ is contradictory.

A conditional $(D|C)$ is a *subconditional* of $(B|A)$, written as $(D|C) \sqsubseteq (B|A)$, iff $CD \models AB$ and $C\overline{D} \models A\overline{B}$, i.e., iff verification resp. falsification of $(D|C)$ implies verification resp. falsification of $(B|A)$. A *weak subconditional* of $(B|A)$ is a subconditional of $(A \Rightarrow B|\top)$ where $\Rightarrow$ means material implication. It is straightforward to show that subconditionals are also weak subconditionals of the respective conditional. For different worlds $\omega_1, \omega_2 \in \Omega$, $(\omega_1|\omega_1 \vee \omega_2)$ is called a *basic conditional* that is verified by ω_1 and falsified by ω_2.

Following [Katsuno and Mendelzon, 1991] and [Darwiche and Pearl, 1997], AGM revisions of epistemic states can be ensured by assuming faithful rankings (i.e., total preorders) that underly the epistemic states such that the revised beliefs can be computed from minimal models according to the ranking. In this paper, we take ordinal conditional functions as representations of such faithful rankings, resp., as representations of epistemic states: *Ordinal conditional functions, OCFs*, (also called *ranking functions*) are functions $\kappa : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$, expressing degrees of plausibility of propositional formulas A by specifying degrees of disbeliefs of their negations $\overline{A}$ (cf. [Spohn, 1988]). More formally, we have $\kappa(A) := \min\{\kappa(\omega) \mid \omega \models A\}$. A proposition A is believed, in symbols $\kappa \models A$, iff $\kappa(\overline{A}) > 0$. Degrees of plausibility can also be assigned to conditionals by setting $\kappa(B|A) = \kappa(AB) - \kappa(A)$. A conditional $(B|A)$ is accepted in the epistemic state represented by κ, written as $\kappa \models (B|A)$, iff $\kappa(AB) < \kappa(A\overline{B})$, i.e. iff AB is more plausible than $A\overline{B}$. In particular, for a basic conditional $(\omega_1|\omega_1 \vee \omega_2)$, we have $\kappa \models (\omega_1|\omega_1 \vee \omega_2)$ iff $\kappa(\omega_1) < \kappa(\omega_2)$. A set $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ is *consistent* iff it has an OCF model, i.e. iff there is an OCF κ such that $\kappa \models (B_i|A_i)$ for all $i, 1 \leqslant i \leqslant n$. κ_u is the *uniform* OCF with $\kappa_u(\omega) = 0$ for all $\omega \in \Omega$. OCF's are semi-qualitative counterparts of probability distributions. Their plausibility degrees may be taken as order-of-magnitude abstractions of probabilities (cf. [Goldszmidt and Pearl, 1996]). Conditionals, epistemic states and AGM revision operators $*$ can be related by the *Ramsey test* [Ramsey, 1950]: $\kappa \models (B|A)$ iff $\kappa * A \models B$.

3 Enforcing revision based on conditional structures

The main contribution of this paper will be the proposal of (iterative) revision operators that take epistemic states (represented by ranking functions κ) and general information sets (consisting of plausible propositions and conditionals) and return revised epistemic states that accept all new pieces of information specified by the general information set and satisfy all major postulates (AGM [Alchourrón *et al.*, 1985], DP-postulates [Darwiche and Pearl, 1997], Independence [Jin and Thielscher, 2007], and Evidence Retainment [Delgrande and Jin, 2008]), when applied to the respective restricted framework (i.e., a single proposition, a single conditional, and sets of propositions).

Since plausible statements can be considered as conditionals with tautological antecedents, the formal specification of the basic technical revision task is given as follows:

> Let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a finite set of conditionals, let κ be a prior epistemic state represented by an OCF. Specify revision operators $*$ that compute a revised OCF $\kappa^* = \kappa * \mathcal{R}$ by taking (somehow, to be further specified by postulates) prior and new information into account.

We will focus on consistent sets $\mathcal{R}$, since revising by an inconsistent set of information should result in an inconsistent state of belief, if no aspects of consolidation or merging are taken into regard which we leave for future work. The crucial points with this task are, first, to take into account that conditionals are logically more complex than propositions and hence may require more sophisticated methods not provided by ideas from propositional revision, and, second, to treat each element of the general information set as an independent piece of information. For both requirements, the theory of conditional structures [Kern-Isberner, 2001; 2004] provides a suitable basis. Conditional structures are kind of footprints that sets of conditionals leave on possible worlds. They base on the three-valued approach to conditionals by De Finetti [DeFinetti, 1974], representing verification, falsification, and non-applicability by abstract symbols that are assigned to conditionals (or propositions) in a general information set. Since different conditionals are assigned different symbols, each conditional is treated as an *independent piece of information*.

More formally, let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a finite set of conditionals, and let $\mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^-$ be distinct algebraic symbols that are used as generators of a (free abelian[1]) group [Fine and Rosenberger, 1999] $\mathcal{F}_\mathcal{R} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$. In short, this group structure provides us with a multiplication (written as juxtaposition) that we can relate to the summation of κ-values, and with a 1 that symbolizes neutrality (i.e. non-applicability). Furthermore, free abelian ensures a parallel handling of the conditionals (without any order of application assumed) as well as independence between different conditionals (by forbidding cancellations between different symbols). For each $i, 1 \leqslant i \leqslant n$, we define a function $\sigma_i = \sigma_{(B_i|A_i)} : \Omega \to \mathcal{F}_\mathcal{R}$ by setting

$$\sigma_i(\omega) := \begin{cases} \mathbf{a}_i^+ & \text{if} \quad \omega \models A_i B_i \quad \text{(verification)} \\ \mathbf{a}_i^- & \text{if} \quad \omega \models A_i \overline{B_i} \quad \text{(falsification)} \\ 1 & \text{if} \quad \omega \models \overline{A_i} \quad \text{(non-applicability)} \end{cases}$$

$\sigma_i(\omega)$ represents the manner in which the conditional $(B_i|A_i)$ applies to the possible world ω. The function $\sigma_\mathcal{R} : \Omega \to \mathcal{F}_\mathcal{R}$,

$$\sigma_\mathcal{R}(\omega) := \prod_{1 \leqslant i \leqslant n} \sigma_i(\omega) = \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \mathbf{a}_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} \mathbf{a}_i^-$$

describes the all-over effect of $\mathcal{R}$ on ω. $\sigma_\mathcal{R}(\omega)$ is called the *conditional structure of* ω *with respect to* $\mathcal{R}$. Since $\mathcal{F}_\mathcal{R}$ is a free (abelian) group, the conditional structures of worlds are uniquely determined by their σ_i-components and hence by their logical relation to each conditional: For any two worlds ω_1, ω_2, we have

$$\sigma_\mathcal{R}(\omega_1) = \sigma_\mathcal{R}(\omega_2) \text{ iff } \sigma_i(\omega_1) = \sigma_i(\omega_2) \; \forall \, 1 \leqslant i \leqslant n. \quad (1)$$

Conditional structures also work for propositions $A_i \equiv (A_i|\top) \in \mathcal{R}$. Propositions can only be verified resp. satisfied, or falsified. There is no non-applicability in this case, so $\sigma_i(\omega) \in \{\mathbf{a}_i^-, \mathbf{a}_i^+\}$ for each ω. Nevertheless, conditional structures are helpful to distinguish between the effects of different pieces of information also in this case.

Example 1 Consider the following scenario described by the set $\mathcal{R} = \{r_1, \ldots, r_9\}$, where penguins ($p$), kiwis ($k$), and doves ($d$) are modelled as disjoint classes of birds (b) which usually fly (f), and typically have wings (w) and eyes (e) :

r_1: $(b\|p)$	r_2: $(\overline{f}\|p)$	r_3: $(w\|b)$
r_4: $(f\|w)$	r_5: $(b\|k)$	r_6: $(b\|d)$
r_7: $(\overline{w}\|k)$	r_8: $(w\|f)$	r_9: $(e\|b)$

Assigning symbols $\mathbf{a}_i^+, \mathbf{a}_i^-$ to any r_i, we compute the following conditional structures for possible worlds:

$$\begin{aligned} \sigma_\mathcal{R}(p\overline{k}\,\overline{d}bfwe) &= \mathbf{a}_1^+\mathbf{a}_2^-\mathbf{a}_3^+\mathbf{a}_4^+\mathbf{a}_8^+\mathbf{a}_9^+ \\ \sigma_\mathcal{R}(p\overline{k}\,\overline{d}bf\overline{w}e) &= \mathbf{a}_1^+\mathbf{a}_2^-\mathbf{a}_3^-\mathbf{a}_8^-\mathbf{a}_9^+ \\ \sigma_\mathcal{R}(\overline{p}k\overline{d}\,\overline{b}\,fw\overline{e}) &= \mathbf{a}_4^+\mathbf{a}_5^-\mathbf{a}_7^-\mathbf{a}_8^+\mathbf{a}_9^-. \end{aligned}$$ ■

Based on conditional structures, a principle of conditional preservation was proposed in [Kern-Isberner, 2004]. It controls conditional belief revision by claiming a balance between conserving conditional relationships in the prior epistemic state and establishing conditional relationships from the information set. To illustrate this, we state a simple consequence of this principle as the following property which will prove to be useful here:

(SCondPres) If for $\omega_1, \omega_2 \in \Omega$, it holds that $\sigma_\mathcal{R}(\omega_1) = \sigma_\mathcal{R}(\omega_2)$, then $\kappa^*(\omega_1) - \kappa(\omega_1) = \kappa^*(\omega_2) - \kappa(\omega_2)$.

(SCondPres) claims that the relative amount of change between prior and posterior epistemic state only depends on the conditional structure of the respective world. As a general consequence of the principle of conditional preservation, the revised epistemic state follows a structure that assigns to each conditional in the new information set a constant parameter that represents the impact of each conditional in the revision process. When revising an ordinal conditional function κ by a set of conditionals $\mathcal{R}$, it is proved in [Kern-Isberner, 2004] that revisions $\kappa^* = \kappa * \mathcal{R}$ of the form

$$\kappa^*(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} \kappa_i^- \quad (2)$$

with non-negative, integral parameters κ_i^- and a normalization constant κ_0, satisfy the full principle of conditional preservation and hence also (SCondPres). Revisions of this kind are called *c-revisions* in [Kern-Isberner, 2004]. The parameters κ_i^- are numerical counterparts of the abstract symbols $\mathbf{a}_i^-$ whereas the $\mathbf{a}_i^+$ have no numerical impact by implicitly setting $\kappa_i^+ = 0$. They have to be set according to general constraints on the outcome of the revision process, i.e. postulates in the spirit of the AGM-theory. A constraint that is typically applied in revision scenarios and will be used here to determine these parameters is *success* on the level of the conditionals in $\mathcal{R}$:

(Success) $\kappa * \mathcal{R} \models \mathcal{R}$, i.e., $\kappa * \mathcal{R} \models (B|A)$ for all $(B|A) \in \mathcal{R}$.

For c-revisions κ^* of the form (2), (Success) is satisfied iff for all $i, 1 \leqslant i \leqslant n$,

$$\begin{aligned} \kappa_i^- \quad > \quad & \min_{\omega \models A_i B_i}(\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \\ & - \min_{\omega \models A_i \overline{B}_i}(\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-). \end{aligned} \quad (3)$$

We will now follow ideas from [Kern-Isberner, 2002] and claim that, ceteris paribus and generally, falsification of conditionals should have a penalizing effect. To implement this, a partial ordering $>$ between conditional structures of worlds is defined in the following way: For any possible worlds $\omega_1, \omega_2 \in \Omega$, we define

$$\sigma_\mathcal{R}(\omega_1) > \sigma_\mathcal{R}(\omega_2)$$

iff for each i, $1 \leqslant i \leqslant n$, $\sigma_i(\omega_1) = \mathbf{a}_i^-$ implies $\sigma_i(\omega_2) = \mathbf{a}_i^-$, and there is at least one i, $1 \leqslant i \leqslant n$, such that $\sigma_i(\omega_1) \in \{\mathbf{a}_i^+, 1\}$ and $\sigma_i(\omega_2) = \mathbf{a}_i^-$. That is, $\sigma_\mathcal{R}(\omega_1) > \sigma_\mathcal{R}(\omega_2)$ iff ω_2 falsifies more conditionals than ω_1. So, for instance, in example 1, we have $\sigma_\mathcal{R}(p\overline{k}\,\overline{d}bfwe) = \mathbf{a}_1^+\mathbf{a}_2^-\mathbf{a}_3^+\mathbf{a}_4^+\mathbf{a}_8^+\mathbf{a}_9^+ > \mathbf{a}_1^+\mathbf{a}_2^-\mathbf{a}_3^-\mathbf{a}_8^-\mathbf{a}_9^+ = \sigma_\mathcal{R}(p\overline{k}\,\overline{d}bf\overline{w}e)$, but $\sigma_\mathcal{R}(p\overline{k}\,\overline{d}bf\overline{w}e)$ and $\sigma_\mathcal{R}(\overline{p}k\overline{d}\,\overline{b}\,fw\overline{e})$ are not comparable with respect to $>$.

Now, both conditional information from the prior epistemic state κ and the new information set $\mathcal{R}$ are used to *enforce* conditional beliefs in the revised epistemic state $\kappa^* = \kappa * \mathcal{R}$:

(Enf) If for $\omega_1, \omega_2 \in \Omega$, it holds that $\sigma_\mathcal{R}(\omega_1) > \sigma_\mathcal{R}(\omega_2)$, then $\kappa(\omega_1) \leqslant \kappa(\omega_2)$ implies $\kappa^*(\omega_1) < \kappa^*(\omega_2)$.

[1]Free abelian groups have no relations except for those induced by commutativity.

This enforcement postulate claims that if ω_2 falsifies more conditionals in $\mathcal{R}$ than ω_1 (ceteris paribus and with respect to set inclusion) and is not more plausible than ω_1 in the prior epistemic state, then it should definitely be less plausible than ω_1 in the revised epistemic state.

We will illustrate how the concept of c-revisions together with the postulates of (Success) and (Enf) yield elaborate revisions in the scenario of example 1. We will also show how the machinery that has been developed so far can be used for inductive default reasoning.

Example 2 We continue example 1, so let $\mathcal{R}$ be given as above. First, we use the method of successful and enforcing c-revision for inductive nonmonotonic reasoning. Let $*$ be a c-revision operator yielding revisions as in (2) and satisfying (Success) and (Enf). Then $\kappa_1 = \kappa_u * \mathcal{R}$ is a new epistemic state of the form (2) that accepts all conditionals in $\mathcal{R}$. The κ_i^- for this example have been calculated as the minimal non-negative integers that satisfy (3) (so that (Success) is ensured) and (Enf) (note that $\kappa_u(\omega) = 0$ for all ω, so the partial order $>$ on conditional structures can be implemented whenever applicable). For instance, we find that $\kappa_3^- = \kappa_4^- = 1$, and $\kappa_1^- = \kappa_2^- = 2$. In κ_1, among others, the following conditionals are accepted: $(\overline{f}|k)$ (kiwis do not fly), $(f|d), (w|d)$ (doves fly and have wings). Please note that no immediate connection between birds and the ability to fly is encoded in $\mathcal{R}$, but just an indirect link via having wings. In particular, κ_1 is undecided concerning penguins having wings, neither $(w|p)$ nor $(\overline{w}|p)$ can be derived because wings are deemed crucial for flying, and penguins can not fly. However, from κ_1 the agent can derive the beliefs that all birds – in particular kiwis and penguins – usually have eyes.

Now, we revise κ_1 by the set $\mathcal{S} = \{s_1 = (\overline{w}|d), s_2 = (f|p)\}$ after the fictitious observations that doves have no wings, and that penguins can fly. We compute $\kappa_2 = \kappa_1 * \mathcal{S}$ in the same way as above. Based on this new epistemic state, the agent believes now that doves can not fly any more ($\kappa_2 \models (\overline{f}|d)$), but that penguins have wings ($\kappa_2 \models (w|p)$). No belief concerning kiwis is affected by the revision, so the agent still holds the same beliefs about kiwis as in κ_1. ■

The following theorem characterizes the semantic (Enf) postulate by exploiting the prior epistemic state and the explicitly given new conditional information:

Theorem 3 *Let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a finite set of conditionals, let κ be an OCF. Let the revision operator $*$ be compatible with propositional AGM revision. Then the revision $\kappa^* = \kappa * \mathcal{R}$ satisfies (Enf) iff it satisfies the* postulate of conditional enforcement*:*

(CondEnf) *For any weak subconditional $(B_i'|A_i')$ of one of the conditionals in $\mathcal{R}$ (i.e. $(B_i'|A_i') \sqsubseteq (A_i \Rightarrow B_i|\top)$ for some i) such that*

1. *no non-trivial subconditional $(B_i''|A_i'')$ of $(B_i'|A_i')$ is rejected by κ, i.e., $\kappa \not\models (\overline{B_i''}|A_i'')$ for all $(B_i''|A_i'') \sqsubseteq (B_i'|A_i')$, and*
2. *for all other conditionals $(B_j|A_j)$ in $\mathcal{R}$, verification of $(B_i'|A_i')$ implies satisfaction of $(B_j|A_j)$, or falsification of $(B_i'|A_i')$ implies falsification of $(B_j|A_j)$, i.e., for all $j \neq i$, it holds that $A_i'B_i' \models A_j \Rightarrow B_j$ or $A_i'\overline{B_i'} \models A_j\overline{B_j}$,*

it holds that $\kappa * \mathcal{R} \models (B_i'|A_i')$.

The conditional enforcement postulate claims that any weak subconditional of any $(B_i|A_i) \in \mathcal{R}$ (and hence any subconditional obtained by strengthening the antedent) is enforced in the revised epistemic state $\kappa * \mathcal{R}$ iff there is no evidence to the contrary neither due to the prior epistemic state κ (as specified by condition (1)) nor from the other conditionals in $\mathcal{R}$ (as specified by (2)). Note that Theorem 3 holds for general condtional revision operators, not only for c-revisions.

Proof of Theorem 3. First, we assume that (Enf) holds. Let $(B_i'|A_i')$ be a weak subconditional of $(B_i|A_i) \in \mathcal{R}$ as specified in the theorem. We have to show $\kappa * \mathcal{R} \models (B_i'|A_i')$, i.e., $\kappa * \mathcal{R}(A_i'B_i') < \kappa * \mathcal{R}(A_i'\overline{B_i'})$. Let $\omega_1, \omega_2 \in \Omega$ be such that $\omega_1 \models A_i'B_i'$ and $\omega_2 \models A_i'\overline{B_i'}$. Then $(\omega_1|\omega_1 \vee \omega_2)$ is a (non-trivial) subconditional of $(B_i'|A_i')$. By condition (1) of (CondEnf), we have $\kappa \not\models (\overline{\omega_1}|\omega_1 \vee \omega_2)$ which is equivalent to $\kappa(\omega_1) \leqslant \kappa(\omega_2)$. Consider now $\sigma_{\mathcal{R}}(\omega_1)$ and $\sigma_{\mathcal{R}}(\omega_2)$. Due to $\omega_1 \models A_i'B_i'$ and $\omega_2 \models A_i'\overline{B_i'}$, we have $\sigma_i(\omega_1) = \mathbf{a}_i^+$ and $\sigma_i(\omega_2) = \mathbf{a}_i^-$. For all other $j \neq i$, if $A_i'B_i' \models A_j \Rightarrow B_j$, then $\sigma_i(\omega_1) \neq \mathbf{a}_i^-$. If $A_i'B_i' \not\models A_j \Rightarrow B_j$, then $A_i'\overline{B_i'} \models A_j\overline{B_j}$ by condition (2) of (CondEnf). In this latter case, $\sigma_i(\omega_2) = \mathbf{a}_i^-$. So, $\sigma_{\mathcal{R}}(\omega_1) > \sigma_{\mathcal{R}}(\omega_2)$. Together with $\kappa(\omega_1) \leqslant \kappa(\omega_2)$, (Enf) now implies $\kappa * \mathcal{R}(\omega_1) \leqslant \kappa * \mathcal{R}(\omega_2)$, and therefore $\kappa * \mathcal{R}(A_i'B_i') < \kappa * \mathcal{R}(A_i'\overline{B_i'})$.

For the other direction, we presuppose that (CondEnf) holds, and we have to show the validity of (Enf). So, let $\omega_1, \omega_2 \in \Omega$ such that $\sigma_{\mathcal{R}}(\omega_1) > \sigma_{\mathcal{R}}(\omega_2)$ and $\kappa(\omega_1) \leqslant \kappa(\omega_2)$. In order to show $\kappa^*(\omega_1) < \kappa^*(\omega_2)$, we prove the equivalent statement $\kappa^* \models (\omega_1|\omega_1 \vee \omega_2)$. By exploiting the prerequisite $\sigma_{\mathcal{R}}(\omega_1) > \sigma_{\mathcal{R}}(\omega_2)$, we find in particular, that there must be an index i such that $\sigma_{\mathcal{R}}(\omega_1)$ must be $\mathbf{a}_i^+$ or 1, and $\sigma_{\mathcal{R}}(\omega_2) = \mathbf{a}_i^-$. Hence $\omega_1 \models A_i \Rightarrow B_i$ and $\omega_2 \models A_i\overline{B_i}$, so $(\omega_1|\omega_1 \vee \omega_2)$ is a weak subconditional of $(B_i|A_i) \in \mathcal{R}$. Since $(\omega_1|\omega_1 \vee \omega_2)$ has no non-trivial subconditionals except for itself, and $\kappa \not\models (\overline{\omega_1}|\omega_1 \vee \omega_2)$ due to $\kappa(\omega_1) \leqslant \kappa(\omega_2)$, (1) of (CondEnf) is satisfied. Next, let $j \neq i$. Again thanks to $\sigma_{\mathcal{R}}(\omega_1) > \sigma_{\mathcal{R}}(\omega_2)$, we have that whenever $\sigma_j(\omega_1) = \mathbf{a}_j^-$ (in which case $\omega_1 \not\models A_j \Rightarrow B_j$), then $\sigma_j(\omega_2) = \mathbf{a}_j^-$ must also hold, i.e., $\omega_2 \models A_j\overline{B_j}$ is ensured by $\sigma_{\mathcal{R}}(\omega_1) > \sigma_{\mathcal{R}}(\omega_2)$ in this case. So, $(\omega_1|\omega_1 \vee \omega_2)$ is a conditional that fulfills all requirements specified in the prerequisites of (CondEnf), and so (CondEnf) postulates $\kappa^* \models (\omega_1|\omega_1 \vee \omega_2)$. □

In summary, as a powerful set of postulates that should guide the revision of epistemic states (given by ordinal conditional functions) by sets of (propositional and conditional) beliefs, we propose a principle of conditional preservation, as specified in [Kern-Isberner, 2004] and implemented by c-revisions of the form (2), together with (Success) and (Enf) (resp. (CondEnf)). A straightforward consequence of Theorem 3 is the following proposition:

Proposition 4 *Let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a finite set of conditionals, let κ be an OCF. If the revision $\kappa^* = \kappa * \mathcal{R}$ satisfies (Enf) then for any conditional $(B_i|A_i)$ in $\mathcal{R}$, and for any strengthening $A \models A_i$ of its* *antecedent such that $\kappa * A' \not\models \overline{B_i}$ for any further strengthening $A' \models A$, and such that for all $j \neq i$, $AB_i \models A_j \Rightarrow B_j$ or $A\overline{B_i} \models A_j\overline{B_j}$, it holds that $(\kappa * \mathcal{R}) * A \models B_i$.*

This proposition has interesting consequences for investigating monotonicity in nonmonotonic reasoning. Applying a successful, enforcing c-revision operator $*$ to the uniform OCF κ_u (with $\kappa_u(\omega) = 0$ for all $\omega \in \Omega$) and to a set of conditionals $\mathcal{R}$ gives rise to an inductive default inference relation similar to system Z [Goldszmidt and Pearl, 1996], setting $C \mid\!\sim_{\mathcal{R}} D$ iff $(\kappa_u * \mathcal{R}) * C \models D$, or $(\kappa_u * \mathcal{R}) \models (D|C)$, respectively (see example 2). In particular, (Success) ensures that $A_i \mid\!\sim_{\mathcal{R}} B_i$ for all conditionals $(B_i|A_i) \in \mathcal{R}$. Since κ_u does not accept any nontrivial conditional beliefs, the first presupposition in proposition 4 which a strengthening A of A_i has to satisfy is vacuous (since $\mathcal{R}$ is assumed to be consistent). So only the second presupposition applies, stating that also $A \mid\!\sim_{\mathcal{R}} B_i$ for $A \models A_i$ if for all $j \neq i$, $AB_i \models A_j \Rightarrow B_j$, or $A\overline{B_i} \models A_j\overline{B_j}$. For instance, in example 2, the inference $dbwe \mid\!\sim_{\mathcal{R}} f$ (dove birds with wings and eyes can also fly) can be directly derived from the knowledge base $\mathcal{R}$. Note that the inference relation $\mid\!\sim_{\mathcal{R}}$ also complies with system P and rational monotonicity [Goldszmidt and Pearl, 1996], since it is based on an ordinal conditional function, and satisfies advanced postulates like *specificity*, *irrelevance*, *property inheritance* (hence drowning problems are avoided), *ambiguity preservation*, and *duplication stability*, since it implements reasoning based on conditional structures (for further information, please see [Kern-Isberner, 2002]).

4 Related work

As one of the first and most important approaches to extend AGM-theory towards epistemic states Ψ and iterated revision operators $*$, and to also take issues of conditional preservation into regard, [Darwiche and Pearl, 1997] proposed four new postulates (*DP-postulates*) for iterated revision. In [Kern-Isberner, 2004] it is proved that c-revisions are AGM revisions (when applied to a single proposition) and satisfy all DP-postulates.

Three of these DP-postulates have been widely accepted but the second one (C2) has been intensely debated:

(C2) If $B \models \neg A$ then $(\Psi * A) * B \equiv \Psi * B$.

For instance, Delgrande and Jin [Delgrande and Jin, 2008] found this postulate to be the culprit for so-called "drowning problems" which are caused by the suppression of effects of some pieces of information by other pieces of information so that they can not have a proper impact on the resulting belief state. In particular, these authors argue that the generalization of (C2) for sets of propositions, as proposed by [Zhang, 2004] is completely inappropriate. However, exploring (C2) for conditional beliefs, its relevance becomes apparent: (C2) claims that if $B \models \neg A$ then $(\Psi * A) * B \models C$ iff $\Psi * B \models C$, so using the connection between revision and conditionals via the Ramsey test, we have $(\Psi * A) \models (C|B)$ iff $\Psi \models (C|B)$ – revising by A should leave conditional beliefs intact the antecedent of which is disjoint from A, which makes perfect sense. (C2) has been generalized to revision by a single conditional belief in [Kern-Isberner, 2001].

The problems that Delgrande and Jin encountered in [Delgrande and Jin, 2008] are principally caused by taking the information sets $\{A, B\}$ and $\{A \wedge B\}$ to be equivalent, ignoring the independence of each piece of information. So, they propose an approach to revise epistemic states $\mathcal{K}$ by sets of propositional formulas S, called parallel belief revision and denoted by $\mathcal{K} \otimes S$. The aim of their framework is to allow for the explicit representation of individual items of new information. For instance, in $\mathcal{K} \otimes \{\alpha, \beta\}$, α and β represent individual items of information while in $\mathcal{K} \otimes \{\alpha \wedge \beta\}$, their coincidence is taken as one piece of information. This means that in that framework, $b \in Bel(\mathcal{K} \otimes \{a, b\} \otimes \{\neg a\})$ but $b \notin Bel(\mathcal{K} \otimes \{a \wedge b\} \otimes \{\neg a\})$. Delgrande and Jin demand that all formulas of S should be established individually such that in case of further revision all items in S against which no evidence exists should be retained. This demand is formalized by their postulate of *evidence retainment*, the equivalent semantical condition of which on a faithful ranking $\preceq_{\mathcal{K}}$ is given as follows:

(Ret$^{\otimes}$R) If $S|\omega_2 \subset S|\omega_1$ then $\omega_1 \preceq_{\mathcal{K}} \omega_2$ implies $\omega_1 \prec_{\mathcal{K}\otimes S} \omega_2$

with $S|\omega = \{\alpha \in S \mid \omega \models \alpha\}$. The *independence* postulate of [Jin and Thielscher, 2007] covers the case $|S| = 1$. In addition to (Ret$^{\otimes}$R), Delgrande and Jin extend two of the DP-postulates, (C3) and (C4) [Darwiche and Pearl, 1997], for parallel belief revision in the following semantical characterization:

(PC3$^{\otimes}$R) If $S|\omega_2 \subseteq S|\omega_1$ then $\omega_1 \prec_{\mathcal{K}} \omega_2$ implies $\omega_1 \prec_{\mathcal{K}\otimes S} \omega_2$

(PC4$^{\otimes}$R) If $S|\omega_2 \subseteq S|\omega_1$ then $\omega_1 \preceq_{\mathcal{K}} \omega_2$ implies $\omega_1 \preceq_{\mathcal{K}\otimes S} \omega_2$

Delgrande and Jin conclude that an adequate operator for parallel belief revision should satisfy the AGM postulates generalized to sets, (Ret$^{\otimes}$R), (PC3$^{\otimes}$R) and (PC4$^{\otimes}$R).

To match the framework of Delgrande and Jin, we consider the propositional case of our proposed conditional revision operator. Thus, let $\mathcal{R} = \{(A_1|\top), \ldots, (A_n|\top)\}$ with $A_i \equiv (A_i|\top)$. In this restricted setting we have $\sigma_i(\omega) = \mathbf{a}_i^+$ if $\omega \models A_i$ and $\sigma_i(\omega) = \mathbf{a}_i^-$ if $\omega \models \overline{A_i}$. In the notation of Delgrande and Jin, the correspondence to our framework is easily seen by $\sigma_{\mathcal{R}}(\omega) = \prod_{A_i \in \mathcal{R}|\omega} \mathbf{a}_i^+ \prod_{A_i \in \mathcal{R}\setminus\mathcal{R}|\omega} \mathbf{a}_i^-$, so that we have $\sigma_{\mathcal{R}}(\omega_1) > \sigma_{\mathcal{R}}(\omega_2)$ iff $R|\omega_2 \subset R|\omega_1$. With this direct connection to the approach of Delgrande and Jin we obtain the following result.

Proposition 5 *Let $\mathcal{R} = \{(A_1|\top), \ldots, (A_n|\top)\}$ be a finite set of propositional conditionals, let κ be an OCF, and let $\kappa * \mathcal{R}$ be a c-revision that satisfies (Enf). Then $\kappa * \mathcal{R}$ satisfies (Ret$^{\otimes}$R), (PC3$^{\otimes}$R) and (PC4$^{\otimes}$R).*

Proof of Proposition 5. The satisfaction of (Ret$^{\otimes}$R) and the proper subset case for the precondition of (PC3$^{\otimes}$R) and (PC4$^{\otimes}$R) follow directly from (Enf). For the case of $\mathcal{R}|\omega_1 = \mathcal{R}|\omega_2$ for (PC3$^{\otimes}$R) and (PC4$^{\otimes}$R) it follows from the satisfaction of (SCondPres) by c-revisions that $\kappa^*(\omega_1) - \kappa(\omega_1) = \kappa^*(\omega_2) - \kappa(\omega_2)$ and especially that $\kappa(\omega_1) \leqslant \kappa(\omega_2)$ implies $\kappa^*(\omega_1) \leqslant \kappa^*(\omega_2)$ and $\kappa(\omega_1) < \kappa(\omega_2)$ implies $\kappa^*(\omega_1) < \kappa^*(\omega_2)$. Hence (PC3$^{\otimes}$R) and (PC4$^{\otimes}$R) are satisfied. ☐

The approach we present in this paper is therefore a constructive approach and a conditional generalization of a parallel, and iterated, belief revision operator that satisfies all desirable properties of parallel belief revision. Actually, both *evidence retainment* and *independence* postulate more than simple independence: They claim that each piece of information should have a noticeable impact on the resulting epistemic state (that is largely independent of the impact of other pieces of information). So, they are both in conflict with a postulate called *stability* that has been considered in [Kern-Isberner, 2001], requiring that $\Psi * A = \Psi$ if $\Psi \models A$. *Stability* claims that no difference should be made between implicitly derived and explicitly stated beliefs. Hence, the postulate of *independence* [Jin and Thielscher, 2007] rather addresses the difference between these two types of belief, stating that one expects the explicit beliefs to be enforced if there is no reason against it. *Evidence retainment* does exactly the same for pieces of information that are independent in that they are given by different propositions. Conditional structures represent each piece of explicitly given information by algebraic symbols, therefore implementing both independence and explicitness of information, but leaving the exploration of logical and numerical (or preferential) constraints to the principle of conditional preservation (which is able to handle much more complicated relationships between conditional structures than those shown in (SCondPres)). The enforcement postulate (CondEnf) comes on top of that, as an independent postulate, claiming the thorough establishing of conditional beliefs with respect to subconditionals. It is immediately seen from its semantical characterization in (Enf) that (CondEnf) generalizes *independence* and *evidence retainment* for conditional beliefs. Analogous to *independence* and *evidence retainment*, this principle makes a strong difference between explicit and implicit beliefs.

5 Conclusion

In this paper, we present principles and a constructive approach to iterative belief revision of epistemic states by sets of (propositional or conditional) sentences. In more detail, exploiting ideas from [Kern-Isberner, 2002] concerning default reasoning for the tasks of belief revision gives rise to an *enforcement postulate* that demands for establishing conditional beliefs also for strengthenings of the antecedents. We combine this with a principle of conditional preservation that has been proposed in [Kern-Isberner, 2004] and yields a schema for revision operations as the basis for our constructive approach. In this way, our approach implements both an *in-depth control of conditional dependencies*, due to the principle of conditional preservation, and an *in-depth enforcement of conditional dependencies*. Actually, it is this latter principle that underlies the ideas of *independence* [Jin and Thielscher, 2007] and *evidence retainment* [Delgrande and Jin, 2008] set up for revising by (sets of) propositions. Together with the success postulate, the postulates of *conditional preservation* and *enforcement* are shown to cover also the advanced postulates for iterated revision proposed in [Darwiche and Pearl, 1997; Jin and Thielscher, 2007; Delgrande and Jin, 2008]. As part of our future work, we will extend the proposed framework by considering inconsistent sets of information that will be consolidated by merging operations before revision.

References

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

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

[DeFinetti, 1974] B. DeFinetti. *Theory of Probability*, volume 1,2. John Wiley and Sons, New York, 1974.

[Delgrande and Jin, 2008] James Delgrande and Yi Jin. Parallel belief revision. In *Proceedings of the 23rd national conference on Artificial intelligence - Volume 1*, pages 430–435. AAAI Press, 2008.

[Fine and Rosenberger, 1999] B. Fine and G. Rosenberger. *Algebraic Generalizations of Discrete Groups*. Dekker, New York, Basel, 1999.

[Goldszmidt and Pearl, 1996] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84:57–112, 1996.

[Jin and Thielscher, 2007] Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.

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

[Kern-Isberner, 2001] G. Kern-Isberner. *Conditionals in nonmonotonic reasoning and belief revision*. Springer, Lecture Notes in Artificial Intelligence LNAI 2087, 2001.

[Kern-Isberner, 2002] G. Kern-Isberner. A structural approach to default reasoning. In *Proceedings KR'2002*, pages 147–157, San Francisco, Ca., 2002. Morgan Kaufmann.

[Kern-Isberner, 2004] G. Kern-Isberner. A thorough axiomatization of a principle of conditional preservation in belief revision. *Annals of Mathematics and Artificial Intelligence*, 40(1-2):127–164, 2004.

[Kern-Isberner, 2008] G. Kern-Isberner. Linking iterated belief change operations to nonmonotonic reasoning. In*Proceedings KR'2008*, pages 166–176, Menlo Park, CA, 2008. AAAI Press.

[Ramsey, 1950] F.P. Ramsey. General propositions and causality. In R.B. Braithwaite, editor, *Foundations of Mathematics and other logical essays*, pages 237–257. Routledge and Kegan Paul, New York, 1950.

[Spohn, 1988] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics, II*, pages 105–134. Kluwer Academic Publishers, 1988.

[Zhang, 2004] D. Zhang. Properties of iterated multiple belief revision. In *Proceedings LPNMR-04*, pages 314–325. Springer, 2004.