# A conceptual framework for (iterated) revision, update, and nonmonotonic reasoning

Gabriele Kern-Isberner

Department of Computer Science
University of Dortmund, Germany
gabriele.kern-isbernercs.uni-dortmund.de

**Abstract.** This paper makes a foundational contribution to the discussions on the very nature of belief change operations. Belief revision and belief update are investigated within an abstract framework of epistemic states and (qualitative or quantitative) conditionals. Moreover, we distinguish between background knowledge and contextual information in order to analyse belief change more appropriately. The rich epistemic representation framework allows us to make a clear conceptual distinction between revision and update on the one side, while revealing structural similarities on the other side. We propose generic postulates for revision and update that also apply to iterated change. Furthermore, we complete the unifying picture by introducing universal inference operations as a proper counterpart in nonmonotonic reasoning to iterated belief change.

## 1 Introduction

Nonmonotonic reasoning and belief change theory are closely related in that they both deal with reasoning under uncertainty and try to reveal sensible lines of reasoning in response to incoming information. The crucial difference between both areas is the role of the current belief state which is only implicit in nonmonotonic reasoning, but explicit and in fact in the focus of interest in belief revision. So the correspondences between axioms of belief change and properties of nonmonotonic inference operations are usually elaborated only in the case that revisions are based on a fixed theory (cf. [MG91]), and very little work has been done to incorporate iterated belief revision in that framework.

However, belief revision theory is not homogeneous in itself. Genuine belief revision following the AGM-theory [AGM85] is different from belief update, as defined by Katsuno and Mendelzon [KM91], although it seems to be difficult to draw a clear line between these two change operations. Usually, belief revision means the process of adjusting prior beliefs to new information in a static world, whereas belief update should be used to adjust prior beliefs to new information in a possibly changing world. This distinction has raised considerably numerous and deep discussions in the past, and it is not clear how nonmonotonic reasoning as the process of deriving plausible beliefs from given knowledge or beliefs, respectively, can be linked to either change operation.

To sum up, it is well-known, that

Dagstuhl Seminar Proceedings 07351
Formal Models of Belief Change in Rational Agents
http://drops.dagstuhl.de/opus/volltexte/2007/1208

– belief revision and belief update are closely related, but may yield different results;
– belief revision and nonmonotonic reasoning are also closely related, but provide different views and options.

In this paper, we will present a unifying framework that allows us to understand and explore these relationships thoroughly. Basically, we will follow the traditional line to relate belief change and nonmonotonic reasoning via the Ramsey test [Ram50]

$$A \mid\!\sim_{(\Psi)} B \quad \text{iff} \quad \Psi * A \models B,$$

with $\Psi$ denoting some epistemic state. That is to say that $B$ is plausibly derived from $A$, given the epistemic background $\Psi$, iff incorporating $A$ into $\Psi$ yields belief in $B$. This last statement is often seen to be equivalent to saying that the conditional belief $(B|A)$ is accepted in $\Psi$, $\Psi \models (B|A)$. As revision strategies, conditionals are of major importance when dealing with iterated belief change.

Therefore, we will explore the relationships between nonmonotonic reasoning and general belief change by considering epistemic states and sets of conditionals instead of theories and propositional beliefs. We will provide a more general framework that not only allows a more accurate representation of belief change via nonmonotonic formalisms, but also gives, vice versa, an important impetus to handle iterated changes. So we will generalize the notion of an inference operation and introduce *universal inference operations* in nonmonotonic logics as a suitable counterpart to (full) change operators. And instead of taking a (propositional) theory as a reference point for inferences and revisions, we will make use of the more comprehensive notion of an epistemic state to base inferences on. In a purely qualitative environment, a preorder might be enough to represent an epistemic state, but one might also choose more sophisticated representation frameworks, such as possibility distributions, ranking functions, and probability distributions. In this paper, $\Psi$ will denote an abstract epistemic state. For the purpose of illustration, we will use a probabilistic environment as a particularly rich epistemic structure but most of the ideas are applicable in qualitative frameworks as well.

Leaving the classical framework also allows a more accurate view on iterated change operations by differentiating between *simultaneous* and *successive revision*. The former will be seen to handle genuine revisions appropriately, while the latter may also model *updating*. This distinction is based on clearly separating background or generic knowledge from evidential or contextual knowledge, a feature that is listed in [DP96] as one of three basic requirements a *plausible exception-tolerant inference system* has to meet.

The new information the agent is going to incorporate into his beliefs is assumed to provide information about the context the agent is focusing on. Since this context may be complex, the information can be complex as well, consisting of a simple fact, some uncertain evidence, or even conditional beliefs. For instance, being situated in some region that has been devastated by a natural disaster, the agent might find that the rules of daily life have changed drastically, and adopting these new conditional beliefs is essential for him to survive. The new information may encompass different chunks of information all pertaining to the given context, so we assume it to be specified by a set $\mathcal{R}$ containing conditionals as the most general items in our framework that may express information. This includes also facts, since facts are considered as degenerate conditionals with tautological antecedents.

The problems we will be dealing with in this paper can be stated as follows:

> Given some epistemic state $\Psi$, possibly consisting of background knowledge and contextual prior information, how to change $\Psi$ so as to incorporate a set of (conditional) beliefs $\mathcal{R}$? (1)
>
> How can plausible inferences drawn on the base of the posterior epistemic state be related to those drawn from the prior state? (2)

These problems can be split up further into the following subproblems:

- Technically, given some epistemic state $\Psi$ and some set of conditional beliefs $\mathcal{R}$, how to compute $\Psi * \mathcal{R}$?
- Conceptually, how to handle background knowledge as well as prior contextual and new contextual information?
- How to take epistemic background explicitly into account in plausible inference?

The outline of the paper is as follows: In the next section, we will introduce some formal notations for epistemic states and conditionals. Section 3 is dedicated to universal inference operations that allow to deal with nonmonotonic inference operations based on different epistemic backgrounds. In section 4, we first discuss update as a basic, imperative change operation. A list of postulates is given providing formal properties of (iterated) belief update. Section 5 links update to universal inference operations. Finally, belief revision is dealt with in section 6, and its connection to universal inference operations and belief update is elaborated in section 7. We conclude in section 8 with a summary and a brief outlook on future work.

## 2 Epistemic states and conditionals

We will use propositions $A, B, \ldots$ as the basic atoms of a propositional logical language $\mathcal{L}$ with the junctors $\wedge$ and $\neg$. The $\wedge$-junctor will mostly be omitted, so that $AB$ stands for $A \wedge B$, and $\neg$ will usually be indicated by barring the corresponding proposition, i.e. $\overline{A}$ means $\neg A$. $\mathcal{L}$ is extended to a conditional language $(\mathcal{L} \mid \mathcal{L})$ by introducing a conditional operator $|$:

$$(\mathcal{L} \mid \mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}.$$

$(\mathcal{L} \mid \mathcal{L})$ is a flat conditional language, no nesting of conditionals is allowed.

Conditionals are usually considered within richer structures such as *epistemic states*. Besides certain knowledge, epistemic states also allow the representation of preferences, beliefs, assumptions etc of an intelligent agent. In a purely qualitative setting, epistemic states can be represented by systems of spheres [Lew73], or simply by a pre-ordering on $\mathcal{L}$ (which is mostly induced by a pre-ordering on worlds). In a (semi-)quantitative setting, also degrees of plausibility, probability, possibility, necessity etc can be expressed. As illustrative examples, we will briefly describe two well-known representations of epistemic states, probability distributions and ordinal conditional functions.

*Probability distributions* in a logical environment can be identified with probability functions $P : \Omega \rightarrow [0,1]$ with $\sum_{\omega \in \Omega} P(\omega) = 1$. The probability of a formula $A \in \mathcal{L}$ is given by $P(A) = \sum_{\omega \models A} P(\omega)$. Conditionals are interpreted via conditional probabilities, so that $P(B|A) = \frac{P(AB)}{P(A)}$ for $P(A) > 0$, and $P \models (B|A)\,[x]$ iff $P(A) > 0$ and $P(B|A) = x$ ($x \in [0,1]$).

*Ordinal conditional functions (OCFs)*, (also called *ranking functions*) $\kappa : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$, were introduced first by Spohn [Spo88]. They express degrees of plausibility of propositional formulas $A$ by specifying degrees of disbeliefs of their negations $\overline{A}$. More formally, we have $\kappa(A) := \min\{\kappa(\omega) \mid \omega \models A\}$, so that $\kappa(A \vee B) = \min\{\kappa(A), \kappa(B)\}$. Hence, due to $\kappa^{-1}(0) \neq \emptyset$, at least one of $\kappa(A), \kappa(\overline{A})$ must be 0. A proposition $A$ is believed if $\kappa(\overline{A}) > 0$ (which implies particularly $\kappa(A) = 0$). 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 $\kappa$, or $\kappa$ *satisfies* $(B|A)$, written as $\kappa \models (B|A)$, iff $\kappa(AB) < \kappa(A\overline{B})$, i.e. iff $AB$ is more plausible than $A\overline{B}$. We can also specify a numerical degree of plausibility of a conditional by defining $\kappa \models (B|A)\,[n]$ iff $\kappa(AB) + n < \kappa(A\overline{B})$ ($n \in \mathbb{N}$). OCF's are the qualitative counterpart of probability distributions. Their plausibility degrees may be taken as order-of-magnitude abstractions of probabilities (cf. [GMP93,GP96]).

Both probability distributions and ordinal conditional functions belong to the class of so-called *conditional valuation functions* which were introduced in [KI01] to abstract from numbers and reveal more clearly and uniformly the way in which (conditional) knowledge may be represented and treated within epistemic states. As an adequate formal structure, we assume an algebra $\mathcal{A} = (\mathcal{A}, \leqslant_{\mathcal{A}}, \oplus, \odot, 0^{\mathcal{A}}, 1^{\mathcal{A}})$ of (real) numbers to be equipped with two operations, $\oplus$ and $\odot$, such that

- $(\mathcal{A}, \oplus)$ is an associative and commutative structure with neutral element $0^{\mathcal{A}}$;
- $(\mathcal{A} - \{0^{\mathcal{A}}\}, \odot)$ is a commutative group with neutral element $1^{\mathcal{A}}$;
- the rule of distributivity holds, i.e. $x \odot (y \oplus z) = (x \odot y) \oplus (x \odot z)$ for $x, y, z \in \mathcal{A}$;
- $\mathcal{A}$ is totally ordered by $\leqslant_{\mathcal{A}}$ with minimum $0^{\mathcal{A}}$ and maximum $1^{\mathcal{A}}$, such that $\leqslant_{\mathcal{A}}$ is compatible with $\oplus$ and $\odot$ in that $x \leqslant_{\mathcal{A}} y$ implies both $x \oplus z \leqslant_{\mathcal{A}} y \oplus z$ and $x \odot z \leqslant_{\mathcal{A}} y \odot z$ for all $x, y, z \in \mathcal{A}$.

So $\mathcal{A}$ is close to be an ordered field, except that the elements of $\mathcal{A}$ need not be invertible with respect to $\oplus$. Conditional valuation functions will make use of two different operations in $\mathcal{A}$ to distinguish between the handling of purely propositional information and conditionals, respectively.

**Definition 1 (conditional valuation function).** *A* conditional valuation function *is a (partial) function $V : \mathcal{L} \cup (\mathcal{L} \mid \mathcal{L}) \to \mathcal{A}$ from the sets of formulas and conditionals into the algebra $\mathcal{A}$ satisfying the following conditions:*

1. *$V|_{\mathcal{L}}$ is total such that $V(\bot) = 0^{\mathcal{A}}, V(\top) = 1^{\mathcal{A}}$, and for exclusive formulas $A, B$ (i.e. $AB \equiv \bot$), it holds that $V(A \vee B) = V(A) \oplus V(B)$;*
2. *for each conditional $(B|A) \in (\mathcal{L} \mid \mathcal{L})$ with $V(A) \neq 0^{\mathcal{A}}$,*

$$V(B|A) = V(AB) \odot V(A)^{-1}$$

*where $V(A)^{-1}$ is the $\odot$-inverse element of $V(A)$ in $\mathcal{A}$; for $V(A) = 0^{\mathcal{A}}$, $V(B|A)$ is undefined.*

Conditional valuation functions assign degrees of certainty, plausibility, possibility etc to propositional formulas and to conditionals. Making use of two operations, they provide a framework for considering and treating conditional knowledge as substantially different from propositional knowledge, a point that is stressed by various authors and that seems to be indispensable for representing epistemic states adequately (cf. [DP97]). There is, however, a close relationship between propositions and conditionals – propositions may be considered as conditionals of a degenerate form by identifying $A$ with $(A|\top)$: Indeed, we have $V(A|\top) = V(A) \odot (1^V)^{-1} = V(A)$. Therefore, conditionals should be regarded as extending propositional knowledge by a new dimension. For further discussions on conditional valuation functions, please see [KI01,KI04].

In the following, let $\Psi$ be any epistemic state, specified e.g. by a preorder, some kind of conditional valuation function, or some other structure that is found appropriate to express conditional beliefs, qualitatively or quantitatively, via a suitable language $(\mathcal{L} \mid \mathcal{L})^*$ and an entailment relation $\models$ between epistemic states and conditionals; basically, $\Psi \models (B|A)^*$ means that $(B|A)^*$ is accepted in $\Psi$. For instance, for probability distributions and ordinal conditional functions, we take $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})^{prob}$ and $(\mathcal{L} \mid \mathcal{L})^{OCF}$, respectively, and in a purely qualitative setting, we assume $(\mathcal{L} \mid \mathcal{L})^* = (\mathcal{L} \mid \mathcal{L})$. Usually, we will assume that epistemic states are *complete* in that they are uniquely characterized by all conditional beliefs which they accept (up to equivalence). Let $Bel(\Psi) \subseteq \mathcal{L}$ the belief set associated to $\Psi$, containing the most plausible (propositional) beliefs.

Let $\mathcal{E}^* = \mathcal{E}^*_{\mathcal{V}}$ denote the set of all such epistemic states using $(\mathcal{L} \mid \mathcal{L})^*$ for representation of (conditional) beliefs. For an epistemic state $\Psi \in \mathcal{E}^*$, we have

$$Th^*(\Psi) = \{\phi \in (\mathcal{L} \mid \mathcal{L})^* \mid \Psi \models \phi\}.$$

In particular, epistemic states are considered as models of sets of conditionals $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$:

$$Mod^*(\mathcal{R}) = \{\Psi \in \mathcal{E}^* \mid \Psi \models \mathcal{R}\}$$

This allows us to extend semantical entailment to sets of conditionals by setting

$$\mathcal{R}_1 \models^* \mathcal{R}_2 \quad \text{iff} \quad Mod^*(\mathcal{R}_1) \subseteq Mod^*(\mathcal{R}_2),$$

and to define a (monotonic) consequence operation $Cn^* : 2^{(\mathcal{L}|\mathcal{L})^*} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^*}$ by

$$Cn^*(\mathcal{R}) = \{\phi \in (\mathcal{L} \mid \mathcal{L})^* \mid \mathcal{R} \models^* \phi\}, \tag{3}$$

in analogy to classical consequence. Two sets of conditionals $\mathcal{R}_1, \mathcal{R}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^*$ are called *(epistemically) equivalent* iff $Mod^*(\mathcal{R}_1) = Mod^*(\mathcal{R}_2)$. As usual, $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ is *consistent* iff $Mod^*(\mathcal{R}) \neq \emptyset$, i.e. iff there is an epistemic state representing $\mathcal{R}$.

Before discussing belief revision and update in a fully epistemic and conditional environment, we will first develop a framework for general, mostly nonmonotonic inference using epistemic states as background knowledge.

## 3 Universal inference operations

From quite a general, abstract point of view, inference operations $C$ map sets of formulas to sets of formulas – given a set of formulas, $C$ is to return which formulas can be derived from this set with respect to some classical or commonsense logic. In this paper, conditionals are assumed to be a basic, most general logical means to express knowledge or beliefs. Therefore, we will consider *(conditional) inference operations*

$$C : 2^{(\mathcal{L}|\mathcal{L})^*} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^*} \tag{4}$$

associating with each set of conditionals a set of inferred conditionals. This also covers the notion of propositional inference operations, taking propositional facts as degenerated conditionals.

A straightforward example for a conditional inference operation is given by the classical operation $Cn^*$ defined by (3) above.

In order to link conditional inference to epistemic states, complete inference operations are of particular interest.

**Definition 2.** *A conditional inference operation $C$ is called* complete *iff it specifies for each set $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ with $C(\mathcal{R}) \neq \emptyset, C(\mathcal{R}) \neq (\mathcal{L} \mid \mathcal{L})^*$, a complete epistemic state $\Psi_{\mathcal{R}}$, i.e. iff there is an epistemic state $\Psi_{\mathcal{R}}$ such that $C(\mathcal{R}) = Th^*(\Psi_{\mathcal{R}})$.*

Complete inference operations realize inductive, model based inference, i.e. they make the information given by some (consistent, non-vacuous) set of conditionals complete.

As we are going to study inference operations which are based on *different* epistemic backgrounds, we make this background explicit and define a formal structure to cover all possible backgrounds, i.e. all epistemic states in $\mathcal{E}^*$.

**Definition 3.** *A* universal inference operation $\mathbf{C}$ *assigns a complete conditional inference operation*

$$C_\Psi : 2^{(\mathcal{L}|\mathcal{L})^*} \rightarrow 2^{(\mathcal{L}|\mathcal{L})^*}$$

*to each epistemic state $\Psi \in \mathcal{E}^*$:*

$$\mathbf{C} : \Psi \mapsto C_\Psi.$$

$\mathbf{C}$ *is said to be* reflexive (idempotent, cumulative) *iff all its involved inference operations have the corresponding property.*

If $\mathbf{C} : \Psi \mapsto C_\Psi$ is a universal inference operation, $C_\Psi$ is complete for each $\Psi \in \mathcal{E}^*$. That means, for each set $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, $C_\Psi(\mathcal{R})$ is either $\emptyset$ or $(\mathcal{L} \mid \mathcal{L})^*$, or it specifies completely (up to equivalence) an epistemic state $\Phi_{\Psi,\mathcal{R}}$:

$$C_\Psi(\mathcal{R}) = Th^*(\Phi_{\Psi,\mathcal{R}}) \tag{5}$$

Define the set of all such epistemic states by

$$\mathcal{E}^*(C_\Psi) = \{\Phi \in \mathcal{E}^* \mid \exists\, \mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^* : C_\Psi(\mathcal{R}) = Th^*(\Phi)\}.$$

A basic property of model-based inference operations is that they produce reasonable outcomes whenever the input set $\mathcal{R}$ is consistent.

**Definition 4.** *A universal inference operation* $\mathbf{C}$ preserves consistency *iff for each epistemic state* $\Psi \in \mathcal{E}^*$ *and for each consistent set* $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, $C_\Psi(\mathcal{R}) \neq \emptyset$ *and* $C_\Psi(\mathcal{R}) \neq (\mathcal{L} \mid \mathcal{L})^*$. *In a quantitative setting, when* $\Psi$ *is represented by a conditional valuation function* $V$, *we further presuppose* $\mathcal{R}$ *to be* $V$*-consistent, i.e.* $\mathcal{R}$ *must respect the* $0^{\mathcal{A}}$*-values in* $V$ *(for a more formal definition of* $V$*-consistency, see [KI01]).*

We are now going to study more interesting properties of universal inference operations.

**Definition 5.** *A universal inference operation* $\mathbf{C}$ *is* founded *iff for each epistemic state* $\Psi$ *and for any* $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, $\Psi \models \mathcal{R}$ *implies* $C_\Psi(\mathcal{R}) = Th^*(\Psi)$.

The property of foundedness establishes a close and intuitive relationship between an epistemic state $\Psi$ and its associated inference operation $C_\Psi$, distinguishing $\Psi$ as its stable starting point. In particular, if $\mathbf{C}$ is founded then $C_\Psi(\emptyset) = Th^*(\Psi)$.

As to the universal inference operation $\mathbf{C}$, foundedness ensures injectivity, as can be proved easily.

**Proposition 1.** *If* $\mathbf{C}$ *is founded, then it is injective.*

In standard nonmonotonic reasoning, as it was developed in [Mak94] and [KLM90], cumulativity occupies a central and fundamental position, claiming the inferences from a set $\mathcal{S}$ that "lies in between" another set $\mathcal{R}$ and its nonmonotonic consequences $C(\mathcal{R})$ to coincide with $C(\mathcal{R})$.

To establish a similar well-behavedness of $\mathbf{C}$ with respect to *epistemic states*, we introduce suitable relations to compare epistemic states with one another.

**Definition 6.** *Let* $\mathbf{C} : \Psi \mapsto C_\Psi$ *be a universal inference operation. For each epistemic state* $\Psi$, *define a relation* $\sqsubseteq_\Psi$ *on* $\mathcal{E}^*(C_\Psi)$ *by setting*

$$\Phi_1 \sqsubseteq_\Psi \Phi_2$$

*iff there are sets* $\mathcal{R}_1 \subseteq \mathcal{R}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^*$ *such that*

$$Th^*(\Phi_1) = C_\Psi(\mathcal{R}_1) \quad \textit{and} \quad Th^*(\Phi_2) = C_\Psi(\mathcal{R}_2)$$

For founded universal inference operations, we have in particular $C_\Psi(\emptyset) = Th^*(\Psi)$ for all $\Psi \in \mathcal{E}^*$, so $\Psi$ is a minimal element of $\mathcal{E}^*(C_\Psi)$ with respect to $\sqsubseteq_\Psi$:

**Proposition 2.** *If* $\mathbf{C}$ *is founded, then for all* $\Psi \in \mathcal{E}^*$ *and for all* $\Phi \in \mathcal{E}^*(C_\Psi)$*, it holds that* $\Psi \sqsubseteq_\Psi \Phi$*.*

We will now generalize the notion of cumulativity for universal inference relations:

**Definition 7.** *A universal inference operation* $\mathbf{C}$ *is called* strongly cumulative *iff for each* $\Psi \in \mathcal{E}^*$ *and for any epistemic states* $\Phi_1, \Phi_2 \in \mathcal{E}^*(C_\Psi)$*,* $\Psi \sqsubseteq_\Psi \Phi_1 \sqsubseteq_\Psi \Phi_2$ *implies: whenever* $\mathcal{R}_1 \subseteq \mathcal{R}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^*$ *such that* $Th^*(\Phi_1) = C_\Psi(\mathcal{R}_1)$ *and* $Th^*(\Phi_2) = C_\Psi(\mathcal{R}_2)$*, then* $Th^*(\Phi_2) = C_\Psi(\mathcal{R}_2) = C_{\Phi_1}(\mathcal{R}_2)$*.*

Strong cumulativity describes a relationship between inference operations based on different epistemic states, thus linking up the inference operations of $\mathbf{C}$. In the definition above, $\Phi_1$ is an epistemic state intermediate between $\Psi$ and $\Phi_2$, with respect to the relation $\sqsubseteq_\Psi$, and strong cumulativity claims that the inferences based on $\Phi_1$ coincide with the inferences based on $\Psi$ within the scope of $\Phi_2$.

The next proposition is immediate:

**Proposition 3.** *Let* $\mathbf{C}$ *be a universal inference operation which is strongly cumulative. Suppose* $\Psi \in \mathcal{E}^*$*,* $\Phi \in \mathcal{E}^*(C_\Psi)$ *such that* $Th^*(\Phi) = C_\Psi(\mathcal{R})$*,* $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$*. Then*

$$C_\Psi(\mathcal{S}) = C_\Phi(\mathcal{S})$$

*for any* $\mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ *with* $\mathcal{R} \subseteq \mathcal{S}$*.*

The following theorem justifies the name "strong cumulativity": It states that strong cumulativity actually generalizes cumulativity for an important class of universal inference operations:

**Theorem 1.** *If* $\mathbf{C}$ *is founded, then strong cumulativity implies cumulativity.*

Universal inference operations will prove to be an adequate formal counterpart of iterated change operations on the other side of the coin (for this metaphor, cf. [Gär92]). Before elaborating this in more detail, we will first develop a comparable formal machinery for belief change.

## 4 Updating epistemic states by conditional beliefs

Basically, from our point of view, the difference between revision and updating is mainly due to different belief change scenarios. Apparently, this is not overwhelmingly new, but elaborating this view in all its consequences has not been done before properly. It implies that the change process should be distinguished from scenario assumptions. So, first of all, we think of $*$ to be a basic imperative change operator that solves our problem (1) successfully:

$$\Psi * \mathcal{R} \in \mathcal{E}^* \text{ such that } \Psi * \mathcal{R} \models \mathcal{R}. \tag{6}$$

It is crucial to point out that $*$ is a full change operator taking *two* entries, namely an epistemic state $\Psi$ on its left and a (compatible and consistent) set of conditionals on its right. Classical theories in nonmonotonic reasoning and belief revision usually focus on handling its right entry, while considering its left entry – i.e. the theory inferences are based upon – to be given.

Typical situations for updating occur when knowledge about a prior world is to be adapted to more recent information (e.g. a demographic model gained from statistical data of past periods should be brushed up by new data, see, for instance, the Example 2 below).

Beyond success, what can be expected from $*$? Since the epistemic state $\Psi$ is assumed to be completely characterized by its set of accepted conditional beliefs, $\mathcal{R}$ can be consistent with $\Psi$ only in a trivial way, i.e., in case of $\Psi \models \mathcal{R}$.

Due to this completeness of knowledge we demand for representing epistemic states, there is few or no room, respectively, to model ignorance. To check whether new information is consistent with an epistemic state thus generally comes down to check whether this information is already represented. On the other hand, this implies that incorporating $\mathcal{R}$ might change lots of (factual) beliefs. In a probabilistic framework, for instance, changing a distribution so as to assimilate to new information will usually change every single atomic probability (even if the change operation is as simple as conditioning). The propositional AGM view of obeying minimal change principles which are based on set inclusion does not make any sense here. This has been recognized already when dealing with conditional beliefs under revision (cf. e.g. [DP97]). We would rather need some kind of distance measure between epistemic states that helps us finding posterior epistemic states closest to prior epistemic states. In the probabilistic framework, cross entropy [JS83] would be a proper candidate for that, and similar ideas could be realized in other frameworks, too.

But resorting to such quantitative measure is not completely satisfactory when the aim is to model belief change processes in human minds, as no human being is able to calculate such things without the help of machinery. This is not at all to be taken as a general argument against such information measures, as implemented logical models of reasoning might be very different from cognitive human models, but nevertheless successful and helpful realizations outside the human brain. However, there still should be some qualitative close connection between prior and posterior epistemic state. The key idea here is to postulate that the reasoning structures underlying $\Psi$ should also be effective in $\Psi * \mathcal{R}$ to a largest possible degree, that is, if no new information in $\mathcal{R}$ force them to change. Since in our framework, such reasoning structures are implemented by conditionals, we should focus on conserving conditional beliefs, hence following a line of thought similar to that of Darwiche and Pearl in their work on iterated belief revision [DP97]. In several papers, we have made precise how a *principle of conditional preservation* can be realized in different frameworks [KI99,KI02,KI04].

Here in this paper, we will generalize these ideas to propose reasonable postulates for an *ideal operator* $*$ changing epistemic states by sets of conditionals. The attribute *ideal* is to be taken in several meanings: First, as we do not stick to one representation framework, the postulates are quite general. Second, we are willing to accept that not all of them can be fulfilled in any framework, nevertheless suggesting them as postulates a "good" change operator should strive to satisfy. Third, even in frameworks where these postulates do not fully apply, we might think of some ideal change operator that the operator under investigation can be based upon (e.g. as a projection), hence deriving from the postulates reasonable properties of belief change in simpler frameworks.

**Postulates for updating epistemic states by sets of conditionals:**
Let $\Psi$ be an epistemic state, and let $\mathcal{R} \subset (\mathcal{L} \mid \mathcal{L})^*$; let $\Psi * \mathcal{R}$ denote the result of updating $\Psi$ by $\mathcal{R}$:

**(CSR1)** **Success:** $\Psi * \mathcal{R} \models \mathcal{R}$.
**(CSR2)** **Stability:** If $\Psi \models \mathcal{R}$ then $\Psi * \mathcal{R} = \Psi$.
**(CSR3)** **Semantical equivalence:** If $\mathcal{R}_1$ and $\mathcal{R}_2$ are (semantically) equivalent, then $\Psi * \mathcal{R}_1 = \Psi * \mathcal{R}_2$.
**(CSR4)** **Reciprocity:** If $\Psi * \mathcal{R}_1 \models \mathcal{R}_2$ and $\Psi * \mathcal{R}_2 \models \mathcal{R}_1$ then $\Psi * \mathcal{R}_1 = \Psi * \mathcal{R}_2$.
**(CSR5)** **Logical coherence:** $\Psi * (\mathcal{R}_1 \cup \mathcal{R}_2) = (\Psi * \mathcal{R}_1) * (\mathcal{R}_1 \cup \mathcal{R}_2)$.

Postulates (CSR1) implements *success*, a crucial property of an imperative change operator. (CSR2) guarantees *stability* if the information represented by $\mathcal{R}$ is not new but already believed. As *success*, also *stability* might be debatable in cases where we would like new confirming information to have a strengthening effect. Again, this is not a matter of pure change processes, but involves different processes that might better be discussed in a merging framework. (CSR3) means that for the basic change operation, only the semantics of new pieces of information matter. This should, however, not be confused with the idea of *syntax independence* which is sometimes used to indicate refraining from using explicit beliefs in belief bases. (CSR4) states that two changing procedures with respect to sets $\mathcal{R}_1$ and $\mathcal{R}_2$ should result in the same epistemic state if each revision represents the new information of the other. This property is called *reciprocity* in the framework of nonmonotonic logics (cf. [Mak94]) and appears as axiom (U6) in the work of Katsuno and Mendelzon [KM91]. (CSR5) is the only seemingly extraordinary and new axiom here. It demands that adjusting any intermediate epistemic state $\Psi * \mathcal{R}_1$ to the full information $\mathcal{R}_1 \cup \mathcal{R}_2$ should result in the same epistemic state as adjusting $\Psi$ by $\mathcal{R}_1 \cup \mathcal{R}_2$ in one step. The rationale behind this axiom is that if the information about the new world drops in in parts, changing any intermediate state of belief by the full information should result unambigously in a final belief state. So, it guarantees the change process to be *logically coherent*. Listing (CSR5) here may be innovative, but the axiom itself is not new. In fact, it is a set-theoretical version of axiom (C1) in [DP97], and it has proved to be a crucial property for the characterization of probabilistic belief change via cross entropy (see [KI01]), but actually goes back to [SJ81].

(CSR5) clearly goes beyond any classical approach to belief change since it is a postulate on iterated belief change, explicitly linking iterated belief change to one-step belief change. Note that (CSR5) does not claim $(\Psi * \mathcal{R}_1) * \mathcal{R}_2$ and $(\Psi * \mathcal{R}_1) * (\mathcal{R}_1 \cup \mathcal{R}_2)$ to be the same, so it does not boil down iterated belief change to simple belief change. Just to the contrary – these epistemic states will be expected to differ in general, because the first is not supposed to maintain prior contextual information, $\mathcal{R}_1$, whereas the second should do so, according to axiom (CSR1).

As an example of a binary update operator which satisfies all postulates (CSR1-CSR5), we briefly present the probabilistic update operation via the *principle of minimum cross-entropy*.

Given two probability distributions $Q$ and $P$, the *cross-entropy* between them is defined by

$$R(Q, P) = \sum_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)}$$

(with $0 \log \frac{0}{0} = 0$ and $Q(\omega) \log \frac{Q(\omega)}{0} = \infty$ for $Q(\omega) \neq 0$) Cross-entropy is a well-known information-theoretic measure of dissimilarity between two distributions and has been studied extensively (see, for instance, [Csi75,HHJ92,Jay83,Kul68]; for a brief, but informative introduction and further references, cf. [Sho86]; see also [SJ81]). Cross-entropy is also called *directed divergence* since it lacks symmetry, i.e. $R(Q, P)$ and $R(P, Q)$ differ in general, so it is not a metric. But cross-entropy is *positive*, that means we have $R(Q, P) \geqslant 0$, and $R(Q, P) = 0$ iff $Q = P$ (cf. [Csi75,HHJ92,Sho86]).

Consider the probabilistic belief revision problem

**($*_{\mathbf{prob}}$)** Given a (prior) distribution $P$ and some set of probabilistic conditionals $\mathcal{R} = \{(B_1|A_1)\,[x_1], \ldots, (B_n|A_n)\,[x_n]\} \subseteq (\mathcal{L} \mid \mathcal{L})^{prob}$, how should $P$ be modified to yield a (posterior) distribution $P^*$ with $P^* \models \mathcal{R}$?

When solving ($*_{prob}$), the paradigm of *informational economy*, i.e. of minimal loss of information (see [Gär88, p. 49]), is realized in an intuitive way by following the *principle of minimum cross-entropy*

$$\begin{aligned} \min R(Q, P) = \sum_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)} \\ \text{s.t. } Q \text{ is a probability distribution with } Q \models \mathcal{R} \end{aligned} \tag{7}$$

For a distribution $P$ and some set $\mathcal{R}$ of probabilistic conditionals compatible with $P$ (cf. [KI01] for the details) there is a (unique) distribution $P_{ME} = P_{ME}(P, \mathcal{R})$ that satisfies $\mathcal{R}$ and has minimal relative entropy to the prior $P$ (cf. [Csi75]), i.e. $P_{ME}$ solves (7) and thereby ($*_{prob}$). Note that ($*_{prob}$) exceeds the framework of the classical AGM-theory with regard to several aspects: an *epistemic state* ($P$) is to be revised by a *set of conditionals* representing *uncertain knowledge*. The *ME-change operator* $*_{ME}$ is defined by

$$P *_{ME} \mathcal{R} = P_{ME}(P, \mathcal{R}). \tag{8}$$

This operator will serve to illustrate some ideas presented in this paper. It can be shown that $*_{ME}$ satisfies all postulates (CSR1-5) [KI01].

## 5 Universal inference operations and belief update

A straightforward relationship between universal inference operations $\mathbf{C} : \Psi \mapsto C_\Psi$ and binary update operators $*$ can now be established by setting

$$\Psi * \mathcal{R} \equiv \Phi_{\Psi,\mathcal{R}}$$

(cf. (5) above) for $\Psi \in \mathcal{E}^*, \mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, that is

$$C_\Psi(\mathcal{R}) = Th(\Psi * \mathcal{R}). \tag{9}$$

We will use this relationship as a formal vehicle to make correspondences between the postulates for updates stated in section 4 and properties of universal inference operations defined in section 3 explicit.

**Proposition 4.** *Suppose update is being realized via a universal inference operation as in (9). Then the following coimplications hold:*

*(i)* $*$ *satisfies (CSR1) iff* $\mathbf{C}$ *is reflexive.*
*(ii)* $*$ *satisfies (CSR2) iff* $\mathbf{C}$ *is founded.*
*(iii)* $*$ *satisfies (CSR3) iff* $\mathbf{C}$ *satisfies left logical equivalence.*
*(iv) Assuming reflexivity resp. the validity of (CSR1),* $*$ *satisfies (CSR4) iff* $\mathbf{C}$ *is cumulative.*
*(v) Assuming foundedness resp. the validity of (CSR2),* $*$ *satisfies (CSR5) iff* $\mathbf{C}$ *is strongly cumulative.*

The proofs are immediate. From this proposition, a representation result follows in a straightforward manner:

**Theorem 2.** *If* $*$ *is defined by (9), it satisfies all of the postulates (CSR1)-(CSR5) iff the universal inference operation* $\mathbf{C}$ *is reflexive, founded, strongly cumulative and satisfies left logical equivalence.*

So, in particular, the properties of foundedness and strong cumulativity turn out to be crucial to control iterated change operations.

## 6 Belief bases and belief revision

Finally, we will turn our attention to belief revision. The usual view on belief revision is that currently held beliefs are revised by new information, where conflicts are solved in favor of the new information. If we presuppose that the revision agent is neither confused nor deceived or stupid, the observations he makes are correct, and conflicts may only arise between new information and plausibly derived beliefs. So, we are going to distinguish between background knowledge and evidential knowledge, and all evidential information the agent collects about the "static world" under consideration – which defines the relevant *context* – must be consistent.

As epistemic states are not supposed to reveal history, and do not allow one to distinguish between explicit and implicit knowledge, or between generic and evidential knowledge, we first take (properly defined) *belief bases* as primitive representations of epistemic knowledge, from which epistemic states may be derived.

**Definition 8.** *A* belief base *is a pair* $(\Psi, \mathcal{R})$*, where* $\Psi$ *is an epistemic state (*background knowledge*), and* $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ *is a set of conditionals representing* contextual *(or* evidential*)* knowledge*.*

The transition from belief bases to epistemic states is assumed to be achieved by the binary belief change operator $*$:

$$*(\Psi, \mathcal{R}) := \Psi * \mathcal{R} \tag{10}$$

considered in section 4. So, flux of knowledge is modelled quite naturally: Prior knowledge serves as a base for obtaining an adequate full description of the present context which may be used again as background knowledge for further change operations.

In the following, we will develop postulates for revising belief bases $(\Psi, \mathcal{R})$ by new conditional information $\mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$, yielding a new belief base $(\Psi, \mathcal{R}) \circ \mathcal{S}$, in the sense of the AGM-postulates.

Due to distinguishing background knowledge from context information, we are able to compare the knowledge stored in different belief bases:

**Definition 9.** *A pre-ordering* $\sqsubseteq$ *on belief bases is defined by*

$$(\Psi_1, \mathcal{R}_1) \sqsubseteq (\Psi_2, \mathcal{R}_2) \quad \textit{iff} \quad \Psi_1 = \Psi_2 \textit{ and } \mathcal{R}_2 \models^* \mathcal{R}_1$$

$(\Psi_1, \mathcal{R}_1)$ *and* $(\Psi_2, \mathcal{R}_2)$ *are* $\sqsubseteq$-equivalent*,*

$$(\Psi_1, \mathcal{R}_1) \equiv_{\sqsubseteq} (\Psi_2, \mathcal{R}_2),$$

*iff* $(\Psi_1, \mathcal{R}_1) \sqsubseteq (\Psi_2, \mathcal{R}_2)$ *and* $(\Psi_2, \mathcal{R}_2) \sqsubseteq (\Psi_1, \mathcal{R}_1)$*.*

Therefore $(\Psi_1, \mathcal{R}_1) \equiv_{\sqsubseteq} (\Psi_2, \mathcal{R}_2)$ iff $\Psi_1 = \Psi_2$ and $\mathcal{R}_1$ and $\mathcal{R}_2$ are semantically equivalent, i.e. iff both belief bases reflect the same epistemic (background and contextual) knowledge.

The following postulates do not make use of the basic change operation $*$, but are to characterize pure belief base revision by the revision operator $\circ$:

**Postulates for conditional base revision:**

Let $\Psi$ be an epistemic state, and let $\mathcal{R}, \mathcal{R}_1, \mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ be sets of conditionals.

**(CBR1)** $(\Psi, \mathcal{R}) \circ \mathcal{S}$ is a belief *base.*

**(CBR2)** If $(\Psi, \mathcal{R}) \circ \mathcal{S} = (\Psi, \mathcal{R}_1)$ then $\mathcal{R}_1 \models^* \mathcal{S}$.
**(CBR3)** $(\Psi, \mathcal{R}) \sqsubseteq (\Psi, \mathcal{R}) \circ \mathcal{S}$.
**(CBR4)** $(\Psi, \mathcal{R}) \circ \mathcal{S}$ is a minimal belief base (with respect to $\sqsubseteq$) among all belief bases satisfying (PR1)-(PR3).

(CBR1) is the most fundamental axiom and coincides with the demand for *categorical matching* (cf. [GR94]). (CBR2) is the *success postulate* here: the new context information is now represented (up to epistemic equivalence). (CBR3) states that revision should preserve prior knowledge. Thus it is crucial for revision in contrast to update. Finally, (CBR4) is in the sense of *informational economy* (cf. [Gär88]): No unnecessary changes should occur. Admittedly, our postulates are much simpler than those proposed by Hansson (see, for instance, [Han89,Han91], and [GR94, p. 61]). They are, however, not based upon classical logic. So, they are more adequate in the framework of general epistemic states.

The following characterization may be proved easily:

**Theorem 3.** *The revision operator $\circ$ satisfies the axioms (CBR1) – (CBR4) iff*

$$(\Psi, \mathcal{R}) \circ \mathcal{S} \equiv_{\sqsubseteq} (\Psi, \mathcal{R} \cup \mathcal{S}). \tag{11}$$

So, from (CBR1)-(CBR4), other properties of the revision operator also follow in a straightforward manner which are usually found among characterizing postulates:

**Proposition 5.** *Suppose the revision operator $\circ$ satisfies (11). Then it fulfills the following properties:*

*(i)* *If* $\mathcal{R} \models^* \mathcal{S}$, *then* $(\Psi, \mathcal{R}) \circ \mathcal{S} \equiv_{\sqsubseteq} (\Psi, \mathcal{R})$;
*(ii)* *If* $(\Psi_1, \mathcal{R}_1) \sqsubseteq (\Psi_2, \mathcal{R}_2)$ *then* $(\Psi_1, \mathcal{R}_1) \circ \mathcal{S} \sqsubseteq (\Psi_2, \mathcal{R}_2) \circ \mathcal{S}$;
*(iii)* $((\Psi, \mathcal{R}) \circ \mathcal{S}_1) \circ \mathcal{S}_2 \equiv_{\sqsubseteq} (\Psi, \mathcal{R}) \circ (\mathcal{S}_1 \cup \mathcal{S}_2)$,

*where* $(\Psi, \mathcal{R}), (\Psi_1, \mathcal{R}_1), (\Psi_2, \mathcal{R}_2)$ *are belief bases and* $\mathcal{S}, \mathcal{S}_1, \mathcal{S}_2 \subseteq (\mathcal{L} \mid \mathcal{L})^*$.

(i) shows a minimality of change, while (ii) is stated in [Gär88] as a monotonicity postulate. (iii) deals with the handling of non-conflicting iterated revisions.

Here we investigate revision merely under the assumption that the new information is compatible with what is already known. Belief revision based on classical logics is nothing but expansion in this case, and Theorem 3 indeed shows that revision of belief bases should reasonably mean expanding contextual knowledge.

Note that revising a belief base $(\Psi, \mathcal{R})$ by $\mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})^*$ also induces a change of the corresponding belief state $\Psi^* = \Psi * \mathcal{R}$ to $(\Psi^*)' = *((\Psi, \mathcal{R}) \circ \mathcal{S})$. So, revision of epistemic states is realized here by making use of base revision and update, or universal inference operations, respectively. According to Theorem 3, if the involved update operation satisfies (CSR3), i.e. if the (underlying) universal inference operation $\mathbf{C}$ satisfies left logical equivalence, then the only reasonable revision operation (as specified by (CBR1)-(CBR4)) is given on the belief state level by

$$*((\Psi, \mathcal{R}) \circ \mathcal{S}) = \Psi * (\mathcal{R} \cup \mathcal{S}), \tag{12}$$

and therefore,

$$Th^*(*((\Psi, \mathcal{R}) \circ \mathcal{S})) = C_\Psi(\mathcal{R} \cup \mathcal{S})$$

This parallels the result for the classical belief revision theory, with the inference operation $C_\Psi$ replacing the classical consequence operation (cf. [Gär88]).

Nevertheless, we prefer using the more general term "revision" to "expansion" here. For, if we consider the epistemic states generated by the two belief bases $\Psi * \mathcal{R}$ and $*((\Psi, \mathcal{R}) \circ \mathcal{S}) = \Psi * (\mathcal{R} \cup \mathcal{S})$, we see that the epistemic status that $\Psi * \mathcal{R}$ assigns to (conditional) beliefs occurring in $\mathcal{S}$ will normally differ from those in $\Psi$ as well as from those in $\Psi * (\mathcal{R} \cup \mathcal{S})$ with expanded contextual knowledge. So the belief in the conditionals in $\mathcal{S}$ is actually *revised*.

## 7 Revision and update

Investigating belief change in the generalized framework of epistemic states and conditionals allows a deeper insight into the mechanisms underlying the belief change process. As a crucial difference to propositional belief change, it is possible to distinguish between revising *simultaneously* and *successively*: In general, we have

$$\Psi * (\mathcal{R} \cup \mathcal{S}) \neq (\Psi * \mathcal{R}) * \mathcal{S}; \tag{13}$$

instead, we may only postulate strong cumulativity or logical coherence, respectively,

$$\Psi * (\mathcal{R} \cup \mathcal{S}) = (\Psi * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}),$$

which is essentially weaker.

The distinction between revision and update can now be made clear on a conceptual level. Revision is to process pieces of contextual information $\mathcal{R}_1, \ldots, \mathcal{R}_n$ pertaining to the same background knowledge $\Psi$, e.g. the epistemic state of the agent at a given time. Hence, revision should be performed by simultaneous belief change. A proper example for this is the process of making up a consistent mental picture on some event, like e.g. a plane crash or a crime which marks a clear point on the time line. Any reliable information an agent obtains on that event clearly has to be handled on the same level. So, if he receives two pieces of information, $\mathcal{R}_1$ and $\mathcal{R}_2$, on that event, one after the other, then he should revise $\Psi * \mathcal{R}_1$ by $\mathcal{R}_2$ to obtain $\Psi * (\mathcal{R}_1 \cup \mathcal{R}_2)$. On the other hand, update is a kind of successive belief change which is able to override any previously held beliefs, just using the current epistemic state as background knowledge. In the scenario with two pieces of information, $\mathcal{R}_1$ and $\mathcal{R}_2$, coming in one after the other, and $\Psi$ as a starting point, the agent should update $\Psi * \mathcal{R}_1$ by $\mathcal{R}_2$ to obtain $(\Psi * \mathcal{R}_1) * \mathcal{R}_2$, which is, as we pointed out, in general different from $\Psi * (\mathcal{R}_1 \cup \mathcal{R}_2)$. The distinction between update and revision thus becomes primarily a question of making a proper decision for the situation at hand, not of techniques.

*Example 1.* A robot is moving around a room, gathering information about that room via its sensors. Depending on his actual position, it comes to know the shape and size of the room, the position of windows, doors and furnitures. Maybe, according to the task it has to fulfill, it might also be interested in which persons are in the room, and the state of the paper basket. The pieces of information clearly come in one after another, at different time points. Nevertheless, the decision whether to use revision or update here depends on the modelling scenario. If the robot takes his own position into account, it clearly has to use an update procedure to process the pieces of information, since its world (where its position is a crucial parameter) constantly changes. On the contrary, if it abstracts from its actual position, e.g. by noticing objects according to their relative position in the room ("door1 is between window2 and cupboard1"), revision might be the proper candidate, as long as we do not expect the room and its object to change during the time of investigation. Making use of the ideas presented above, even a mixture of revision and update is possible, e.g. revising by information about static aspects of the room and updating by information concerning persons in the room or the state of the paper basket.

This makes clear that in our conceptual framework of belief change, incorporating several pieces of new information into our stock of belief can be achieved in different ways by using the same change operator. This allows a more accurate view on iterated belief change by differentiating between simultaneous and successive revision. This means, having to deal with different pieces of information, the crucial question is not whether one information is more recent than others, but which pieces of information should be considered to be on the same level (which may, but is not restricted to be, of temporal type), or to be more precise, which pieces of information refer to the same context. Basically, informations on the same level are assumed to be compatible with one another, so simple set union will return a consistent set of formulas. Otherwise, more complex merging operations have to be considered. Informations on different levels do not have to be consistent, here later ones may override those on previous levels.

In the sequel, we will try to get a clearer view on formal parallels and differences, respectively, between revision and updating. For an adequate comparison, we have to observe the changes of belief *states* that are induced by revision of belief *bases*. Observing (10), (CBR2) and (CBR3) translate into

**(CBR2')** $*((\Psi, \mathcal{R}) \circ \mathcal{S}) \models \mathcal{S}$.
**(CBR3')** $*((\Psi, \mathcal{R}) \circ \mathcal{S}) \models \mathcal{R}$.

While (CBR2') parallels (CSR1), (CBR3') establishes a crucial difference between revision and updating: revision is supposed to preserve prior knowledge while updating does not, neither in a classical nor in a generalized framework.

The intended effects of revision and updating on a belief state $\Psi * \mathcal{R}$ that is generated by a belief base $(\Psi, \mathcal{R})$ are made obvious by – informally! – writing

$$(\Psi * \mathcal{R}) \circ \mathcal{S} = \Psi * (\mathcal{R} \cup \mathcal{S}) \neq (\Psi * \mathcal{R}) * \mathcal{S} \qquad (14)$$

(cf. (12)). This reveals clearly the difference, but also the relationship between revision and updating: Revising $\Psi * \mathcal{R}$ by $\mathcal{S}$ results in the same state of belief as updating $\Psi$ by (the full contextual information) $\mathcal{R} \cup \mathcal{S}$.

The representation of an epistemic state by a belief base, however, is not unique, different belief bases may generate the same belief state (the same holds for classical belief bases, cf. [Han89], [GR94, p. 48]). So we could not define genuine epistemic revision on belief states, but had to consider belief bases in order to separate background and context knowledge unambigously. It is interesting to observe, however, that the logical coherence property (CSR5) of the basic change operator $*$ ensures at least a convenient independence of revision from background knowledge: If two belief bases $(\Psi_1, \mathcal{R}), (\Psi_2, \mathcal{R})$ with different prior knowledge but the same contextual knowledge give rise to the same belief state

$$\Psi_1 * \mathcal{R} = \Psi_2 * \mathcal{R},$$

then, assuming logical coherence to hold,

$$\begin{aligned} \Psi_1 * (\mathcal{R} \cup \mathcal{S}) &= (\Psi_1 * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}) \\ &= (\Psi_2 * \mathcal{R}) * (\mathcal{R} \cup \mathcal{S}) \\ &= \Psi_2 * (\mathcal{R} \cup \mathcal{S}). \end{aligned}$$

So logical coherence guarantees a particular well-behavedness with respect not only to updating, but also to revision.

In the following example, we will illustrate revision and updating in a probabilistic environment, using the ME operator $*_{ME}$ as implementing a proper change operation.

*Example 2.* A psychologist has been working with addicted people for a couple of years. His experiences concerning the propositions

$\mathcal{V}$ : $a$ : addicted to <u>a</u>lcohol
$d$ : addicted to <u>d</u>rugs
$y$ : being <u>y</u>oung

may be summarized by the following distribution $P$ that expresses his belief state probabilistically:

| $\omega$ | $P(\omega)$ | $\omega$ | $P(\omega)$ | $\omega$ | $P(\omega)$ | $\omega$ | $P(\omega)$ |
|---|---|---|---|---|---|---|---|
| $ady$ | 0.050 | $\overline{a}dy$ | 0.333 | $ad\overline{y}$ | 0.053 | $\overline{a}d\overline{y}$ | 0.053 |
| $a\overline{d}y$ | 0.093 | $\overline{a}\overline{d}y$ | 0.102 | $a\overline{d}\overline{y}$ | 0.225 | $\overline{a}\overline{d}\overline{y}$ | 0.091 |

The following probabilistic conditionals may be entailed from $P$:

$$\begin{array}{ll} (d|a)[0.242] & (\text{i.e.} \quad P(d|a) = 0.242) \\ (d|\overline{a})[0.666] & (\text{i.e.} \quad P(d|\overline{a}) = 0.666) \\ (a|y)[0.246] & (a|\overline{y})[0.660] \\ (d|y)[0.662] & (d|\overline{y})[0.251] \end{array}$$

Now the psychologist is going to change his job: He will be working in a clinic for people addicted only to alcohol and/or drugs. He is told that the percentage of persons addicted to alcohol, but also addicted to drugs, is higher than usual and may be estimated by 40 %.

So the information the psychologist has about the "new world" is represented by

$$\mathcal{R} = \{a \vee d[1], (d|a)[0.4]\}.$$

The distribution $P$ from above is now supposed to represent background or prior knowledge, respectively. So the psychologist updates $P$ by $\mathcal{R}$ using ME-change and obtains $P^* = P *_{ME} \mathcal{R}$ as new belief state:

| $\omega$ | $P^*(\omega)$ | $\omega$ | $P^*(\omega)$ | $\omega$ | $P^*(\omega)$ | $\omega$ | $P^*(\omega)$ |
|---|---|---|---|---|---|---|---|
| $ady$ | 0.099 | $\overline{a}dy$ | 0.425 | $ad\overline{y}$ | 0.105 | $\overline{a}d\overline{y}$ | 0.066 |
| $a\overline{d}y$ | 0.089 | $\overline{a}\overline{d}y$ | 0.0 | $a\overline{d}\overline{y}$ | 0.216 | $\overline{a}\overline{d}\overline{y}$ | 0.0 |

After having spent a couple of days in the new clinic, the psychologist realized that this clinic was for young people only. So he had to revise his knowledge about his new sphere of activity and arrived at the revised belief state $*_{ME}((P, \mathcal{R}) \circ y[1]) = P *_{ME} (\mathcal{R} \cup y[1]) =: P_1^*$ shown in the following table:

| $\omega$ | $P_1^*(\omega)$ | $\omega$ | $P_1^*(\omega)$ | $\omega$ | $P_1^*(\omega)$ | $\omega$ | $P_1^*(\omega)$ |
|---|---|---|---|---|---|---|---|
| $ady$ | 0.120 | $\overline{a}dy$ | 0.700 | $ad\overline{y}$ | 0.0 | $\overline{a}d\overline{y}$ | 0.0 |
| $a\overline{d}y$ | 0.180 | $\overline{a}\overline{d}y$ | 0.0 | $a\overline{d}\overline{y}$ | 0.0 | $\overline{a}\overline{d}\overline{y}$ | 0.0 |

This distribution obtained by *revision* is different from that one the psychologist would have obtained by *focusing* his knowledge represented by $P^* = P *_{ME} \mathcal{R}$ on a young patient, which can be computed via ME update as $P^* *_{ME} \{y[1]\} = P^*(\cdot|y) =: P_y^*$ (please note that the context has changed!):

| $\omega$ | $P_y^*(\omega)$ | $\omega$ | $P_y^*(\omega)$ | $\omega$ | $P_y^*(\omega)$ | $\omega$ | $P_y^*(\omega)$ |
|---|---|---|---|---|---|---|---|
| $ady$ | 0.162 | $\overline{a}dy$ | 0.693 | $ad\overline{y}$ | 0.0 | $\overline{a}d\overline{y}$ | 0.0 |
| $a\overline{d}y$ | 0.145 | $\overline{a}\ \overline{d}y$ | 0.0 | $a\overline{d}\overline{y}$ | 0.0 | $\overline{a}\overline{d}\overline{y}$ | 0.0 |

## 8 Conclusion

In this paper, we presented a formal, unifying framework for iterated belief change and nonmonotonic inference operations with explicit epistemic background knowledge. We made a clear conceptual distinction between belief revision and belief update and showed how this distinction can be realized technically. Both revision and update operations could be iterated, with different outcomes. General postulates served as cornerstones to classify change operations, in the tradition of the AGM-theory. Most of these postulates were related to similar statements in revision and update theories. The postulate (CSR5) of Logical Coherence, however, is new and refers explicitly to iterated belief change.

It proved to be very strong, and is shown to correspond to strong cumulativity of universal inference operations.

We illustrated our ideas by the probabilistic change operation that is induced by the principle of minimum cross entropy. Up to date, it is the only change operation which is known to satisfy all requirements listed in this paper. It is part of our ongoing work to explore which other frameworks are rich enough to comply with the postulates and properties presented in this paper.

## References


[AGM85] 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.

[Csi75] I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Ann. Prob.*, 3:146–158, 1975.

[DP96] D. Dubois and H. Prade. Non-standard theories of uncertainty in plausible reasoning. In G. Brewka, editor, *Principles of Knowledge Representation.* CSLI Publications, 1996.

[DP97] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge, Mass., 1988.

[Gär92] P. Gärdenfors. Belief revision and nonmonotonic logic: Two sides of the same coin? In *Proceedings European Conference on Artificial Intelligence, ECAI'92*, pages 768–773. Pitman Publishing, 1992.

[GMP93] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):220–232, 1993.

[GP96] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84:57–112, 1996.

[GR94] P. Gärdenfors and H. Rott. Belief revision. In D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, pages 35–132. Oxford University Press, 1994.

[Han89] S.O. Hansson. New operators for theory change. *Theoria*, 55:114–132, 1989.

[Han91] S.O. Hansson. *Belief base dynamics.* PhD thesis, Uppsala University, 1991.

[HHJ92] P. Hajek, T. Havranek, and R. Jirousek. *Uncertain Information Processing in Expert Systems.* CRC-Press, Boca Raton, Florida, 1992.

[Jay83] E.T. Jaynes. *Papers on Probability, Statistics and Statistical Physics.* D. Reidel Publishing Company, Dordrecht, Holland, 1983.

[JS83] R.W. Johnson and J.E. Shore. Comments on and correction to "Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy". *IEEE Transactions on Information Theory*, IT-29(6):942–943, 1983.

[KI99] G. Kern-Isberner. Postulates for conditional belief revision. In *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI-99*, pages 186–191. Morgan Kaufmann, 1999.

[KI01] G. Kern-Isberner. *Conditionals in nonmonotonic reasoning and belief revision.* Springer, Lecture Notes in Artificial Intelligence LNAI 2087, 2001.

[KI02] G. Kern-Isberner. The principle of conditional preservation in belief revision. In *Proceedings of the Second International Symposium on Foundations of Information and Knowledge Systems, FoIKS 2002*, pages 105–129. Springer LNCS 2284, 2002.

[KI04] 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.

[KLM90] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[KM91] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings Second International Conference on Principles of Knowledge Representation and Reasoning, KR'91*, pages 387–394, San Mateo, Ca., 1991. Morgan Kaufmann.

[Kul68] S. Kullback. *Information Theory and Statistics*. Dover, New York, 1968.

[Lew73] D. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Mass., 1973.

[Mak94] D. Makinson. General patterns in nonmonotonic reasoning. In D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 35–110. Oxford University Press, 1994.

[MG91] D. Makinson and P. Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In *Proceedings Workshop The Logic of Theory Change, Konstanz, Germany, 1989*, pages 185–205, Berlin Heidelberg New York, 1991. Springer.

[Ram50] 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.

[Sho86] J.E. Shore. Relative entropy, probabilistic inference and AI. In L.N. Kanal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 211–215. North-Holland, Amsterdam, 1986.

[SJ81] J.E. Shore and R.W. Johnson. Properties of cross-entropy minimization. *IEEE Transactions on Information Theory*, IT-27:472–482, 1981.

[Spo88] 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.