

Belief Change in a Preferential Non-Monotonic Framework

Giovanni Casini
Université du Luxembourg
Luxembourg
giovanni.casini@uni.lu

Thomas Meyer
CAIR-CSIR
University of Cape Town
South Africa
tmeyer@cs.uct.ac.za

Abstract

Belief change and non-monotonic reasoning are usually viewed as two sides of the same coin, with results showing that one can formally be defined in terms of the other. In this paper we show that we can also integrate the two formalisms by studying belief change within a (preferential) non-monotonic framework. This integration relies heavily on the identification of the *monotonic core* of a non-monotonic framework. We consider belief change operators in a non-monotonic propositional setting with a view towards preserving consistency. These results can also be applied to the preservation of coherence—an important notion within the field of logic-based ontologies. We show that the standard AGM approach to belief change can be adapted to a preferential non-monotonic framework, with the definition of expansion, contraction, and revision operators, and corresponding representation results. Surprisingly, preferential AGM belief change, as defined here, can be obtained in terms of classical AGM belief change.

1 Introduction

Both belief change and non-monotonic reasoning frameworks provide mechanisms for dealing with conflicting information. For example, suppose I know that vertebrate red blood cells have a nucleus ($v \rightarrow n$), that mammalian red blood cells are vertebrate red blood cells ($m \rightarrow v$), but that mammalian red blood cells don't have a nucleus ($m \rightarrow \neg n$). The existence of mammalian red blood cells (m) then renders my knowledge base inconsistent. Belief change operators modify the existing knowledge base to make it consistent again, while non-monotonic reasoning typically resolves the issue by introducing a notion of *defeasibility*—for example, by being able to state that vertebrate red blood cells *usually* have a nucleus.

In this paper our aim is to integrate the two frameworks by studying belief change in a non-monotonic setting. This approach has not received much attention thus far, mostly for two reasons. Firstly, the standard approaches to belief change usually assume an underlying Tarskian consequence relation which is explicitly monotonic [Alchourrón *et al.*, 1985]. Secondly, at a first glance belief change seems to be superfluous

in non-monotonic settings—a revision operator can simply be replaced with expansion, with the non-monotonic machinery then ensuring consistency of some kind. What this simplistic view does not take into account is that non-monotonic frameworks contain a mix of defeasible and classical (non-defeasible) information. In our example above, for instance, the statement that mammalian red blood cells are vertebrate red blood cells should probably not be defeasible, while the statement about vertebrate red blood cells having a nucleus probably should be. The challenge thus becomes one of defining belief change for the *monotonic* part of the formalism, while simultaneously ensuring that the non-monotonic part remains well-behaved.

In addition to *consistency preservation* we also consider belief change for the preservation of a restricted version of *coherence*, as it is intended in the field of logic-based ontologies, that is, that every class that has been introduced in the language can in principle be populated [Qi and Hunter, 2007]. In our example above the (non-defeasible) statements $\{v \rightarrow n, m \rightarrow v, m \rightarrow \neg n\}$ cause the knowledge base to be incoherent w.r.t. m (there cannot be any mammalian red blood cells), but the knowledge base only becomes inconsistent when m is added. For our purposes here it is sufficient to note that, in a propositional setting, the incoherence of a knowledge base w.r.t. an atom (or any propositional formula A , for that matter), corresponds to the statement $\neg A$ being a consequence of our knowledge base.

We focus on belief change in the preferential approach to non-monotonic reasoning developed by Kraus, Lehmann and Magidor [1990] in which *defeasible* conditionals of the form $C \sim D$ are added to the language of propositional logic (with C and D being classical propositional formulas). In this enriched language the information about vertebrate blood cells usually having a nucleus could be expressed as $v \sim n$.

In Section 2 we briefly summarise the AGM framework and preferential reasoning. Section 3 describes the non-monotonic framework we consider and introduces the notion of a *monotonic core*. Section 4 analyses a contraction operator and two revision operators for the monotonic core. Section 5 investigates the connections between contraction of the monotonic core and classical propositional contraction. Section 6 analyses two revision operators for non-monotonic entailment. Section 7 is the concluding section.

2 Background

We consider a finitely generated propositional language \mathcal{L} with lower case letters denoting atoms, and capital letters denoting elements of \mathcal{L} . We adopt the standard semantics for propositional logic. The set V of valuations v are functions from the set of atoms in \mathcal{L} to $\{0, 1\}$ denoting truth and falsity. Satisfaction is denoted by \models , and entailment by \models .

AGM Belief Change. AGM belief change assumes an underlying logic with a propositional language and a Tarskian entailment relation Cn that is compact and satisfies *disjunction in the premises* ([Alchourrón et al., 1985], p.511-512). A knowledge base K is assumed to be a set of formulas *closed under entailment*. AGM is concerned with three types of operators: expansion, contraction, and revision. Expansion is simply defined as adding a formula and closing under entailment: $K + A = Cn(K \cup \{A\})$. The intuition associated with the contraction of K by A is that it should result in a knowledge base $K - A$ in which A is not entailed. Dually, a revision of K by A should result in a consistent knowledge base $K * A$ in which A is entailed. Formally, AGM contraction is required to satisfy the following basic postulates:

- (-1) $K_A^- = Cn(K_A^-)$ (*- closure*)
- (-2) $K_A^- \subseteq K$ (*- inclusion*)
- (-3) If $A \notin K$, then $K_A^- = K$ (*- vacuity*)
- (-4) If $\not\models A$, then $A \notin K_A^-$ (*- success*)
- (-5) If $A \equiv B$ then $K_A^- = K_B^-$ (*- extensionality*)
- (-6) $K \subseteq (K_A^-)^+$ (*- recovery*)

AGM contraction operators can be constructed using so-called *partial meet functions* ([Alchourrón et al., 1985], Observation 2.5). Let $K \downarrow A$ be the *remainder set*, containing the maximal subsets K' of K s.t. $A \notin K'$. That is, $K' \in K \downarrow A$ iff (i) $K' \subseteq K$, (ii) $A \notin K'$, and (iii) there is no set K'' s.t. $K' \subset K'' \subseteq K$ and $A \notin K''$. Let pm be a partial meet function defined over $K \downarrow A$ s.t. $pm(K \downarrow A) \subseteq K \downarrow A$ and if $K \downarrow A \neq \emptyset$, then $pm(K \downarrow A) \neq \emptyset$. A partial meet contraction operator $-$ is defined as: $K_A^- = \bigcap pm(K \downarrow A)$.

Theorem 1 [Alchourrón et al., 1985] *A contraction operator is an AGM contraction operator (satisfying (-1)-(-6)) iff it is a partial meet operator.*

Similarly, AGM revision is required to satisfy the following postulates:

- (*1) $K_A^* = Cn(K_A^*)$ (** closure*)
- (*2) $K_A^* \subseteq K_A^+$ (** inclusion*)
- (*3) If $\neg A \notin K$ then $K \subseteq K_A^+ \subseteq K_A^*$ (** vacuity*)
- (*4) $A \in K_A^*$ (** success*)
- (*5) If $A \equiv B$ then $K_A^* = K_B^*$ (** extensionality*)
- (*6) If $A \not\models \perp$ then $K_A^* \not\models \perp$ (** consistency*)

It is well-known that revision can be defined in terms of contraction and expansion via the Levi Identity [Levi, 1977]: $K_A^* = (K_{-A}^-)^+$.

Theorem 2 [Alchourrón et al., 1985] *Revision defined in terms of a contraction operator – via the Levi Identity is an AGM revision operator iff – is an AGM contraction operator.*

Intuitively the Levi Identity states that, to revise a knowledge base with A , we first contract with its negation, thereby “making room” for A , and then simply expand with A .

Preferential Entailment. To introduce defeasibility we consider the language \mathcal{L}^\sim consisting of defeasible conditionals of the form $A \sim B$ (for $A, B \in \mathcal{L}$), that are read as ‘Typically, if A then B ’. The semantics of \mathcal{L}^\sim is based on the notion of a *preferential interpretation*: triples of the form $\langle W, l, < \rangle$ where W is a set of objects (*states*), l is a function from W to V (mapping states into valuations), and $<$ is a strict partial ordering on W [Kraus et al., 1990]. $w < v$ is interpreted as indicating that the state w represents a more typical situation than the state v . For $A \in \mathcal{L}$, we let $[A]^W = \{w \in W \mid l(w) \models A\}$. A defeasible conditional $A \sim B$ is satisfied in a preferential interpretation \mathcal{P} , denoted as $\mathcal{P} \models A \sim B$, iff $\min_{<} [A]^W \subseteq [B]^W$. Observe that a propositional formula A is satisfied in all the elements of W (i.e., $l(w) \models A$ for every $w \in W$) iff $\mathcal{P} \models \neg A \sim \perp$. This means that any classical propositional formula A can be represented as the defeasible conditional $\neg A \sim \perp$. Indeed every conditional of the form $A \sim \perp$ is *not* defeasible, and actually represents classical propositional information. E.g., the statement that vertebrate red blood cells have a nucleus ($v \rightarrow n$) will be represented as $\neg(v \rightarrow n) \sim \perp$. Because of this we sometimes abuse notation by referring to $\neg A \sim \perp$ as the propositional formula A .

Let \mathcal{B} indicate a finite set of defeasible conditionals. The set of *preferential models* of \mathcal{B} , preferential interpretations satisfying \mathcal{B} , is denoted by $[[\mathcal{B}]]$. The obvious notion of Tarskian entailment associated with this semantics is known as *preferential entailment* [Lehmann and Magidor, 1992], represented as \models_{pr} , where $\mathcal{B} \models_{pr} A \sim B$ iff $[[\mathcal{B}]] \subseteq [[\{A \sim B\}]]$. The closure operator associated with preferential entailment is known as *preferential closure*: $\mathcal{P}(\mathcal{B}) = \{A \sim B \mid \mathcal{B} \models_{pr} A \sim B\}$. We use the notation $\mathcal{K}, \mathcal{K}'$ etc. to refer to conditional *knowledge bases* closed under preferential entailment. \mathcal{B} is *preferentially inconsistent* iff $\mathcal{B} \models_{pr} \top \sim \perp$, and \mathcal{B} is *preferentially incoherent* w.r.t. a propositional formula A iff $\mathcal{B} \models_{pr} A \sim \perp$. Two bases are *preferentially equivalent* iff they have the same preferential models.

In what follows below, the notion of *exceptionality* and the rank of a formula (and a defeasible conditional) is important ([Lehmann and Magidor, 1992], Section 2.6). $A \in \mathcal{L}$ is said to be *exceptional* w.r.t. \mathcal{B} iff $\mathcal{B} \models_{pr} \top \sim \neg A$ (that is, in every model of \mathcal{B} A is falsified in all the most typical states), and $A \sim B \in \mathcal{L}^\sim$ is said to be *exceptional* w.r.t. \mathcal{B} iff A is exceptional w.r.t. \mathcal{B} . We make use of exceptionality to provide a notion of rank for defeasible conditionals. $\mathcal{E}_0^\mathcal{B} = \mathcal{B}$ and for $i \geq 0$, $\mathcal{E}_i^\mathcal{B} = \{A \sim B \in \mathcal{E}_{i-1}^\mathcal{B} \mid A \sim B \text{ is exceptional w.r.t. } \mathcal{E}_{i-1}^\mathcal{B}\}$. There is a smallest integer, say n , such that $\mathcal{E}_n^\mathcal{B} = \mathcal{E}_{n+1}^\mathcal{B}$ (it may be the case that $\mathcal{E}_n^\mathcal{B} = \emptyset$). The *rank* of $A \in \mathcal{L}$ w.r.t. \mathcal{B} , denoted as $r_\mathcal{B}(A)$, is the smallest integer k for which A is not exceptional w.r.t. $\mathcal{E}_k^\mathcal{B}$. If A is exceptional w.r.t. $\mathcal{E}_i^\mathcal{B}$ for all i , then $r_\mathcal{B}(A) = \infty$. The rank of a defeasible conditional $A \sim B$ w.r.t. \mathcal{B} , denoted as $r_\mathcal{B}(A \sim B)$, is equal to $r_\mathcal{B}(A)$. Intuitively, the rank of a defeasible conditional provides an indication of the extent to which we would be willing to override it: the lower the

rank, the more likely we are to discard it. Conditionals with rank ∞ will never be overridden and can thus be viewed as classical assertions w.r.t. \mathcal{B} . For $i \in \{0, \dots, n-1, \infty\}$ we let $\mathcal{B}^i = \{A \sim B \in \mathcal{B} \mid r_{\mathcal{B}}(A \sim B) = i\}$, and we let $\mathcal{B}^{<\infty} = \mathcal{B} \setminus \mathcal{B}^\infty$. From a computational perspective, it is useful to note that exceptionality checking can be reduced to classical propositional reasoning. The *materialisation* $\overrightarrow{\mathcal{B}}$ of \mathcal{B} is defined as follows: $\overrightarrow{\mathcal{B}} = \{A \rightarrow B \mid A \sim B \in \mathcal{B}\}$. It is known that exceptionality w.r.t. \mathcal{B} ($\mathcal{B} \models_{pr} \top \sim \neg A$) can be reduced to checking whether $\overrightarrow{\mathcal{B}} \models \neg A$ ([Lehmann and Magidor, 1992], Section 5.8). In particular, \mathcal{B}^∞ is the biggest subset of \mathcal{B} s.t. $\mathcal{B}^\infty = \{A \sim B \in \mathcal{B} \mid \overrightarrow{\mathcal{B}^\infty} \models \neg A\}$, and, for every $A \in \mathcal{L}$, $A \sim \perp \in \mathcal{P}(\mathcal{B})$ iff $\overrightarrow{\mathcal{B}^\infty} \models \neg A$. All the notions defined above for conditional bases \mathcal{B} also apply to every knowledge base with the property of having a well-founded model. Base-generated knowledge bases have such a property ([Lehmann and Magidor, 1992], Lemma 2.24); it can be proved that the operators defined here preserve it as well.

3 The Monotonic Core

It is widely recognised that preferential entailment is too weak to be an appropriate form of entailment for a non-monotonic framework ([Kraus *et al.*, 1990], pp. 4,34). This is primarily because the preferential entailment relation itself is monotonic (non-monotonicity occurs on the object-level, within defeasible conditionals). At the same time there is sufficient consensus that any acceptable form of non-monotonic entailment will be an extension of preferential entailment [Kraus *et al.*, 1990; Lehmann and Magidor, 1992]. Our final focus in this paper is on such an extended set of non-monotonic entailment relations. Below we make this idea more precise by defining the class of *supra-preferential cumulative* entailment relations.

Consider a closure operator \mathcal{C} , which is defined as follows: $A \sim B \in \mathcal{C}(\mathcal{B})$ iff $\mathcal{P} \Vdash A \sim B$ for every preferential model $\mathcal{P} \in c(\mathcal{B})$, where c is a choice function s.t. (i) $c(\mathcal{B}) \subseteq [[\mathcal{B}]]$, (ii) if $[[\mathcal{B}]] \neq \emptyset$ then $c(\mathcal{B}) \neq \emptyset$, and (iii) $c(\mathcal{B}) = c(c(\mathcal{B}))$. The closure operators definable using such choice functions are the ones that extend preferential closure and respect *consistency preservation* (if $\top \sim \perp \notin \mathcal{P}(\mathcal{B})$, then $\top \sim \perp \notin \mathcal{C}(\mathcal{B})$), and we will indicate them as *supra-classical*. We also require *cumulativity*: for every $\mathcal{B}, \mathcal{B}'$, if $\mathcal{B} \subseteq \mathcal{B}' \subseteq \mathcal{C}(\mathcal{B})$, then $\mathcal{C}(\mathcal{B}') = \mathcal{C}(\mathcal{B})$. This is captured by a choice function c as follows: if $c(\mathcal{B}) \subseteq [[\mathcal{B}']] \subseteq [[\mathcal{B}]]$, then $c(\mathcal{B}') = c(\mathcal{B})$. We refer to the closure operators satisfying these properties as *spc*-operators (as *supra-preferential cumulative*). Most of the prominent non-monotonic closure operators proposed in the preferential framework (e.g. [Kraus *et al.*, 1990; Lehmann and Magidor, 1992; Lehmann, 1995; Casini and Straccia, 2013]) are *spc*-operators. Our goal is to analyse belief change for the class of *spc*-operators.

As mentioned in Section 1, belief change in a non-monotonic framework seems superfluous at a first glance, since the non-monotonic machinery usually takes care of ‘readjusting’ the inferences in order to preserve consistency and coherence facing new unexpected evidence. But that’s not always the case. Let’s look at some examples.

Example 1 Consider the example about the red blood cells, but adding the information that avian red blood cells are vertebrate red blood cells ($a \rightarrow v$). Let our base contain the conditional counterpart of the classical formalisation. I.e., $\mathcal{B} = \{\neg(v \rightarrow n) \sim \perp, \neg(m \rightarrow v) \sim \perp, \neg(m \rightarrow \neg n) \sim \perp, \neg(a \rightarrow v) \sim \perp\}$. Consider the following situations:

1. $m \sim \perp$ is in $\mathcal{C}(\mathcal{B})$. But mammalian red blood cells exist, and we want to enforce such information. In propositional belief change we would remove some piece of information, presumably either $m \rightarrow v$, $v \rightarrow n$, or $m \rightarrow \neg n$. In the framework we propose we can resolve the situation by introducing defeasibility. Contracting $m \sim \perp$ from \mathcal{B} we would like to end up, for example, with $\mathcal{B}' = \{v \sim n, \neg(m \rightarrow v) \sim \perp, \neg(m \rightarrow \neg n) \sim \perp, \neg(a \rightarrow v) \sim \perp\}$ in which, instead of eliminating $\neg(v \rightarrow n) \sim \perp$, we have just made it defeasible.

2. If \mathcal{C} is a well-behaved non-monotonic closure operator, we should have $a \sim n \in \mathcal{C}(\mathcal{B}')$, since, with the information we have, we can treat avian cells as typical vertebrate cells. Now assume we are informed that $a \sim \neg n$: $\mathcal{B}'' = \{v \sim n, \neg(m \rightarrow v) \sim \perp, \neg(m \rightarrow \neg n) \sim \perp, \neg(a \rightarrow v) \sim \perp, a \sim \neg n\}$. In this case, since $a \sim n$ is a presumptive conclusion made by the non-monotony machinery, the entailment relation itself will take care of eliminating such a conclusion once faced with conflicting evidence. In such a case we want the introduction of $a \sim \neg n$ to correspond to a simple expansion.

3. We are then informed that $a \sim n$ actually holds, information that directly conflicts with $a \sim \neg n$ in our base. This kind of conflict cannot be managed by the non-monotonic machinery, since $a \sim \neg n$ is a (trivial) necessary conclusion from \mathcal{B}'' . So, we would be forced to conclude that avian red blood cell do not exist ($a \sim \perp$). We have two choices: either we are interested just in general consistency (not deriving $\top \sim \perp$), and in such a case the addition of $a \sim n$ is just an expansion and we conclude that birds do not have red blood cells. Or we perform a revision in order to ‘make room’ for $a \sim n$ without being forced to conclude $a \sim \perp$.

As shown in the examples, in order to preserve consistency and coherence we need to distinguish between what follows necessarily from the information at our disposal and what the inference operator concludes ‘presumptively’. That is, we need to identify the monotonic part of our reasoning and operate only on it, while for the non-monotonic part the entailment relation should automatically take care of the appropriate adjustments. A closure operator $\mathcal{C}l$ is called the *monotonic core* of a non-monotonic operator \mathcal{C} if, for every $\mathcal{B}, \mathcal{B}'$, (i) $\mathcal{B} \subseteq \mathcal{B}'$ implies $\mathcal{C}l(\mathcal{B}) \subseteq \mathcal{C}l(\mathcal{B}')$; (ii) $\mathcal{C}l(\mathcal{B}) \subseteq \mathcal{C}(\mathcal{B})$; (iii) for every closure operator \mathcal{C}' satisfying (i) and (ii), $\mathcal{C}'(\mathcal{B}) \subseteq \mathcal{C}l(\mathcal{B})$. Given a non-monotonic entailment relation, the existence of a monotonic core needs to be proved.

Proposition 1 *The monotonic core of a supra-preferential closure operator is the preferential closure \mathcal{P} .*

Proof: (Sketch). Assume that for some $A \sim B \in \mathcal{C}(\mathcal{B}) \setminus \mathcal{P}(\mathcal{B})$ there is a monotonic \mathcal{C}' contained in \mathcal{C} s.t. $A \sim B \in \mathcal{C}'(\mathcal{B})$, and obtain a contradiction. \square

Hence, the monotonic core of the class of *spc*-operators is the preferential closure.

4 Preferential Belief Change

In this section we start our investigation with a study of belief change for the monotonic core of *spc*-operators—that is, belief change for preferential entailment. We apply an AGM-like approach to the preferential closure \mathcal{K} of defeasible knowledge bases \mathcal{B} . Expansion is then defined as: $\mathcal{K}_{A \sim B}^+ = \mathcal{P}(\mathcal{K} \cup \{A \sim B\})$. The postulates for contraction are as follows (where $\equiv_{\mathcal{P}}$ refers to preferential equivalence):

- (\div 1) $\mathcal{K}_{A \sim B}^{\div} = \mathcal{P}(\mathcal{K}_{A \sim B}^{\div})$ (\div closure)
- (\div 2) $\mathcal{K}_{A \sim B}^{\div} \subseteq \mathcal{K}$ (\div inclusion)
- (\div 3) If $\mathcal{K} \not\models_{pr} A \sim B$, then $\mathcal{K}_{A \sim B}^{\div} = \mathcal{K}$ (\div vacuity)
- (\div 4) If $\not\models_{pr} A \sim B$, then $A \sim B \notin \mathcal{K}_{A \sim B}^{\div}$ (\div success)
- (\div 5) If $A \sim B \equiv_{\mathcal{P}} A' \sim B'$, then $\mathcal{K}_{A \sim B}^{\div} = \mathcal{K}_{A' \sim B'}$ (\div ext.)
- (\div 6) $\mathcal{K} \subseteq (\mathcal{K}_{A \sim B}^{\div})_{A \sim B}^+$ (\div recovery)

Partial meet functions for preferential entailment are defined as expected. $\mathcal{K}' \in (\mathcal{K} \downarrow A \sim B)$ iff (i) $\mathcal{K}' \subseteq \mathcal{K}$, (ii) $A \sim B \notin \mathcal{K}'$, and (iii) there is no \mathcal{K}'' s.t. $\mathcal{K}' \subset \mathcal{K}'' \subseteq \mathcal{K}$ and $A \sim B \notin \mathcal{K}''$. Let *pm* be a partial meet function defined over $(\mathcal{K} \downarrow A \sim B)$ s.t. $pm(\mathcal{K} \downarrow A \sim B) \subseteq (\mathcal{K} \downarrow A \sim B)$ and if $(\mathcal{K} \downarrow A \sim B) \neq \emptyset$, then $pm(\mathcal{K} \downarrow A \sim B) \neq \emptyset$. A partial meet contraction operator \div is defined as: $\mathcal{K}_{A \sim B}^{\div} = \bigcap pm(\mathcal{K} \downarrow A \sim B)$. And as expected, we have a representation result, linking the postulates and partial meet functions.

Theorem 3 *A contraction operator \div for preferential entailment satisfies (\div 1) to (\div 6) iff it is a partial meet contraction operator.*

The proof of this result is postponed to the end of Section 5. Turning now to coherence preservation, observe firstly that this can be attained using preferential contraction. More precisely, to ensure coherence of \mathcal{K} w.r.t. a formula A , we perform a contraction of \mathcal{K} by $A \sim \perp$ (see Example 1.1).

Moving to revision, in the present framework we can define two kinds of revision: one to preserve consistency (adding $A \sim B$ while avoiding to conclude $\top \sim \perp$) and one to preserve coherence (adding $A \sim B$ while avoiding to conclude $A \sim \perp$, see Example 1.3). Regarding the former, we will make use of the following result.

Proposition 2 $\mathcal{K} \cup \{A \sim B\} \models_{pr} \top \sim \perp$ iff $\mathcal{K} \models_{pr} \top \sim \neg(A \rightarrow B)$.

Proof: From left to right, assume that $\mathcal{K} \not\models_{pr} \top \sim A \wedge \neg B$. Then there is a model P s.t. $P \not\models \top \sim A \wedge \neg B$. That is, there is a valuation v that is minimal in P s.t. $v \Vdash A \rightarrow B$. Let P' be a model consisting of v only. P' is a model of $\mathcal{K} \cup \{A \sim B\}$. The other direction comes from $A \sim B \models_{pr} \top \sim A \rightarrow B$ ([Kraus *et al.*, 1990], Lemma 19) and the monotonicity of \models_{pr} . \square

The postulates for this type of revision are the following:

- (\bullet 1) $\mathcal{K}_{A \sim B}^{\bullet} = \mathcal{P}(\mathcal{K}_{A \sim B}^{\bullet})$ (\bullet closure)
- (\bullet 2) $\mathcal{K}_{A \sim B}^{\bullet} \subseteq \mathcal{K}_{A \sim B}^+$ (\bullet inclusion)
- (\bullet 3) If $\mathcal{K} \cup \{A \sim B\} \not\models_{pr} \top \sim \perp$, then $\mathcal{K}_{A \sim B}^+ \subseteq \mathcal{K}_{A \sim B}^{\bullet}$ (\bullet vacuity)

(\bullet 4) $A \sim B \in \mathcal{K}_{A \sim B}^{\bullet}$ (\bullet success)

(\bullet 5) If $A \sim B \equiv_{\mathcal{P}} A' \sim B'$, then $\mathcal{K}_{A \sim B}^{\bullet} = \mathcal{K}_{A' \sim B'}$ (\bullet ext.)

(\bullet 6) If $A \sim B \not\models_{pr} \top \sim \perp$, then $\mathcal{K}_{A \sim B}^{\bullet} \not\models_{pr} \top \sim \perp$ (\bullet consistency)

(\bullet +) $\mathcal{K}_{A \sim B}^{\bullet} = (\mathcal{K}_{\top \sim A \rightarrow B}^{\bullet})_{A \sim B}^+$ (\bullet vacuity+)

Observe that for revision we require, in addition to the translation of the standard postulates, the extra postulate (\bullet +) . The reason for the need of this postulate relates to the result in Proposition 2, plus the fact that $A \sim B \models_{pr} \top \sim A \rightarrow B$. Proposition 2 implies that to “make room” for $A \sim B$ and preserve consistency we need to “make room” only for the weaker conditional $\top \sim A \rightarrow B$.

A direct translation of the Levi Identity into the preferential framework is not possible, since the negation of a conditional $A \sim B$ cannot be expressed directly. But recall that in classical AGM, contraction with the negation of A should result in a knowledge base that is compatible with A . Proposition 2 gives us the right contraction to apply to preserve consistency, and we rephrase the Levi Identity as follows:

$$\mathcal{K}_{A \sim B}^{\bullet} = (\mathcal{K}_{\top \sim A \wedge \neg B}^{\div})_{A \sim B}^+ \quad (1)$$

From this we get the following representation result.

Theorem 4 *A revision operator \bullet satisfies the postulates (\bullet 1) – (\bullet 6) and (\bullet +) iff it can be defined, via (1), from a contraction operator satisfying the postulates (\div 1) – (\div 6).*

Proof: (Sketch). From right to left, it is easy to check that if \div satisfies (\div 1) – (\div 6), \bullet satisfies (\bullet 1) – (\bullet 6), (\bullet +) . From left to right, define $\mathcal{K}_{A \sim B}^{\div}$ as follows: if $\models A \equiv \top$, then $\mathcal{K}_{A \sim B}^{\div} := (\mathcal{K}_{\top \sim \neg B}^{\bullet} \cap \mathcal{K})$ (a reformulation of the Harper Identity [Harper, 1976]); if $\not\models A \equiv \top$, then $\mathcal{K}_{A \sim B}^{\div} := \bigcap pm(\mathcal{K} \downarrow A \sim B)$ for some choice function *pm*. We prove that \div generates \bullet : $(\mathcal{K}_{\top \sim A \wedge \neg B}^{\div})_{A \sim B}^+ = (\mathcal{K}_{\top \sim A \rightarrow B}^{\bullet} \cap \mathcal{K})_{A \sim B}^+$; $(\mathcal{K}_{\top \sim A \rightarrow B}^{\bullet} \cap \mathcal{K})_{A \sim B}^+ = ((\mathcal{K}_{\top \sim A \rightarrow B}^{\bullet})_{A \sim B}^+ \cap \mathcal{K}_{A \sim B}^+)$, that by (\bullet +) is equivalent to $(\mathcal{K}_{A \sim B}^{\bullet} \cap \mathcal{K}_{A \sim B}^+)$, that by (\bullet 2) is equivalent to $\mathcal{K}_{A \sim B}^{\bullet}$. Now we have to check that \div satisfies (\div 1) – (\div 6). For $\not\models A \equiv \top$, use Theorem 3. If $\models A \equiv \top$, the only difficulty is to prove (\div 6): $\mathcal{K} \subseteq (\mathcal{K}_{\top \sim B}^{\div})_{\top \sim B}^+$ iff $\mathcal{K} \subseteq (\mathcal{K}_{\top \sim \neg B}^{\bullet} \cap \mathcal{K})_{\top \sim B}^+$ iff $\mathcal{K} \subseteq ((\mathcal{K}_{\top \sim \neg B}^{\bullet})_{\top \sim B}^+ \cap (\mathcal{K})_{\top \sim B}^+)$ iff $\mathcal{K} \subseteq (\mathcal{K}_{\top \sim \neg B}^{\bullet})_{\top \sim B}^+$, that holds since $\top \sim \perp \in (\mathcal{K}_{\top \sim \neg B}^{\bullet})_{\top \sim B}^+$. \square

If we want to add a conditional $A \sim B$ preserving not only consistency but also coherence w.r.t. A , we can model a revision operator \circ as

$$\mathcal{K}_{A \sim B}^{\circ} := (\mathcal{K}_{A \sim \neg B}^{\div})_{A \sim B}^+ \quad (2)$$

We characterise such a class of operators with a set of postulates (\circ 1) – (\circ 6) where (\circ 1), (\circ 2), (\circ 4), (\circ 5) are the reformulation for \circ of the corresponding \bullet -postulates, while (\circ 3) and (\circ 6) are:

(\circ 3) If $\mathcal{K} \not\models_{pr} A \sim \neg B$, then $\mathcal{K}_{A \sim B}^+ \subseteq \mathcal{K}_{A \sim B}^{\circ}$ (\bullet vacuity)

(\circ 6) If $A \sim B \models_{pr} A \sim \perp$, then $\mathcal{K}_{A \sim B}^{\circ} \not\models_{pr} A \sim \perp$ (\bullet coherence)

If we use a partial meet contraction $\dot{\dashv}$ in the definition of \circ , it is easy to show that we obtain a revision operator that satisfies $(\circ 1) - (\circ 6)$, but in order to characterise such a class of revision operators, we refer to a more general class of contraction operations: those satisfying $(\dot{\dashv} 1) - (\dot{\dashv} 5)$.

Proposition 3 *A revision operator \circ satisfies $(\circ 1) - (\circ 6)$ iff it can be defined, via (2), using a contraction $\dot{\dashv}$ satisfying $(\dot{\dashv} 1) - (\dot{\dashv} 5)$.*

Proof: (Sketch). If $\dot{\dashv}$ satisfies $(\dot{\dashv} 1) - (\dot{\dashv} 5)$, it is easy to check that $(\circ 1) - (\circ 6)$ are satisfied. For $(\circ 6)$ we need $(\dot{\dashv} 4)$ and Lemma (1) (see below). Now let \circ be a revision operator satisfying $(\circ 1) - (\circ 6)$, and define $\dot{\dashv}$ as $\mathcal{K}_{A \sim B}^{\dot{\dashv}} := (\mathcal{K}_{A \sim B}^{\circ} \cap \mathcal{K})$. First, we have to check that $\dot{\dashv}$ generates \circ . Start with $\mathcal{K}_{A \sim B}^{\circ} = (\mathcal{K}_{A \sim B}^{\circ} \cap \mathcal{K})_{A \sim B}^+$, that is equivalent to $\mathcal{K}_{A \sim B}^{\circ} = ((\mathcal{K}_{A \sim B}^{\circ})_{A \sim B}^+ \cap \mathcal{K}_{A \sim B}^+)$. By $(\circ 2)$ that's equivalent to $\mathcal{K}_{A \sim B}^{\circ} = (\mathcal{K}_{A \sim B}^{\circ})_{A \sim B}^+$, that by $(\circ 6)$, $(\circ 4)$, and $(\circ 3)$ is equivalent to $\mathcal{K}_{A \sim B}^{\circ} = \mathcal{K}_{A \sim B}^{\circ}$. Proving that $\dot{\dashv}$ satisfies $(\dot{\dashv} 1) - (\dot{\dashv} 5)$ is immediate. \square

Contraction operations satisfying $(\dot{\dashv} 1) - (\dot{\dashv} 5)$ can be easily characterised using a choice function that, for the contraction of $A \sim B$ from \mathcal{K} , just pick a subtheory of \mathcal{K} that does not imply $A \sim B$.

5 A Reduction to Classical Contraction

In this section we investigate the option of defining contraction for preferential entailment in terms of classical propositional contraction. As we shall see, it turns out that (i) there is an intuitive way for doing so and (ii) perhaps surprisingly, there is a precise correspondence between classical AGM contraction and preferential AGM contraction. To do so, we'll make use of the following result.

Lemma 1 ([Lehmann and Magidor, 1992], Corollary 2.9) $\mathcal{K} \models_{pr} A \sim B$ iff $\mathcal{K} \cup \{A \sim \neg B\} \models_{pr} A \sim \perp$.

Recall from Section 2 that a knowledge base \mathcal{K} is partitioned into ranks \mathcal{K}^i for $i \in \{0, \dots, n-1, \infty\}$. It is not hard to prove that the entailment of formulas of the form $A \sim \perp$ is determined solely by the conditionals of infinite rank.

Lemma 2 $\mathcal{K} \models_{pr} A \sim \perp$ iff $\mathcal{K}^{\infty} \models_{pr} A \sim \perp$

For brevity, we use \mathcal{K}' to denote $\mathcal{K} \cup \{A \sim \neg B\}$, $\mathcal{K}_{A \sim \neg B}^{\infty}$ to denote $\mathcal{K}'^{\infty} \cap \mathcal{K}$ (that is, the set of conditionals in \mathcal{K} that have infinite rank in \mathcal{K}'), while $\mathcal{K}_{A \sim \neg B}^{<\infty} := \mathcal{K} \setminus \mathcal{K}_{A \sim \neg B}^{\infty}$ (the conditionals in \mathcal{K} that do not have infinite rank in \mathcal{K}'). Then we get from Lemmas 1 and 2 that $\mathcal{K} \models_{pr} A \sim B$ iff $\mathcal{K}'^{\infty} \models_{pr} A \sim \perp$. Also, if $\mathcal{K}' \models_{pr} A \sim \perp$ it must hold that $A \sim \neg B \in \mathcal{K}'^{\infty}$. Therefore $\mathcal{K}'^{\infty} = \mathcal{K}_{A \sim \neg B}^{\infty} \cup \{A \sim \neg B\}$ and

$$\mathcal{K} \models_{pr} A \sim B \text{ iff } \mathcal{K}_{A \sim \neg B}^{\infty} \models_{pr} A \sim B. \quad (3)$$

So, to obtain a preferential partial meet contraction operator w.r.t. the contraction of $A \sim B$ we only consider $\mathcal{K}_{A \sim \neg B}^{\infty}$, leaving $\mathcal{K}_{A \sim \neg B}^{<\infty}$ untouched, as Propositions 4 and 5 imply.

Proposition 4 *If $\mathcal{K}^* \in (\mathcal{K} \downarrow A \sim B)$, then $\mathcal{K}_{A \sim \neg B}^{<\infty} \subseteq \mathcal{K}^*$.*

Proof: (Sketch). If $\mathcal{K}_{A \sim \neg B}^{<\infty} \not\subseteq \mathcal{K}^* \subseteq \mathcal{K}$ and $\mathcal{K}^* \not\models_{pr} A \sim B$, then, by (3), $\mathcal{K}^* \cup \mathcal{K}_{A \sim \neg B}^{<\infty} \not\models_{pr} A \sim B$. \square

Moreover, since $\mathcal{K}'^{\infty} \models_{pr} A \sim \perp$ iff $\overline{\mathcal{K}'^{\infty}} \models \neg A$, and $\overline{\mathcal{K}'^{\infty}} = \overline{\mathcal{K}_{A \sim \neg B}^{\infty}} \cup \{A \rightarrow \neg B\}$, we obtain

$$\mathcal{K} \models_{pr} A \sim B \text{ iff } \overline{\mathcal{K}_{A \sim \neg B}^{\infty}} \models A \rightarrow B. \quad (4)$$

We can conclude the following.

Proposition 5 $\mathcal{K}^* \in (\mathcal{K} \downarrow (A \sim B))$ iff it is a maximal set s.t. $\mathcal{K}_{A \sim \neg B}^{<\infty} \subseteq \mathcal{K}^* \subseteq \mathcal{K}$ and $\overline{\mathcal{K}^* \setminus \mathcal{K}_{A \sim \neg B}^{<\infty}} \in (\overline{\mathcal{K}_{A \sim \neg B}^{\infty}} \downarrow (A \rightarrow B))$.

Proof: (Sketch). From (4) we have, for $\mathcal{K}^* \subseteq \mathcal{K}$, $\mathcal{K}^* \not\models_{pr} A \sim B$ iff $\overline{\mathcal{K}^* \setminus \mathcal{K}_{A \sim \neg B}^{<\infty}} \not\models A \rightarrow B$. Together with the definition of a remainder set and Proposition 4 this proves the proposition. \square

Given a classical propositional contractor operator $-$, we can now justify the following definition of preferential contraction in terms of $-$:

Definition 1 $\mathcal{K}_{A \sim B}^{\dot{\dashv}} := \mathcal{K}_{A \sim B}^{<\infty} \cup \{C \sim D \in \mathcal{K}_{A \sim \neg B}^{\infty} \mid C \rightarrow D \in \text{Cn}(\overline{\mathcal{K}_{A \sim \neg B}^{\infty}})_{A \rightarrow B}\}$

The following is the main result of this section.

Theorem 5 $\dot{\dashv}$ is a partial meet preferential contraction iff it can be defined, in terms of Definition 1, using a partial meet propositional contraction $-$.

Proof: (Sketch). It is easily derivable from Proposition 5. For one direction let $\dot{\dashv}$ be defined as $\bigcap pm(\mathcal{K} \downarrow A \sim B)$ for some pm , and then introduce a propositional choice function pm' as $pm'(\overline{\mathcal{K}_{A \sim \neg B}^{\infty}} \downarrow A \rightarrow B) := \{\mathcal{K}^* \setminus \mathcal{K}_{A \sim \neg B}^{<\infty} \mid \mathcal{K}^* \in pm(\mathcal{K} \downarrow A \sim B)\}$ to define $-$. For the other direction let $-$ be defined using a function pm' and use $pm(\mathcal{K} \downarrow A \sim B) := \{\mathcal{K}_{A \sim \neg B}^{<\infty} \cup \{C \sim D \in \mathcal{K}_{A \sim \neg B}^{\infty} \mid C \rightarrow D \in \mathcal{A}\} \mid \mathcal{A} \in pm'(\overline{\mathcal{K}_{A \sim \neg B}^{\infty}} \downarrow A \rightarrow B)\}$ to define $\dot{\dashv}$. \square

Theorem 5 shows that there is a precise correspondence between classical AGM contraction and preferential AGM contraction: the latter can be determined using the former, preceded by the ranking process described in Section 2.

We conclude this section with a proof of Theorem 3 in Section 5, made easy by the availability of Theorem 5.

Proof of Theorem 3: (Sketch). From left to right it is sufficient to check that every postulate is satisfied. $(\dot{\dashv} 1) - (\dot{\dashv} 5)$ are easy. $(\dot{\dashv} 6)$ comes easily from Theorem 5 and (-6) . From right to left, from $\mathcal{K}_{A \sim B}^{\dot{\dashv}}$ define a partial meet function pm as $\mathcal{K}^* \in pm(\mathcal{K} \downarrow A \sim B)$ iff $\mathcal{K}^* \in (\mathcal{K} \downarrow A \sim B)$ and $\mathcal{K}_{A \sim \neg B}^{\dot{\dashv}} \subseteq \mathcal{K}^*$.

We need to prove that $\mathcal{K}_{A \sim B}^{\dot{\dashv}} = \bigcap pm(\mathcal{K} \downarrow A \sim B)$. It is provable defining a propositional partial meet function pm' over $(\overline{\mathcal{K}_{A \sim \neg B}^{\infty}} \downarrow A \rightarrow B)$ as in the proof of Theorem 5, then defining from it a partial meet propositional contraction $-$, and then showing that $\dot{\dashv}$ can be defined using $-$ in Definition 1. By Theorem 5, $\dot{\dashv}$ is a partial meet preferential contraction. \square

Note that, applying any partial meet operator in Example 1.1 to contract $m \sim \perp$, we preserve the conditional $v \sim n$, since it does not have an infinite rank.

6 Revision w.r.t. Non-monotonic Closure

In the previous sections we investigated belief change for the monotonic core of the (non-monotonic) *spc*-operators. Here we build on those results to venture into the area of belief

change for the *spc*-operators. So consider an *spc*-operator \mathcal{C} . As in the previous sections, our focus will be on operators for *closed* knowledge bases—closed under \mathcal{C} in this case. We use \mathcal{K} to refer to knowledge bases closed under \mathcal{C} . Since revision for \mathcal{C} will be defined in terms of its monotonic core, we need to distinguish between \mathcal{K} and its monotonic core. We refer to the latter as $\mathcal{K}^{\mathcal{P}}$. Here are the postulates for revision w.r.t. \mathcal{C} :

- (*1) $\mathcal{K}_{A \sim B}^* = \mathcal{C}(\mathcal{K}_{A \sim B}^*)$ (* closure)
- (*2) $\exists \mathcal{K}'$ s.t. $\mathcal{C}(\mathcal{K}') = \mathcal{C}(\mathcal{K}_{A \sim B}^*)$ and $\mathcal{K}' \subseteq (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+$ (* generator inclusion)
- (*3) If $(\mathcal{K}^{\mathcal{P}})_{A \sim B}^+ \not\models_{pr} \top \sim \perp$, then $\mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^+) \subseteq \mathcal{K}_{A \sim B}^*$ (* vacuity)
- (*4) $A \sim B \in \mathcal{K}_{A \sim B}^*$ (* success)
- (*5) If $A \sim B \equiv_{\mathcal{P}} A' \sim B'$, then $\mathcal{K}_{A \sim B}^* = \mathcal{K}_{A' \sim B'}^*$ (* extensional.)
- (*6) If $A \sim B \not\models_{pr} \top \sim \perp$, then $\mathcal{K}_{A \sim B}^* \not\models_{\mathcal{C}} \top \sim \perp$ (* consistency)

The only postulate that needs some explanation is (*2). Given that \mathcal{C} is non-monotonic, standard inclusion $(\mathcal{K}_{A \sim B}^* \subseteq \mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^+))$ is not applicable. What (*2) conveys is that it should be possible to generate the non-monotonic closure $\mathcal{K}_{A \sim B}^*$ from the \mathcal{C} -closure of a set of conditionals \mathcal{K}' that is a subset of $(\mathcal{K}^{\mathcal{P}})_{A \sim B}^+$: our revision on the non-monotonic level should be generated by a revision in the monotonic core.

As Theorem 6 shows, the revision for \mathcal{C} can be defined in terms of the revision of the monotonic core.

Theorem 6 ** is a revision operator satisfying (*1) – (*6) iff there is a preferential revision \bullet satisfying (•1) – (•6) s.t. $\mathcal{K}_{A \sim B}^* = \mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^{\bullet})$.*

Proof: From right to left it is easily provable. In the other direction, define \bullet as $(\mathcal{K}^{\mathcal{P}})_{A \sim B}^{\bullet} := (\mathcal{K}_{A \sim B}^* \cap (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+)$. We need to prove that $\mathcal{K}_{A \sim B}^* = \mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^{\bullet})$. $\mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^{\bullet}) = \mathcal{C}(\mathcal{K}_{A \sim B}^* \cap (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+)$ and, by (*2), there is a \mathcal{K}' s.t. $\mathcal{K}_{A \sim B}^* = \mathcal{C}(\mathcal{C}(\mathcal{K}') \cap (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+)$ and $\mathcal{K}' \subseteq (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+$. Hence we have that $\mathcal{K}' \subseteq (\mathcal{C}(\mathcal{K}') \cap (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+) \subseteq \mathcal{C}(\mathcal{K}')$, and, \mathcal{C} being cumulative, we obtain that $\mathcal{C}((\mathcal{C}(\mathcal{K}') \cap (\mathcal{K}^{\mathcal{P}})_{A \sim B}^+)) = \mathcal{C}(\mathcal{K}')$ that in turn equals $\mathcal{K}_{A \sim B}^*$. Check that \bullet satisfies (•1) – (•6). \square At the moment we do not have a representation result that also incorporates the postulate (•+).

We can define revision for coherence preservation in a similar way, except that we have to add a constraint to \mathcal{C} . That is, the monotonic part of our information is determined by the preferential closure: for every A and every \mathcal{K} , $A \sim \perp \in \mathcal{K}$ iff $A \sim \perp \in \mathcal{K}^{\mathcal{P}}$. This property was initially proposed by Lehmann and Magidor ([1992], Sections 5.4-5.5). It is satisfied by most of the prominent non-monotonic entailment relations proposed in the preferential framework [Kraus *et al.*, 1990; Lehmann and Magidor, 1992; Lehmann, 1995; Casini and Straccia, 2013]. The revision postulates for coherence preservation are (⊗1), (⊗2), (⊗4), and (⊗5), that are direct reformulations of ⊗ of (*1), (*2), (*4), and (*5), and:

- (⊗3) If $\mathcal{K}^{\mathcal{P}} \not\models_{pr} A \sim \neg B$, then $\mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^+) \subseteq \mathcal{K}_{A \sim B}^{\otimes}$ (⊗-vacuity)

(⊗6) If $A \sim B \not\models_{pr} A \sim \perp$, then $\mathcal{K}_{A \sim B}^{\otimes} \not\models_{\mathcal{C}} A \sim \perp$ (⊗-coherence)

We obtain the analogous of Th. 6 for coherence preservation.

Theorem 7 *⊗ is a revision operator satisfying (⊗1) – (⊗6) iff there is a preferential revision \circ satisfying (◦1) – (◦6) s.t. $\mathcal{K}_{A \sim B}^{\otimes} = \mathcal{C}((\mathcal{K}^{\mathcal{P}})_{A \sim B}^{\circ})$.*

Proof: The proof is similar to that of Theorem 6. In the first half, to prove that (⊗6) hold we have to refer to the property that $A \sim \perp \in \mathcal{K}$ iff $A \sim \perp \in \mathcal{K}^{\mathcal{P}}$. \square

As (*3) indicates, the operator is such that if the conflicts do not arise at the level of preferential entailment, the revision resolves in a simple expansion, as indicated in Example 1.2.

7 Conclusion

In this paper we have laid the foundations for defining belief change in (preferential) non-monotonic framework. Following the well-known AGM approach, we defined and characterised AGM expansion, contraction and revision (for both consistency and coherence preservation) on the monotonic core of a class of non-monotonic closure operators—the *spc*-operators. We also showed that preferential AGM contraction can be reduced to classical AGM contraction. Then we showed how to define revision for the class of *spc*-operators.

As far as we know, just a few works deal with belief change in a nonmonotonic framework. Booth and Paris [1998] deal with conditional bases, but also containing negative conditionals, which can be thought of as related to the contraction of conditionals. Casini and Meyer [2016] are concerned with the contraction of strict conditionals $A \sim \perp$ from finite conditional bases in order to restore coherence. Hunter [2016] deals with the revision of highly implausible conditionals, while Delgrande and others [2013] analyse belief change in Answer Set Programming. Conditional belief change was previously addressed by Kern-Isberner [1999], who proposes a set of rationality postulates; the relations between her approach and ours will be object of further analysis.

For future work, an obvious next step is to define a form of base change for defeasible bases, but in such a way that it adheres to the type of change defined for the closed knowledge bases. For example, for a base contraction operator $-$ of this kind, we would require that $\mathcal{P}(\mathcal{B}_{A \sim B}^-) = \mathcal{P}(\mathcal{B}_{A \sim B}^-)$. We have started to work along these lines where the base contraction operator $-$ is built on top of similar propositional base contraction operators [Nebel, 1989; Meyer, 2001]. We also intend to extend our work to belief change for Description Logics. Recently, there has been a lot of interest in developing non-monotonic formalisms for this family of logics, especially in the preferential framework [Britz *et al.*, 2008; 2011; Casini and Straccia, 2010; 2013; Giordano *et al.*, 2013; 2015; Lukasiewicz, 2008]. Given the importance of coherence in the field, we expect such an extension to be very useful.

Acknowledgments

The work of Giovanni Casini has been supported by the Fonds National de la Recherche, Luxembourg, and cofunded by the Marie Curie Actions of the European Commission (FP7-COFUND) (AFR/9181001). Thomas Meyer has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agr. No. 690974.

References

- [Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [Booth and Paris, 1998] Richard Booth and Jeff B. Paris. A note on the rational closure of knowledge bases with both positive and negative knowledge. *Journal of Logic, Language and Information*, 7(2):165–190, 1998.
- [Britz *et al.*, 2008] Katarina Britz, Johannes Heidema, and Thomas Meyer. Semantic preferential subsumption. In J. Lang and G. Brewka, editors, *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning, KR 2008*, pages 476–484. AAAI Press/MIT Press, 2008.
- [Britz *et al.*, 2011] Katarina Britz, Thomas Meyer, and Ivan Varzinczak. Semantic foundation for preferential Description Logics. In D. Wang and M. Reynolds, editors, *Proceedings of the 24th Australasian Joint Conference on Artificial Intelligence*, number 7106 in LNAI, pages 491–500. Springer, 2011.
- [Casini and Meyer, 2016] Giovanni Casini and Thomas Meyer. Using defeasible information to obtain coherence. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning, KR 2016*, pages 537–540, 2016.
- [Casini and Straccia, 2010] Giovanni Casini and Umberto Straccia. Rational closure for defeasible Description Logics. In T. Janhunen and I. Niemelä, editors, *Proceedings of the 12th European Conference on Logics in Artificial Intelligence (JELIA)*, number 6341 in LNCS, pages 77–90. Springer-Verlag, 2010.
- [Casini and Straccia, 2013] Giovanni Casini and Umberto Straccia. Defeasible inheritance-based Description Logics. *J. Artif. Intell. Res. (JAIR)*, 48:415–473, 2013.
- [Delgrande *et al.*, 2013] James Delgrande, Torsten Schaub, Hans Tompits, and Stefan Woltran. A model-theoretic approach to belief change in answer set programming. *ACM Trans. Comput. Logic*, 14(2):1–46, 2013.
- [Giordano *et al.*, 2013] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. A non-monotonic Description Logic for reasoning about typicality. *Artif. Intell.*, 195:165–202, 2013.
- [Giordano *et al.*, 2015] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. Semantic characterization of rational closure: From propositional logic to Description Logics. *Artif. Intell. J.*, 226:1–33, 2015.
- [Harper, 1976] William L. Harper. Rational conceptual change. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1976:462–494, 1976.
- [Hunter, 2016] Aaron Hunter. *Nearly Counterfactual Revision*, pages 263–269. Springer International Publishing, Cham, 2016.
- [Kern-Isberner, 1999] Gabriele Kern-Isberner. Postulates for conditional belief revision. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'99*, pages 186–191, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
- [Kraus *et al.*, 1990] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.*, 44:167–207, 1990.
- [Lehmann and Magidor, 1992] Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55:1–60, 1992.
- [Lehmann, 1995] Daniel Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82, 1995.
- [Levi, 1977] Isaac Levi. Subjunctives, dispositions and chances. *Synthese*, 34:423–455, 1977.
- [Lukasiewicz, 2008] Thomas Lukasiewicz. Expressive probabilistic Description Logics. *Artif. Intell.*, 172(6-7):852–883, 2008.
- [Meyer, 2001] Thomas Meyer. Basic infobase change. *Studia Logica*, 67(2):215–242, 2001.
- [Nebel, 1989] Bernhard Nebel. A knowledge level analysis of belief revision. In *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*. Toronto, Canada, May 15-18 1989., pages 301–311, 1989.
- [Qi and Hunter, 2007] Guilin Qi and Anthony Hunter. Measuring incoherence in Description Logic-based ontologies. In *The Semantic Web*, volume 4825 of LNCS, pages 381–394. Springer, 2007.