A Semantic Perspective on Belief Change in a Preferential Non-Monotonic Framework

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

Eduardo Fermé
Universidade da Madeira
NOVA-LINCS
Portugal
ferme@uma.pt

Thomas Meyer
CAIR
(CSIR & Univ. of Cape Town)
South Africa
tmeyer@cs.uct.ac.za

Ivan Varzinczak
CRIL
Univ. Artois & CNRS
France
varzinczak@cril.fr

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 investigate the integration of the two formalisms by studying belief change for a (preferential) non-monotonic framework. We show that the standard AGM approach to belief change can be transferred to a preferential non-monotonic framework in the sense that change operations can be defined on *conditional* knowledge bases. We take as a point of departure the results presented by Casini and Meyer (2017), and we develop and extend such results with characterisations based on semantics and entrenchment relations, showing how some of the constructions defined for propositional logic can be lifted to our preferential non-monotonic framework.[1]

1 Introduction

Both belief change and non-monotonic reasoning deal with the problem of handling conflicting information. For example, suppose we 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 do not have a nucleus ($m \rightarrow \neg n$). The existence of mammalian red-blood cells (m) then renders our knowledge base inconsistent. Belief-change operators modify the existing knowledge base to preserve its consistency (in our example, $v \rightarrow n$ should be weakened). On the other hand, non-monotonic reasoning usually handles such conflicts by introducing *defeasibility* (this amounts to stating that vertebrate red-blood cells *usually* have a nucleus). Both mechanisms are deeply interconnected. Indeed, technically each of them can be considered as a re-formulation of the other (Makinson 1993).

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. This view does not take into account that non-monotonic frameworks contain a mix of defeasible and classical (non-defeasible) information. In our example above 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. 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.

Our goal is the characterisation of belief change in a non-monotonic setting, that is still an open problem, since the standard approaches to belief change usually assume an underlying Tarskian consequence relation which is explicitly monotonic (Alchourrón, Gärdenfors, and Makinson 1985). This is not the first study on belief revision in a conditional framework, but it differs from most of the previous approaches, such as those by Kern-Isberner (1999) and Wobcke (1995). These proposals give to the conditionals a subjunctive interpretation and rely on the already-known correspondence between them and the belief revision operators: the conditionals are interpreted as direct expressions of the belief revision policies of the agent, and usually the main technical problem is the definition of non-trivial revision operators that avoid the well-known Gärdenfors' impossibility result (Gärdenfors 1988, Section. 7.4). Here we work on conditional knowledge bases, and we do not assume a subjunctive interpretation of the conditionals. Rather, the conditional knowledge base could formalise defeasible information of a different nature, describing for example prototypical behaviour or deontic constraints, and the conditionals are the objects of the belief change, not an alternative expression of the belief revision policies. As we shall see, this also implies that Gärdenfors' impossibility result does not apply in our framework. The present paper extends the results on the topic by Casini and Meyer (2017) by giving a semantic characterisation and defining another class of revision operators.

In addition to the classical problem of *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. A knowledge base is coherent if 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 formula A, for that matter), corresponds to the statement $\neg A$ being a consequence of our knowledge base.

We focus on modelling belief change assuming as the underlying logical framework that of non-monotonic reasoning developed by Kraus, Lehmann and Magidor (1990), in which *defeasible* conditionals of the form $C \mathrel{|\!\sim} D$ are added to the language of propositional logic (with C and D being classical propositional formulas).

The paper is structured as follows: In Section 2 we give brief summaries of both the AGM and KLM frameworks, and summarise the main results by Casini and Meyer (2017). In Section 3 we present our first contribution, namely a semantic characterisation of Casini and Meyer's operators. Section 4 is devoted to the characterisation of a new class of change operators.

2 Preliminaries

We consider a propositional language $\mathcal{L}$ generated by some finite set of atoms $\mathcal{A}$, with lower case letters denoting atoms, and capital letters denoting elements of $\mathcal{L}$. We adopt the standard semantics for propositional logic. The set $\mathcal{V}$ of valuations v are functions from the set of atoms in $\mathcal{L}$ to $\{0, 1\}$, denoting truth and falsity. Satisfaction is denoted by $\Vdash$, and entailment by $\models$, while $\overline{v}$ (respectively, $\overline{S}$) indicates a formula uniquely characterising v (a set of valuations $S \subseteq \mathcal{V}$). $[A]$ denotes the set of valuations satisfying A.

2.1 AGM belief change

AGM belief change (Alchourrón, Gärdenfors, and Makinson 1985) assumes an underlying logic with a propositional language and a Tarskian consequence relation $Cn(\cdot)$ that is compact and satisfies *disjunction in the premises* (Alchourrón, Gärdenfors, and Makinson 1985, p. 511–512). A knowledge base K is assumed to be a set of formulas *closed under $Cn(\cdot)$*. AGM is concerned with three types of operations on knowledge bases: expansion, contraction, and revision. Expansion is simply defined as adding a formula and closing under entailment: $\mathsf{K}^+_A \equiv_{\text{def}} Cn(\mathsf{K} \cup \{A\})$. The intuition associated with the contraction of K by A is that it should result in a knowledge base K^-_A not entailing A. Dually, a revision of K by A should result in a consistent knowledge base K^*_A from which A follows.

In the AGM approach, any appropriate contraction operator is required to satisfy six *basic postulates for contraction*: *Closure* $(-_1)$, *Inclusion* $(-_2)$, *Vacuity* $(-_3)$, *Success* $(-_4)$, *Extensionality* $(-_5)$, and *Recovery* $(-_6)$. For details, we refer the reader to one the many publications presenting the AGM theory (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988; Hansson 1999). Additionally, AGM contraction may be required to satisfy the two supplementary postulates about conjunctions:

$(-_7)$ $\mathsf{K}^-_A \cap \mathsf{K}^-_B \subseteq \mathsf{K}^-_{(A \wedge B)}$ (Conjunctive Overlap)

$(-_8)$ If $A \notin \mathsf{K}^-_{(A \wedge B)}$, then $\mathsf{K}^-_{(A \wedge B)} \subseteq \mathsf{K}^-_A$ (Conj. Inclusion)

AGM contraction operators can be constructed using so-called *partial-meet functions* (Alchourrón, Gärdenfors, and

Makinson 1985, Obs. 2.5). Let $\mathsf{K} \!\downarrow\! A$ be the *remainder set* of K w.r.t. A, defined as the set containing the maximal subsets K' of K s.t. $A \notin \mathsf{K}'$. That is, $\mathsf{K}' \in \mathsf{K} \!\downarrow\! A$ iff (*i*) $\mathsf{K}' \subseteq \mathsf{K}$, (*ii*) $A \notin \mathsf{K}'$, and (*iii*) there is no set K'' s.t. $\mathsf{K}' \subset \mathsf{K}'' \subseteq \mathsf{K}$ and $A \notin \mathsf{K}''$. Let $pm(\cdot)$ be a partial-meet function defined over $\mathsf{K} \!\downarrow\! A$ s.t. $pm(\mathsf{K} \!\downarrow\! A) \subseteq \mathsf{K} \!\downarrow\! A$ and, if $\mathsf{K} \!\downarrow\! A \neq \varnothing$, then $pm(\mathsf{K} \!\downarrow\! A) \neq \varnothing$. A partial-meet contraction operator $-$ is defined as: $\mathsf{K}^-_A = \bigcap pm(\mathsf{K} \!\downarrow\! A)$.

Observation 1 (Alchourrón, Gärdenfors, and Makinson 1985, Obs. 2.5) *An operator $-$ on K is a partial meet contraction iff $-$ satisfies $(-_1)$–$(-_6)$.*

Similarly, AGM revision is required to satisfy six basic postulates (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988): *Closure* $(*_1)$, *Inclusion* $(*_2)$, *Vacuity* $(*_3)$, *Success* $(*_4)$, *Extensionality* $(*_5)$, and *Consistency* $(*_6)$. Two supplementary postulates for revision are proposed:

$(*_7)$ $\mathsf{K}^*_{(A \wedge B)} \subseteq (\mathsf{K}^*_A)^+_B$ (Superexpansion)

$(*_8)$ If $\neg B \notin \mathsf{K}^*_A$, then $(\mathsf{K}^*_A)^+_B \subseteq \mathsf{K}^*_{(A \wedge B)}$ (Subexpansion)

Revision can be defined in terms of contraction and expansion via the Levi Identity (Levi 1977): $\mathsf{K}^*_A = (\mathsf{K}^-_{\neg A})^+_A$.

Observation 2 (Alchourrón, Gärdenfors, and Makinson 1985; Fermé 1999) *Let $-$ be an operator on K satisfying $(-_1)$–$(-_5)$, and let $*$ be defined via the Levi Identity. Then:*

1. *$*$ satisfies $(*_1)$–$(*_6)$.*
2. *If $-$ also satisfies $(-_6)$, $(-_7)$ and $(-_8)$, then $*$ also satisfies $(*_7)$ and $(*_8)$.*

There are multiple ways of characterising the class of the contraction operations satisfying the six basic postulates plus $(-_7)$ and $(-_8)$. One way is through *epistemic entrenchment relations* $\leqslant_e$, introduced by Gärdenfors (1988) to define the properties that an order over the sentences of $\mathcal{L}$ should satisfy. He proposed the following set of axioms:

E1 if $C \leqslant_e D$ and $D \leqslant_e E$, then $C \leqslant_e E$ (Transitivity)

E2 if $C \models D$, then $C \leqslant_e D$ (Dominance)

E3 Either $C \leqslant_e C \wedge D$ or $D \leqslant_e C \wedge D$ (Conjunctiveness)

E4 If $\perp \notin \mathsf{K}$, then $C \notin \mathsf{K}$ iff $C \leqslant_e D$ for all D (Minimality)

E5 If $D \leqslant_e C$ for every D, then $\models C$ (Maximality)

Gärdenfors proposed the following connections between orders of epistemic entrenchment and operations of contraction (Gärdenfors 1988):

$$D \in \mathsf{K}^{-e}_C \text{ iff } D \in \mathsf{K} \text{ and either } C <_e C \vee D \text{ or } \models C \quad (1)$$

$$C \leqslant_e D \text{ iff } C \notin \mathsf{K}^-_{C \wedge D} \text{ or } \models C \wedge D \quad (2)$$

Observation 3 (Gärdenfors 1988; Gärdenfors and Makinson 1988) *Given an epistemic entrenchment $\leqslant_e$ on a consistent belief set K, then $-e$, defined via (1), is a contraction operator satisfying $(-_1) - (-_8)$. Vice versa, if we have a contraction operator $-$ satisfying $(-_1) - (-_8)$, there is an epistemic entrenchment $\leqslant_e$ s.t. $-$ corresponds to the contraction $-e$ defined via (2).*

Rott (1991) proposed an alternative class of entrenchment-based contraction operator:

$$D \in \mathsf{K}_C^{-R} \text{ iff } D \in \mathsf{K} \text{ and either } C <_e D \text{ or } \models C \qquad (3)$$

In order to relate $-_e$ and $-_R$, let $\mathsf{K} \Downarrow C$ be a set of preferred remainder sets w.r.t. the epistemic entrenchment relation,

$$\mathsf{K} \Downarrow C \equiv_{\text{def}} \{\mathsf{K}' \in \mathsf{K} \downarrow C \mid \mathsf{K}_C^{-R} \subseteq \mathsf{K}'\}$$

The operator defined in (1) corresponds to the intersection of all the elements of $\mathsf{K} \Downarrow C$.

Proposition 1 *For every K and every $C, D \in \mathcal{L}$, $D \in \mathsf{K}_C^{-e}$ iff $D \in \mathsf{K}'$ for every $\mathsf{K}' \in \mathsf{K} \Downarrow C$.*

Proposition 1, beyond helping in further analysing the relation between Gärdenfors' and Rott's contractions, gives an alternative definition of Gärdenfors' contraction in terms of full meet of the preferred remainder set. We will refer again to this characterisation later on.

The semantics of the AGM model can be characterized by a total pre-order over the set of valuations. A total pre-order $\leqslant_\mathsf{K}$ on valuations, with the strict part $<_\mathsf{K}$ and the symmetric part $\simeq_\mathsf{K}$, is a *faithful assignment* if and only if the following conditions hold (1) If $v \in [\mathsf{K}]$ and $v' \in [\mathsf{K}]$, then $v \simeq_\mathsf{K} v'$ and (2) If $v \in [\mathsf{K}]$ and $v' \notin [\mathsf{K}]$, then $v <_\mathsf{K} v'$.

The notion of a faithful assignment allows us to characterise contraction and revision operations satisfying $(-_1) - (-_8)$ and $(*_1) - (*_8)$, respectively:

Observation 4 (Katsuno and Mendelzon 1991) *Let K be a belief set, and let*

$$\min_{\leqslant_\mathsf{K}}([A]) \equiv_{\text{def}} \{v \in [A] \mid v \leqslant_\mathsf{K} u \text{ for every } u \in [A]\}.$$

An operation $-$ on K satisfies $(-_1) - (-_8)$ if and only if there is a faithful assignment $\leqslant_\mathsf{K}$ for K such that $\mathsf{K}_A^- = \overline{[\mathsf{K}] \cup \min_{\leqslant_K}([\neg A])}$. An operation $*$ on K satisfies $(*_1) - (*_8)$ if and only if there is a faithful assignment $\leqslant_\mathsf{K}$ for K such that $\mathsf{K}_A^* = \overline{\min_{\leqslant_\mathsf{K}}([A])}$.

2.2 Preferential reasoning

To introduce defeasibility we consider the language $\mathcal{L}^{\vert\sim}$ consisting of conditionals of the form $A \vert\sim B$ (for $A, B \in \mathcal{L}$), that can be read as 'typically, if A then B'. The semantics of $\mathcal{L}^{\vert\sim}$ 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 $\mathcal{V}$ (mapping states into valuations), and $<$ is a strict partial order on W, that also satisfy the *smoothness* condition: for every consistent propositional formula A, $\min_<([A]^W) \neq \varnothing$, where $[A]^W \equiv_{\text{def}} \{s \in W \mid l(s) \Vdash A\}$ and $\min_<([A]^W) \equiv_{\text{def}} \{s \in [A]^W \mid \text{ there is no } t \in [A]^W \text{ s.t. } t < s\}$ (Kraus, Lehmann, and Magidor 1990). $s < t$ is interpreted as indicating that the state s represents a more typical situation than the state t. A defeasible conditional $A \vert\sim B$ is satisfied in a preferential interpretation $\mathcal{P}$, denoted as $\mathcal{P} \Vdash A \vert\sim B$, iff $\min_<[A]^W \subseteq [B]^W$. Observe that a propositional formula A is satisfied in all the elements of W (i.e., $l(w) \Vdash A$ for every $w \in W$) iff $\mathcal{P} \Vdash \neg A \vert\sim \bot$. This means that any classical

propositional formula A can be represented as the defeasible conditional $\neg A \vert\sim \bot$. Indeed, every conditional of the form $A \vert\sim \bot$ is *not* defeasible, and actually represents classical propositional information. In our example, the statement that vertebrate red-blood cells have a nucleus ($v \to n$) will be represented as $\neg(v \to n) \vert\sim \bot$. Because of this, we sometimes abuse notation by referring to $\neg A \vert\sim \bot$ as the propositional formula A, or using the strict conditional $A \vdash B$ to indicate $A \wedge \neg B \vert\sim \bot$. Let $\mathfrak{U}$ indicate the set of all the preferential interpretations for our language $\mathcal{L}^{\vert\sim}$, and $\mathfrak{P}_1, \mathfrak{P}_2, \ldots$ be subsets of $\mathfrak{U}$. Given a preferential model $\mathcal{P}$, let $\vert\sim_\mathcal{P}$ be the set of all the conditionals that are satisfied by $\mathcal{P}$ ($\vert\sim_\mathcal{P} \equiv_{\text{def}} \{A \vert\sim B \in \mathcal{L} \mid \mathcal{P} \Vdash A \vert\sim B\}$), while $\vert\sim_\mathfrak{P}$ represents the conditionals satisfied by a set of preferential models $\mathfrak{P}$ ($\vert\sim_\mathfrak{P} \equiv_{\text{def}} \{A \vert\sim B \in \mathcal{L} \mid \text{ for every } \mathcal{P} \in \mathfrak{P}, \mathcal{P} \Vdash A \vert\sim B\}$).

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 $\mathcal{B} \models_{pr} A \vert\sim B$ iff $[\![\mathcal{B}]\!] \subseteq [\![\{A \vert\sim B\}]\!]$. The corresponding closure operator is known as *preferential closure*: $\mathcal{PC}(\mathcal{B}) = \{A \vert\sim B \mid \mathcal{B} \models_{pr} A \vert\sim B\}$. 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 $\mathcal{B} \models_{pr} \top \vert\sim \bot$, and $\mathcal{B}$ is *preferentially incoherent* w.r.t. a propositional formula A iff $\mathcal{B} \models_{pr} A \vert\sim \bot$. Two bases are *preferentially equivalent* iff they have the same preferential models.

It is widely recognised that preferential entailment is too weak to be an appropriate form of entailment for a non-monotonic framework (Kraus, Lehmann, and Magidor 1990, pp. 4, 34). This is primarily because the preferential entailment relation itself is monotonic (non-monotonicity occurs at 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, Lehmann, and Magidor 1990; Lehmann and Magidor 1992). Here we consider a large class of closure operations (entailment relations) extending the preferential closure (preferential entailment), the *supra-preferential cumulative* closure operators (Casini and Meyer 2017, Section 3). A closure operator $\mathcal{C}(\cdot)$, is *supra-preferential cumulative* if it can be defined as follows: $A \vert\sim B \in \mathcal{C}(\mathcal{B})$ iff $\mathcal{P} \Vdash A \vert\sim B$, for every preferential model $\mathcal{P} \in s([\![\mathcal{B}]\!])$, where s is a choice function s.t. (*i*) $s([\![\mathcal{B}]\!]) \subseteq [\![\mathcal{B}]\!]$; (*ii*) if $[\![\mathcal{B}]\!] \neq \varnothing$, then $s([\![\mathcal{B}]\!]) \neq \varnothing$; (*iii*) $s([\![\mathcal{B}]\!]) = s(s([\![\mathcal{B}]\!]))$; and (*iv*) if $s([\![\mathcal{B}]\!]) \subseteq [\![\mathcal{B}']\!] \subseteq [\![\mathcal{B}]\!]$, then $s([\![\mathcal{B}']\!]) = s([\![\mathcal{B}]\!])$. The closure operators definable using such choice functions are the ones that extend preferential closure, respect *consistency preservation* (if $\top \vert\sim \bot \notin \mathcal{PC}(\mathcal{B})$, then $\top \vert\sim \bot \notin \mathcal{C}(\mathcal{B})$), and satisfy *cumulativity*, that is, for every $\mathcal{B}, \mathcal{B}'$, if $\mathcal{B} \subseteq \mathcal{B}' \subseteq \mathcal{C}(\mathcal{B})$, then $\mathcal{C}(\mathcal{B}') = \mathcal{C}(\mathcal{B})$. 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, Lehmann, and Magidor 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.

2.3 Previous results

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. Nevertheless, that is not always the case, as shown in Example 1.

Example 1 *Let* $\mathcal{A} = \{a, m, n, v\}$, *with the propositions standing for, respectively, "(being) an avian red-blood cell", "(being) a mammalian red-blood cell", "(being) a vertebrate red-blood cell", and "(having) a nucleous". Assume a knowledge base* $\mathcal{B} = \{v \vdash n, a \vdash v, m \vdash v, m \vdash \neg n\}$. *Consider the following situations:*

1. $\neg m \in \mathcal{C}(\mathcal{B})$. *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* $m \vdash v$, $v \vdash n$, *or* $m \vdash \neg n$. *In the framework we propose we can resolve the situation by introducing defeasibility. In contracting* $\neg m$, *we would like to end up, for example, with* $\mathcal{B}' = \{v \mathrel{|\sim} n, a \vdash v, m \vdash v, m \vdash \neg n\}$ *in which, instead of eliminating* $v \vdash n$, *we have just made it defeasible.*

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

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

What is shown in the example is that belief change in a non-monotonic framework needs to distinguish between the *certain* portion of our belief state (corresponding to the *monotonic* part of our reasoning) and the *uncertain, defeasible* portion (the *non-monotonic* part): if new information creates a conflict, the dynamics should be different, according to whether the conflict is with the certain portion of our state (hence we need to operate a non-trivial change) or with its uncertain portion (just expand the information, and the non-monotonic machinery will take care of resolving

the conflict). Hence, a main difference between modelling belief change in a Tarskian or in a non-monotonic framework is the one shown in Figure 1: while in the former case we move from a closed theory K, representing our belief state, to a new closed theory K′, in the latter, to describe our belief state, we need to distinguish between its monotonic part, a theory $\mathcal{K}$ closed under a Tarskian operator, and its non-monotonic part, the information obtained closing $\mathcal{K}$ under a non-monotonic operator $\mathcal{C}(\cdot)$. Also, such a distinction should be preserved at the end of the operation.
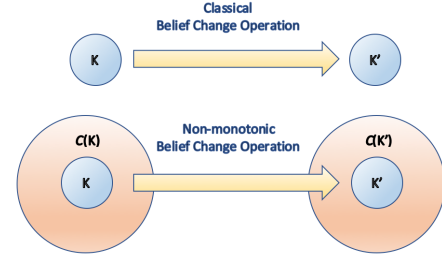

Figure 1: Classical vs. Non-monotonic Belief Change

Hence, it is essential to our approach that, considering some non-monotonic entailment relation, we identify its "monotonic part", that is, the biggest monotonic entailment contained in it. We call a closure operator $Cl(\cdot)$ the *monotonic core* of a non-monotonic operator $\mathcal{C}(\cdot)$ if, for every $\mathcal{B}, \mathcal{B}'$ we have (*i*) $\mathcal{B} \subseteq \mathcal{B}'$ implies $Cl(\mathcal{B}) \subseteq Cl(\mathcal{B}')$; (*ii*) $Cl(\mathcal{B}) \subseteq \mathcal{C}(\mathcal{B})$; (*iii*) for every closure operator $Cl'(\cdot)$ satisfying (*i*) and (*ii*), $Cl'(\mathcal{B}) \subseteq Cl(\mathcal{B})$. It turns out that the preferential closure $\mathcal{PC}$ is the monotonic core of every supra-preferential closure operator (Casini and Meyer 2017, Proposition 1).

Casini and Meyer (2017) provide a characterisation of non-monotonic revision for *spc*-operators. *Preferential expansion* is defined as $\mathcal{K}^+_{A|\sim B} \equiv_{\text{def}} \mathcal{PC}(\mathcal{K} \cup \{A \mathrel{|\sim} B\})$.

Basic preferential contraction $\dot{-}$ is characterised by the following postulates (where $\equiv_{\mathcal{P}}$ refers to preferential equivalence):

($\dot{-}_1$) $\mathcal{K}^{\dot{-}}_{A|\sim B} = \mathcal{PC}(\mathcal{K}^{\dot{-}}_{A|\sim B})$

($\dot{-}_2$) $\mathcal{K}^{\dot{-}}_{A|\sim B} \subseteq \mathcal{K}$

($\dot{-}_3$) If $\mathcal{K} \not\models_{pr} A \mathrel{|\sim} B$, then $\mathcal{K}^{\dot{-}}_{A|\sim B} = \mathcal{K}$

($\dot{-}_4$) If $\not\models_{pr} A \mathrel{|\sim} B$, then $A \mathrel{|\sim} B \notin \mathcal{K}^{\dot{-}}_{A|\sim B}$

($\dot{-}_5$) If $A \mathrel{|\sim} B \equiv_{\mathcal{P}} A' \mathrel{|\sim} B'$, then $\mathcal{K}^{\dot{-}}_{A|\sim B} = \mathcal{K}^{\dot{-}}_{A'|\sim B'}$

($\dot{-}_6$) $\mathcal{K} \subseteq (\mathcal{K}^{\dot{-}}_{A|\sim B})^+_{A|\sim B}$

The class of such operators corresponds to the class of *pm*-contractions defined over the set of the preferential theories (Casini and Meyer 2017, Theorem 3).

Two kinds of revision operations are introduced for the monotonic core $\mathcal{PC}(.)$, one to preserve consistency, the other aimed at the preservation of coherence. *Basic preferential revisions* • for the preservation of *consistency* behave in such a way that, adding a new conditional $A \mathrel{|\sim} B$, we obtain a

new preferential theory $\mathcal{K}^{\bullet}_{A|\sim B}$ that contains $A \mid\sim B$ and is logically consistent ($\top \mid\sim \bot \notin \mathcal{K}^{\bullet}_{A|\sim B}$). This class of operators is defined by the following postulates:

($\bullet_1$) $\mathcal{K}^{\bullet}_{A|\sim B} = \mathcal{PC}(\mathcal{K}^{\bullet}_{A|\sim B})$

($\bullet_2$) $\mathcal{K}^{\bullet}_{A|\sim B} \subseteq \mathcal{K}^{+}_{A|\sim B}$

($\bullet_3$) If $\mathcal{K} \cup \{A \mid\sim B\} \not\models_{pr} \top \mid\sim \bot$, then $\mathcal{K}^{+}_{A|\sim B} \subseteq \mathcal{K}^{\bullet}_{A|\sim B}$

($\bullet_4$) $A \mid\sim B \in \mathcal{K}^{\bullet}_{A|\sim B}$

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

($\bullet_6$) If $A \mid\sim B \not\models_{pr} \top \mid\sim \bot$, then $\mathcal{K}^{\bullet}_{A|\sim B} \not\models_{pr} \top \mid\sim \bot$

($\bullet_+$) $\mathcal{K}^{\bullet}_{A|\sim B} = (\mathcal{K}^{\bullet}_{\top |\sim A \to B})^{+}_{A|\sim B}$

This class of revision operators is characterised using the class of *Basic Preferential Contraction* operators via the following re-formulation of the Levi Identity (Casini and Meyer 2017, Theorem 4):

$$\mathcal{K}^{\bullet}_{A|\sim B} = (\mathcal{K}^{\div}_{\top |\sim A \wedge \neg B})^{+}_{A|\sim B} \tag{4}$$

On the other hand, *basic preferential revision* $\circ$ aimed at the preservation of *coherence* models a change such that, adding a new conditional $A \mid\sim B$, we obtain a new preferential theory $\mathcal{K}^{\circ}_{A|\sim B}$ containing $A \mid\sim B$ and that is coherent w.r.t. A ($A \mid\sim \bot \notin \mathcal{K}^{\circ}_{A|\sim B}$). This class of operators is characterised by 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, whereas $(\circ_3)$ and $(\circ_6)$ are:

($\circ_3$) If $\mathcal{K} \not\models_{pr} A \mid\sim \neg B$, then $\mathcal{K}^{+}_{A|\sim B} \subseteq \mathcal{K}^{\circ}_{A|\sim B}$

($\circ_6$) If $A \mid\sim B \not\models_{pr} A \mid\sim \bot$, then $\mathcal{K}^{\circ}_{A|\sim B} \not\models_{pr} A \mid\sim \bot$

This class of revision operators is characterised using a super-class of the basic preferential contraction operators, the *preferential withdrawals* (satisfying the Postulates $(\div_1)$–$(\div_5)$) via the following re-formulation of the Levi Identity (Casini and Meyer 2017, Proposition 3):

$$\mathcal{K}^{\circ}_{A|\sim B} := (\mathcal{K}^{\div}_{A|\sim \neg B})^{+}_{A|\sim B} \tag{5}$$

The results above serve as a springboard for the definition of belief change for the *spc*-operators. Let $\mathcal{C}(\cdot)$ be an *spc*-operator and let $\mathcal{K}$ now refer to knowledge bases that are closed under $\mathcal{C}(\cdot)$. We define revision for the non-monotonic closure operator $\mathcal{C}(\cdot)$ in terms of its monotonic core. Hence, we distinguish between $\mathcal{K}$ and its monotonic core (see Figure 1) that we indicate as $\mathcal{K}^{\mathcal{PC}}$; $\mathcal{K}^{\mathcal{PC}}$ is a preferential theory s.t. $\mathcal{K} = \mathcal{C}(\mathcal{K}^{\mathcal{PC}})$. The agent must be aware of such a distinction from the beginning, and our operators allow to preserve it at every step; if we are dealing with theories initially generated by a conditional base $\mathcal{B}$, the distinction would be evident, since we would have $\mathcal{K} = \mathcal{C}(\mathcal{B})$ and $\mathcal{K}^{\mathcal{PC}} = \mathcal{PC}(\mathcal{B})$. The postulates for the class of basic revisions w.r.t. $\mathcal{C}(\cdot)$ for the preservation of consistency (Casini and Meyer 2017, Section 6), are:

($*_1$) $\mathcal{K}^{*}_{A|\sim B} = \mathcal{C}(\mathcal{K}^{*}_{A|\sim B})$

($*_2$) There is a $\mathcal{K}'$ s.t. $\mathcal{C}(\mathcal{K}') = \mathcal{C}(\mathcal{K}^{*}_{A|\sim B})$ and $\mathcal{K}' \subseteq (\mathcal{K}^{\mathcal{PC}})^{+}_{A|\sim B}$

($*_3$) If $(\mathcal{K}^{\mathcal{PC}})^{+}_{A|\sim B} \not\models_{pr} \top \mid\sim \bot$, then $\mathcal{C}((\mathcal{K}^{\mathcal{PC}})^{+}_{A|\sim B}) \subseteq \mathcal{K}^{*}_{A|\sim B}$

($*_4$) $A \mid\sim B \in \mathcal{K}^{*}_{A|\sim B}$

($*_5$) If $A \mid\sim B \equiv_{\mathcal{P}} A' \mid\sim B'$, then $\mathcal{K}^{*}_{A|\sim B} = \mathcal{K}^{*}_{A'|\sim B'}$

($*_6$) If $A \mid\sim B \not\models_{pr} \top \mid\sim \bot$, then $\top \mid\sim \bot \notin \mathcal{K}^{*}_{A|\sim B}$

As the following observation shows, revision for $\mathcal{C}(\cdot)$ can be defined in terms of revision of the monotonic core.

Observation 5 (Casini and Meyer 2017, Th. 6) $*$ *is a revision operator satisfying* $(*_1)$–$(*_6)$ *iff there is a preferential revision* $\bullet$ *satisfying* $(\bullet_1)$–$(\bullet_6)$ *s.t.* $\mathcal{K}^{*}_{A|\sim B} = \mathcal{C}((\mathcal{K}^{\mathcal{PC}})^{\bullet}_{A|\sim B})$.

The postulates for the revision operations $\otimes$ for preserving coherence under *spc*-operators are $(\otimes_1)$, $(\otimes_2)$, $(\otimes_4)$ and $(\otimes_5)$, which are direct reformulations of, respectively, $(*_1)$, $(*_2)$, $(*_4)$, and $(*_5)$, plus the following two:

($\otimes_3$) If $\mathcal{K}^{\mathcal{PC}} \not\models_{pr} A \mid\sim \neg B$, then $\mathcal{C}((\mathcal{K}^{\mathcal{PC}})^{+}_{A|\sim B}) \subseteq \mathcal{K}^{\otimes}_{A|\sim B}$

($\otimes_6$) If $A \mid\sim B \not\models_{pr} A \mid\sim \bot$, then $\mathcal{K}^{\otimes}_{A|\sim B} \not\models_{\mathcal{C}} A \mid\sim \bot$

And then we obtain the analogue of Observation 5:

Observation 6 (Casini and Meyer 2017, Theorem 7) *An operator* $\otimes$ *is a revision operator satisfying* $(\otimes_1)$–$(\otimes_6)$ *iff there is a preferential revision* $\circ$ *satisfying* $(\circ_1)$–$(\circ_6)$ *s.t.* $\mathcal{K}^{\otimes}_{A|\sim B} = \mathcal{C}((\mathcal{K}^{\mathcal{PC}})^{\circ}_{A|\sim B})$.

The latter Observation is valid only for the *spc*-operators that satisfy an extra property, that is satisfied by most of the prominent entailment relations in the preferential framework (Casini and Meyer 2017, Sect. 6): for every A and every $\mathcal{K}$, $A \mid\sim \bot \in \mathcal{K}$ iff $A \mid\sim \bot \in \mathcal{K}^{\mathcal{PC}}$.

3 A Semantics for the Basic Operations

In this section we provide a semantic characterisation of the classes of operators presented in Section 2.3. Once we have characterised preferential contraction and expansion, the results for the revision operators will follow immediately from the Observations in the previous section.

As it is usually the case, the semantic characterisation of the expansion operator is straightforward:

$$\mathcal{K}^{+c}_{A|\sim B} \equiv_{\text{def}} \{C \mid\sim D \mid \mathcal{P} \Vdash C \mid\sim D \\ \text{for every } \mathcal{P} \in [\![\mathcal{K}]\!] \cap [\![A \mid\sim B]\!]\}.$$

where $[\![\mathcal{K}]\!]$ and $[\![A \mid\sim B]\!]$ are the sets of all preferential models of $\mathcal{K}$ and $A \mid\sim B$, respectively.

Proposition 2 *For every conditional theory* $\mathcal{K}$ *and every conditional* $A \mid\sim B$, $\mathcal{K}^{+c}_{A|\sim B} = \mathcal{K}^{+}_{A|\sim B}$.

The proof is immediate from the fact that $\mathcal{K}^{+c}_{A|\sim B}$ contains *all* the preferential models satisfying $\mathcal{K}$ and $A \mid\sim B$. Moving to the class of preferential contractions, its semantic characterisation is quite straightforward; characterisations of basic contraction on the same line have been done also in the propositional case (Grove 1988; Hansson 1999). Let $]A \mid\sim B[$ denote the set of the *counter-models* of $A \mid\sim B$:

$$]A \mid\sim B[\equiv_{\text{def}} \{\mathcal{P} \in \mathfrak{U} \mid \mathcal{P} \not\Vdash A \mid\sim B\}$$

Let $c : \mathscr{P}(\mathcal{L}^{|\sim}) \times (\mathcal{L}^{|\sim}) \longrightarrow \mathscr{P}(\mathfrak{U})$ be a choice function. For every $\mathcal{K}$ and every pair of conditionals $A \mid\sim B$, $C \mid\sim D$, $c(\mathcal{K}, A \mid\sim B)$ is such that:

(c1) If $\mathcal{K} \not\models_{pr} A \mid\sim B$, then $c(\mathcal{K}, A \mid\sim B) = \varnothing$;

(c2) If $\mathcal{K} \models_{pr} A \mid\sim B$ and $\not\models_{pr} A \mid\sim B$, then $\varnothing \neq c(\mathcal{K}, A \mid\sim B) \subseteq {]}A \mid\sim B{[}$;

(c3) If $A \mid\sim B \equiv_{\mathcal{P}} C \mid\sim D$, then $\mid\sim_{c(\mathcal{K}, A\mid\sim B)} = \mid\sim_{c(\mathcal{K}, C\mid\sim D)}$;

(c4) If $\models_{pr} A \mid\sim B$, $c(\mathcal{K}, A \mid\sim B) = \varnothing$.

Notice that, under Proposition 4 below, in Property (c3) above we refer to the conditionals holding in the selected interpretations ($\mid\sim_{c(\mathcal{K}, A\mid\sim B)}$), instead of referring to the selected interpretations themselves ($c(\mathcal{K}, A \mid\sim B)$), since distinct sets of models can characterise the same set of conditionals. Using the function c, we can define a semantic contraction operator $\dot{-}c$ as follows:

Definition 1 *Given a knowledge base $\mathcal{K}$ and a conditional $A \mid\sim B$, for every conditional $C \mid\sim D$:*

$$C \mid\sim D \in \mathcal{K}_{A\mid\sim B}^{\dot{-}c} \quad iff \quad \begin{array}{l} \mathcal{P} \Vdash C \mid\sim D \text{ for every} \\ \mathcal{P} \in [\![\mathcal{K}]\!] \cup c(\mathcal{K}, A \mid\sim B) \end{array}$$

The following representation theorem connects $\dot{-}$ to $\dot{-}c$.

Theorem 1 *For every basic preferential contraction operator $\dot{-}$ there is a semantic contraction operator $\dot{-}c$ s.t. for every $\mathcal{K}$ and every conditional $A \mid\sim B$:*

$$\mathcal{K}_{A\mid\sim B}^{\dot{-}} = \mathcal{K}_{A\mid\sim B}^{\dot{-}c}$$

Conversely, for every semantic contraction operator $\dot{-}c$ there is a basic preferential contraction operator $\dot{-}$ s.t., for every $\mathcal{K}$ and every conditional $A \mid\sim B$:

$$\mathcal{K}_{A\mid\sim B}^{\dot{-}c} = \mathcal{K}_{A\mid\sim B}^{\dot{-}}$$

It is also possible to give a characterisation of the class of *preferential withdrawals*, i.e., the contraction operations satisfying $(\dot{-}_1)$–$(\dot{-}_5)$, but not $(\dot{-}_6)$ (recovery). It is sufficient to define a class of choice functions $c^w(\cdot)$ that differ from $c(\cdot)$ in that (c4) is dropped and (c2) is reformulated.

$(c^w 1)$ if $\mathcal{K} \not\models_{pr} A \mid\sim B$, then $c(\mathcal{K}, A \mid\sim B) = \varnothing$;

$(c^w 2)$ if $\mathcal{K} \models_{pr} A \mid\sim B$ and $\not\models_{pr} A \mid\sim B$, then $c(\mathcal{K}, A \mid\sim B) \cap {]}A \mid\sim B{[} \neq \varnothing$.

$(c^w 3)$ if $A \mid\sim B \equiv_{\mathcal{P}} C \mid\sim D$, then $\mid\sim_{c(\mathcal{K}, A\mid\sim B)} = \mid\sim_{c(\mathcal{K}, C\mid\sim D)}$.

Proposition 3 *For every preferential withdrawal operator $\dot{-}^w$ there is a semantic contraction operator $\dot{-}c^w$ s.t. for every $\mathcal{K}$ and every $A \mid\sim B$:*

$$\mathcal{K}_{A\mid\sim B}^{\dot{-}^w} = \mathcal{K}_{A\mid\sim B}^{\dot{-}c^w}$$

Conversely, for every semantic contraction operator $\dot{-}c^w$ there is a basic preferential contraction operator $\dot{-}^w$ s.t. for every $\mathcal{K}$ and every $A \mid\sim B$:

$$\mathcal{K}_{A\mid\sim B}^{\dot{-}c^w} = \mathcal{K}_{A\mid\sim B}^{\dot{-}^w}$$

Once we have representation results for expansion and basic contraction, the analogous results for both classes of preferential revision presented by Casini and Meyer (2017) are straightforward. We just need to consider a semantic difference between the preferential interpretations and the propositional ones: the former ones do not necessarily satisfy *complete* theories, so there are pairs of interpretations $\mathcal{P}_1, \mathcal{P}_2$ s.t. $\mid\sim_{\mathcal{P}_1} \subset \mid\sim_{\mathcal{P}_2}$. It is an immediate consequence of the monotonicity of $\mathcal{PC}$ and the representation result by Kraus et al. (1990, Theorem 3), proving that for every preferential theory there is a preferential model characterising it. As a consequence, we have that a preferential theory is not characterised by a unique set of interpretations.

Proposition 4 *There are sets of preferential models $\mathfrak{R}_1, \mathfrak{R}_2$ such that $\mathfrak{R}_1 \neq \mathfrak{R}_2$ but $\mathfrak{R}_1$ and $\mathfrak{R}_2$ satisfy exactly the same set of conditionals, i.e., $\mid\sim_{\mathfrak{R}_1} = \mid\sim_{\mathfrak{R}_2}$.*

Given this result, we introduce additional notation: for a given a set of interpretations $\mathfrak{P}$, with $[\![\mathfrak{P}]\!]$ we denote the set of all models of the conditionals that are satisfied by all interpretations in $\mathfrak{P}$ ($[\![\mathfrak{P}]\!] \equiv_{\mathrm{def}} [\![\mid\sim_{\mathfrak{P}}]\!]$).

We start by addressing the class of operators $\bullet$, which add a conditional to a knowledge base while preserving logical consistency. We define the corresponding semantic revision operator as follows:

$$\mathcal{K}_{A\mid\sim B}^{\bullet c} \equiv_{\mathrm{def}} (\mathcal{K}_{\top\mid\sim A \wedge \neg B}^{\dot{-}c})_{A\mid\sim B}^{+c} \qquad (6)$$

That is,

$$\mathcal{K}_{A\mid\sim B}^{\bullet c} = [\![([\![\mathcal{K}]\!] \cup c(\mathcal{K}, \top \mid\sim A \wedge \neg B))]\!] \cap [\![A \mid\sim B]\!] \quad (7)$$

Proposition 5 *For every revision operator $\bullet$ there is a semantic revision operator $\bullet c$ s.t. for every $\mathcal{K}$ and $A \mid\sim B$:*

$$\mathcal{K}_{A\mid\sim B}^{\bullet} = \mathcal{K}_{A\mid\sim B}^{\bullet c}$$

Conversely, for every semantic revision operator $\bullet c$ there is a revision operator $\bullet$ s.t. for every $\mathcal{K}$ and every $A \mid\sim B$:

$$\mathcal{K}_{A\mid\sim B}^{\bullet c} = \mathcal{K}_{A\mid\sim B}^{\bullet}$$

Notice that, unlike in the propositional setting, under Proposition 4 we need to apply the closure operation $[\![\cdot]\!]$ on the interpretations selected by the contraction (that is, using $[\![([\![\mathcal{K}]\!] \cup c(\mathcal{K}, \top \mid\sim A \wedge \neg B)]\!]$ instead of simply $[\![\mathcal{K}]\!] \cup c(\mathcal{K}, \top \mid\sim A \wedge \neg B)$) before intersecting with $[\![A \mid\sim B]\!]$. Otherwise we could end up with an empty set of interpretations, violating Postulate ($\bullet_6$). The following example should clarify the issue.

Example 2 *Let $\mathcal{P}_2$ be built by a single propositional valuation w, while $\mathcal{P}_1$ is composed by a pair of distinct valuations w, v, with $w \prec v$. It follows that $\mid\sim_{\mathcal{P}_1} \subset \mid\sim_{\mathcal{P}_2}$. Let $\mathfrak{P}_1 = \{\mathcal{P}_1\}$ and $\mathfrak{P}_2 = \{\mathcal{P}_1, \mathcal{P}_2\}$. $\mathfrak{P}_1 \neq \mathfrak{P}_2$, but $\mid\sim_{\mathfrak{P}_1} = \mid\sim_{\mathfrak{P}_2}$. Let $\mathcal{K} = \mid\sim_{\mathfrak{P}_1} = \mid\sim_{\mathfrak{P}_2}$, and assume we want to expand $\mathcal{K}$ with $\overline{v} \mid\sim \bot$. Both $\mathfrak{P}_1$ and $\mathfrak{P}_2$ characterise the same preferential theory $\mathcal{K}$, and $\overline{v} \mid\sim \bot$ is preferentially consistent with $\mathcal{K}$, since we have $\mathcal{P}_2$ satisfying both $\mathcal{K}$ and $\overline{v} \mid\sim \bot$. However, from the semantical point of view, if we model $\mathcal{K}_{A\mid\sim B}^{+}$ as $\mathfrak{P}_1 \cap [\![A \mid\sim B]\!]$ we obtain a consistent result ($\mathcal{P}_2$), while if we use $\mathfrak{P}_2 \cap [\![A \mid\sim B]\!]$ we obtain $\varnothing$.*

We proceed in a analogous way if we are interested in modelling a revision operation that, while adding $A \mid\sim B$, preserves not only consistency but also coherence w.r.t. the formula A. As seen in Section 2.3, the class of operations $\circ$ was characterised only w.r.t. the class of preferential withdrawals (Casini and Meyer 2017, Proposition 3), not w.r.t. the preferential basic contractions. We reformulate the class of the basic revision operators preserving coherence using semantic contractions: that is, as

$$\mathcal{K}^{\circ c^w}_{A\mid\sim B} \equiv_{\text{def}} [\![(\![\mathcal{B}]\!] \cup c^w(\mathcal{K}, A \mid\sim \neg B)]\!] \cap [\![A \mid\sim B]\!] \quad (8)$$

Proposition 6 *For every revision operator $\circ$ there is a semantic revision operator $\circ c^w$ s.t., for every $\mathcal{K}$ and $A \mid\sim B$,*

$$\mathsf{K}^{\circ}_{A\mid\sim B} = \mathsf{K}^{\circ c^w}_{A\mid\sim B}.$$

Conversely, for every semantic revision operator $\circ c^w$ there is a revision operator $\circ$ s.t., for every $\mathcal{K}$ and every $A \mid\sim B$,

$$\mathsf{K}^{\circ c^w}_{A\mid\sim B} = \mathsf{K}^{\circ}_{A\mid\sim B}.$$

Also, using this semantic characterisation, it becomes very easy to strengthen the previous representation result (Casini and Meyer 2017, Proposition 3), using the class of preferential basic contractions $\dot{-}$ instead of simply using preferential withdrawals. First, we can prove that we can use the class of choice functions c instead of the choice functions c^w, still referring to the same semantic construction:

$$\mathcal{K}^{\circ c}_{A\mid\sim B} \equiv_{\text{def}} [\![(\![\mathcal{B}]\!] \cup c(\mathcal{K}, A \mid\sim \neg B)]\!] \cap [\![A \mid\sim B]\!] \quad (9)$$

Theorem 2 *For every revision operator $\circ$ there is a semantic revision operator $\circ c$ s.t., for every $\mathcal{K}$ and $A \mid\sim B$,*

$$\mathsf{K}^{\circ}_{A\mid\sim B} = \mathsf{K}^{\circ c}_{A\mid\sim B}.$$

Conversely, for every semantic revision operator $\circ c$ there is a revision operator $\circ$ s.t., for every $\mathcal{K}$ and every $A \mid\sim B$,

$$\mathsf{K}^{\circ c}_{A\mid\sim B} = \mathsf{K}^{\circ}_{A\mid\sim B}.$$

An easy consequence of Theorems 1 and 2 is a strenghtening of Proposition 3 by Casini and Meyer (2017), using basic preferential contractions (satisfying $(\dot{-}_1)$–$(\dot{-}_6)$) instead of preferential withdrawals $((\dot{-}_1)$–$(\dot{-}_5))$.

Theorem 3 *For every revision operator $\circ$ there is a preferential basic contraction operator $\dot{-}$ s.t., for every $\mathcal{K}$ and every $A \mid\sim B$,*

$$\mathsf{K}^{\circ}_{A\mid\sim B} = (\mathsf{K}^{\dot{-}}_{A\mid\sim \neg B})^+_{A\mid\sim B}$$

Vice versa, for every preferential basic contraction operator $\dot{-}$ there is a revision operator $\circ$ s.t., for every $\mathcal{K}$ and every $A \mid\sim B$,

$$\mathsf{K}^{\circ}_{A\mid\sim B} = (\mathsf{K}^{\dot{-}}_{A\mid\sim \neg B})^+_{A\mid\sim B}$$

Regarding the class of *spc*-closures, we have seen in Section 2.2 that each closure $\mathcal{C}(\cdot)$ corresponds on the semantic side to a choice function $s_{\mathcal{C}} : \mathscr{P}(\mathfrak{U}) \longrightarrow \mathscr{P}(\mathfrak{U})$ s.t. $A \mid\sim B \in \mathcal{C}(\mathcal{K})$ iff $A \mid\sim B$ is satisfied by every preferential interpretation in $s_{\mathcal{C}}([\![\mathcal{K}]\!])$. Given a theory $\mathcal{K}$, closed

under $\mathcal{C}(\cdot)$, and its monotonic core $\mathcal{K}^{\mathcal{PC}}$, the class of revision operators $*$ for consistency preservation is characterised using the class of preferential revision operators $\bullet$ as $\mathcal{K}^{*}_{A\mid\sim B} = \mathcal{C}((\mathcal{K}^{\mathcal{PC}})^{\bullet}_{A\mid\sim B})$ (see Observation 5). Given (6) and Proposition 5, it is immediate to conclude that the class of non-monotonic revision operators $*$ for a *spc*-closure $\mathcal{C}(\cdot)$ can be characterised via the class of preferential contractions $\dot{-}c$ as

$$\mathcal{K}^{*c}_{A\mid\sim B} \equiv_{\text{def}} \mathcal{C}(((\mathcal{K}^{\mathcal{PC}})^{\dot{-}c}_{\top\mid\sim A \wedge \neg B})^{+c}_{A\mid\sim B}) \quad (10)$$

That is, $\mathcal{K}^{*c}_{A\mid\sim B}$ is equal to:

$$s_{\mathcal{C}}([\![([\![\mathcal{K}^{\mathcal{PC}}]\!] \cup c(\mathcal{K}, \top \mid\sim A \wedge \neg B)]\!] \cap [\![A \mid\sim B]\!]) \quad (11)$$

We proceed analogously for the class of non-monotonic revisions $\otimes$ to preserve coherence, characterised via $\circ$ as $\mathcal{K}^{\otimes}_{A\mid\sim B} = \mathcal{C}((\mathcal{K}^{\mathcal{PC}})^{\circ}_{A\mid\sim B})$ (see Observation 6).

Given (9) and Theorem 2, the class of revision operators $\otimes$ for an *spc*-closure $\mathcal{C}(\cdot)$ can be characterised via the class of preferential contractions $\dot{-}c$ as

$$\mathcal{K}^{\otimes c}_{A\mid\sim B} \equiv_{\text{def}} \mathcal{C}(((\mathcal{K}^{\mathcal{PC}})^{\dot{-}c}_{A\mid\sim \neg B})^{+c}_{A\mid\sim B}) \quad (12)$$

That is,

$$\mathcal{K}^{*c}_{A\mid\sim B} = s_{\mathcal{C}}([\![([\![\mathcal{K}^{\mathcal{PC}}]\!] \cup c(\mathcal{K}, A \mid\sim \neg B))]\!] \cap [\![A \mid\sim B]\!]) \quad (13)$$

4 Beyond the Basic Postulates

We move to explore possible classes of interesting operators beyond the basic forms of contraction and revision. As seen in Section 2, in the AGM approach there are two extra postulates for contraction, $(-_7)$ and $(-_8)$, and for revision, $(*_7)$ and $(*_8)$, that can be associated with an ordering between the formulas (entrenchment) and an order between the interpretations. $(-_7)$ and $(-_8)$ define a way of preferring some contraction operations to others, in particular favouring the contraction operators that respect the relevance of certain pieces of information, corresponding to more entrenched formulas and/or more preferred interpretations. We aim at an analogous characterisation in the conditional framework of operations respecting the relevance of the information.

Since Postulates $(-_7)$ and $(-_8)$, and $(*_7)$ and $(*_8)$ use conjunction, their translation into the conditional framework is not immediate. Gärdenfors' entrenchment contraction (1) uses disjunction, but Proposition 1 gives us an alternative characterisation in terms of partial-meet contraction. Katsuno and Mendelson's (1991) semantic characterisation (see Section 2 here) looks more immediately implementable in the conditional framework as a ranking of preferential models, but we have to take under consideration Proposition 4 in the definition of the ranking.

Given the difficulties in reformulating the postulates $(-_7)$ and $(-_8)$ in the conditional framework, we leave them aside for the moment. Still, our goal here is to define classes of contraction and revision operators that constrain the basic operators in a form that is analogous to what we obtain in the propositional case when adding $(-_7)$ and $(-_8)$. Here we propose to bypass the postulates and investigate the correspondence between an entrenchment relation over the conditionals and a preference relation over the preferential interpretations, reformulating in the conditional environment

Gärdenfors' entrenchment and Katsuno and Mendelson's ranking of interpretations. In this way, we can give a characterisation of a class of contraction operators constraining the class of basic contraction operators in a similar way as postulates $(-_7)$ and $(-_8)$ do in the propositional framework.

Conditional Entrenchment. We define in the conditional framework an entrenchment relation $\geqslant_e \colon \mathcal{L}^{\mid\sim} \times \mathcal{L}^{\mid\sim}$ satisfying the following constraints:

PE1 if $C \mid\sim D \leqslant_e E \mid\sim F$ and $E \mid\sim F \leqslant_e G \mid\sim H$, then $C \mid\sim D \leqslant_e G \mid\sim H$ (Transitivity)

PE2 if $\{C_1 \mid\sim D_1, \ldots, C_n \mid\sim D_n\} \vDash_{pr} E \mid\sim F$, then there is an i, $1 \leqslant i \leqslant n$, s.t. $C_i \mid\sim D_i \leqslant_e E \mid\sim F$ (Dominance)

PE3 Either $C \mid\sim D \leqslant_e E \mid\sim F$ or $E \mid\sim F \leqslant_e C \mid\sim D$ (Connectivity)

PE4 If $\mathcal{K}$ is consistent, then $C \mid\sim D \notin \mathcal{K}$ iff $C \mid\sim D \leqslant_e E \mid\sim F$ for every $E \mid\sim F$ (Minimality)

PE5 If $E \mid\sim F \leqslant_e C \mid\sim D$ for every $E \mid\sim F$, then $\vDash_{pr} C \mid\sim D$ (Maximality)

Let $A \mid\sim B^{>e} \equiv_{\text{def}} \{C \mid\sim D \mid A \mid\sim B <_e C \mid\sim D\}$. By Compactness of $\mathcal{PC}(\cdot)$ and Dominance, $A \mid\sim B^{>e} = \mathcal{PC}(A \mid\sim B^{>e})$, that is, $A \mid\sim B^{>e}$ is a preferential theory.

For lack of expressivity, we substitute *Conjunctiveness* with *Connectivity*. In Gärdenfors' entrenchment the latter is implied by (E1)-(E3) (Gärdenfors 1988, Lemma 4.20). *Connectivity*, (E1) and the property

$$\text{If } C \leqslant_e D \text{ and } C \leqslant_e E, \text{ then } C \leqslant_e D \wedge E$$

together imply *Conjunctiveness* (Gärdenfors 1988, Lemma 4.21). A restricted forms of the above property and of the *Conjunctiveness* property are derivable from **PE1**-**PE3**.

Proposition 7 *The following two properties are a consequence of **PE1**-**PE2**. For every tuple of propositions A, B, C, D, E:*

(a) *If $A \mid\sim B \leqslant_e C \mid\sim D$ and $A \mid\sim B \leqslant_e C \mid\sim E$, then $A \mid\sim B \leqslant_e C \mid\sim D \wedge E$;*

(b) *Either $A \mid\sim B \leqslant_e A \mid\sim B \wedge C$ or $A \mid\sim C \leqslant_e A \mid\sim B \wedge C$.*

Let $\mathsf{K}^{>e}_{C \mid\sim D} \equiv_{\text{def}} \{E \mid\sim F \mid C \mid\sim D <_e E \mid\sim F\}$ and $\mathsf{K} \Downarrow C \mid\sim D$ be the set of the preferred remainder sets w.r.t. the entailment relation, that is:

$$\mathsf{K} \Downarrow A \mid\sim B \equiv_{\text{def}} \{\mathsf{K}' \in \mathsf{K} \downarrow A \mid\sim B \mid \mathsf{K}^{>e}_{A \mid\sim B} \subseteq \mathsf{K}'\}$$

We define the contraction $\dot{-}e$ using preferential entrenchment as follows:

$$\mathsf{K}^{\dot{-}e}_{A \mid\sim B} \equiv_{\text{def}} \bigcap \mathsf{K} \Downarrow A \mid\sim B \tag{14}$$

This entrenchment-based conditional contraction $\dot{-}e$ is a reformulation in the conditional framework of Gärdenfors' entrenchment-based propositional contraction $-e$: $\mathsf{K}^{>e}_{A \mid\sim B}$ corresponds to Rott's contraction (3) in the conditional environment, and consequently (14) is a conditional reformulation of the propositional entrenched-based contraction $-e$, following the version used in Proposition 1.

Semantic Characterisation. On the semantic side, we introduce a ranking of preferential interpretations in the spirit of Katsuno and Mendelzon (1991). With $\mathfrak{U}$ being the set of all the preferential interpretation for language $\mathcal{L}^{\mid\sim}$, let $r \colon \mathfrak{U} \mapsto \mathbb{N}$ be a function associating with every model a natural number, and let $\mathfrak{L}_i$ indicate the i-th layer of the ranking ($\mathfrak{L}_i \equiv_{\text{def}} \{\mathcal{P} \in \mathfrak{U} \mid r(\mathcal{P}) = i\}$). Let $\overline{\mathfrak{L}_i}$ be the set of conditionals satisfied by all the interpretations in $\mathfrak{L}_i$. Given a set of conditionals $\mathcal{B}$ and an interpretation $\mathcal{P}$, let $\mathcal{B}^{\mathcal{P}} \equiv_{\text{def}} \{C \mid\sim D \in \mathsf{K} \mid \mathcal{P} \Vdash C \mid\sim D\}$. We want the ranking to satisfy the following constraints:

(r1) $\forall i \in \mathbb{N}$ and $i > 0$, if $\mathfrak{L}_i \neq \varnothing$, then $\mathfrak{L}_{i-1} \neq \varnothing$;

(r2) $\forall \mathcal{P} \in \mathfrak{U}$, if $\mathcal{P} \Vdash \overline{\mathfrak{L}_0}$, then $\mathcal{P} \in \mathfrak{L}_0$;

(r3) $\forall \mathcal{P}, \mathcal{P}', \mathcal{P}^* \in \mathfrak{U}$ and $i > 0$, if $\mathcal{P} \in \mathfrak{L}_i$ and $\mathcal{P}' \in \mathfrak{L}_{i-1}$, then if $\overline{\mathfrak{L}_0}^{\mathcal{P}^*} \supseteq \overline{\mathfrak{L}_0}^{\mathcal{P}} \cap \overline{\mathfrak{L}_0}^{\mathcal{P}'}$, and $\mathcal{P}^* \notin \mathfrak{L}_j$ for any $j < i$, then $\mathcal{P}^* \in \mathfrak{L}_i$.

Let us consider $\overline{\mathfrak{L}_0}$ as our knowledge base $\mathcal{K}$, and let, for every $\mathcal{P}, \mathcal{P}' \in \mathfrak{U}$, $\mathcal{P} \leqslant_r \mathcal{P}'$ iff $r(\mathcal{P}) \leqslant r(\mathcal{P}')$. It is very easy to check that $\leqslant_r$ is a faithful assignment (as described in Sect. 2.1). That is, (1) if $\mathcal{P} \Vdash \overline{\mathfrak{L}_0}$ and $\mathcal{P}' \Vdash \overline{\mathfrak{L}_0}$, then $\mathcal{P} \simeq_r \mathcal{P}'$; and (2) if $\mathcal{P} \Vdash \overline{\mathfrak{L}_0}$ and $\mathcal{P}' \not\Vdash \overline{\mathfrak{L}_0}$, then $\mathcal{P} <_r \mathcal{P}'$.

(r3) is necessary to deal with the peculiarity of preferential models, as indicated by Proposition 4, and it is equivalent to the combination of the following two properties:

(r3') $\forall \mathcal{P}, \mathcal{P}^* \in \mathfrak{U}$ and $i > 0$, if $\mathcal{P} \in \mathfrak{L}_i$, $\overline{\mathfrak{L}_0}^{\mathcal{P}} \subseteq \overline{\mathfrak{L}_0}^{\mathcal{P}^*}$, and $\mathcal{P}^* \notin \mathfrak{L}_j$ for any $j < i$, then $\mathcal{P}^* \in L_i$.

(r3'') $\forall \mathcal{P}, \mathcal{P}' \in \mathfrak{U}$ and $i > 0$, if $\mathcal{P} \in L_i$ and $R' \in L_j$, $j \leqslant i$, then there is an interpretation $\mathcal{P}^* \in L_i$ s.t. $\overline{\mathfrak{L}_0}^{\mathcal{P}^*} = \overline{\mathfrak{L}_0}^{\mathcal{P}} \cap \overline{\mathfrak{L}_0}^{\mathcal{P}'}$.

Proposition 8 *The condition $(r3)$ is equivalent to the combination of the conditions $(r3')$ and $(r3'')$.*

Given a set of interpretations $\mathfrak{P} \subseteq \mathfrak{U}$, a ranking r, and a preferentially closed set of conditionals $\mathcal{K}$, let $\min_r(\mathfrak{P}) = \{\mathcal{P} \in \mathfrak{P} \cap \mathfrak{L}_i \mid \mathfrak{P} \cap \mathfrak{L}_i \neq \varnothing$ and $\forall j < i, \mathfrak{P} \cap \mathfrak{L}_j = \varnothing\}$ and $\max_{\mathsf{K}}(\mathfrak{P}) = \{\mathcal{P} \in \mathfrak{P} \mid \nexists \mathcal{P}' \in \mathfrak{P}$ s.t. $\mathsf{K}^{\mathcal{P}'} \supset \mathsf{K}^{\mathcal{P}}\}$.

We define the choice function c_r over this ranking as:

$$c_r(\overline{L_0}, C \mid\sim D) \equiv_{\text{def}} \max_{\overline{L_0}}(\min_r(]C \mid\sim D[))$$

That is, given a knowledge base $\mathcal{K}$ and a conditional $C \mid\sim D$ to be contracted, we first pick the preferred models that do not satisfy $C \mid\sim D$ ($\min_r(]C \mid\sim D[)$). Between such models, we choose only the ones that satisfy more conditionals from $\overline{L_0}$ ($\max_{\overline{L_0}}(\min_r(]C \mid\sim D[))$); the use of $\max_{\overline{L_0}}(.)$ is due to the property of preferential interpretations indicated in Proposition 4. Having defined the selection function c_r, we can define contraction in the usual way.

$$C \mid\sim D \in \mathcal{K}^{\dot{-}c_r}_{A \mid\sim B} \text{ iff } \mathcal{P} \Vdash C \mid\sim D \text{ for every}$$
$$\mathcal{P} \in [\mathcal{K}] \cup c_r(\mathcal{K}, A \mid\sim B).$$

Now we can prove the correspondence between the two classes of contraction operations we have defined using the entrenchment and semantic rankings, respectively.

Theorem 4 *For every entrenchment contraction operator $\dot{-}e$ there is a semantic contraction operator $\dot{-}c_r$ s.t., for every KB $\mathcal{K}$ and every conditional $A \mathrel{|\!\sim} B$,*

$$\mathcal{K}_{A|\sim B}^{\dot{-}e} = \mathcal{K}_{A|\sim B}^{\dot{-}c_r}.$$

Conversely, for every semantic contraction operator $\dot{-}c_r$ there is an entrenchment contraction operator $\dot{-}e$ s.t., for every KB $\mathcal{K}$ and every conditional $A \mathrel{|\!\sim} B$,

$$\mathcal{K}_{A|\sim B}^{\dot{-}c_r} = \mathcal{K}_{A|\sim B}^{\dot{-}e}.$$

Our aim was for a full characterisation of AGM-like operations in the framework of preferential conditional logic. Ideally, we would also have provided postulates corresponding to $(-_7)$, $(-_8)$, $(*_7)$, and $(*_8)$, but since they refer to the conjunction of formulae, their translation in a conditional language is not immediate, and we have omitted that. We have considered two alternative and equivalent characterizations of the class of contraction operations satisfying $(\dot{-}_1) - (\dot{-}_8)$, using Gardenfors' entrenchment and Katsuno and Mendelson's rankings, and we have reformulated them accordingly for the conditional framework, showing that the correspondence is preserved, as Theorem 4 proves.

We can easily give a semantic characterisation of the revision operators, using the choice functions c_r defined earlier in this section, instead of the choice functions c (which were defined in Section 3). First the ones for preferential revision. The only constraint is that, starting from a preferential theory $\mathcal{K}$, the ranking r must be defined with $\mathfrak{L}_0 = [\![\mathcal{K}]\!]$.

$$[\![\mathcal{K}_{A|\sim B}^{\bullet c_r}]\!] = [\![([\![\mathcal{K}]\!] \cup c_r(\mathcal{K}, \top \mathrel{|\!\sim} A \wedge \neg B)]\!] \cap [\![A \mathrel{|\!\sim} B]\!]$$

That is,

$$[\![\mathcal{K}_{A|\sim B}^{\bullet c_r}]\!] = [\![([\![\mathcal{K}]\!] \cup \max_{\mathcal{K}}(\min_{r}(]\top \mathrel{|\!\sim} A \wedge \neg B[)))]\!] \cap [\![A \mathrel{|\!\sim} B]\!]$$

which implies

$$[\![\mathcal{K}_{A|\sim B}^{\bullet c_r}]\!] = [\![\max_{\mathcal{K}}(\min_{r}(]\top \mathrel{|\!\sim} A \wedge \neg B[))]\!] \cap [\![A \mathrel{|\!\sim} B]\!]$$

Proceeding analogously for $\circ$, we obtain

$$[\![\mathcal{K}_{A|\sim B}^{\circ c_r}]\!] = [\![([\![\mathcal{K}]\!] \cup \max_{\mathcal{K}}(\min_{r}(]A \mathrel{|\!\sim} \neg B[)))]\!] \cap [\![A \mathrel{|\!\sim} B]\!]$$

Given a theory $\mathcal{K}$ closed under a *spc*-closure $\mathcal{C}$, the choice function $s_{\mathcal{C}}$ associated at $\mathcal{C}$, and the monotonic core $\mathcal{K}^{\mathcal{PC}}$, using the two last equations we can define the two corresponding classes of non-monotonic revision operators. The two classes of operators $*$ and $\otimes$ can then be defined as follows. $\mathcal{K}_{A|\sim B}^{*c_r}$ corresponds to the conditionals satisfied by the set of interpretations

$$s_{\mathcal{C}}([\![\max_{\mathcal{K}^{\mathcal{PC}}}(\min_{r}(]\top \mathrel{|\!\sim} A \wedge \neg B[))]\!] \cap [\![A \mathrel{|\!\sim} B]\!])$$

$\mathcal{K}_{A|\sim B}^{\otimes c_r}$ corresponds to the conditionals satisfies by the set

$$s_{\mathcal{C}}([\![([\![\mathcal{K}]\!] \cup \max_{\mathcal{K}^{\mathcal{PC}}}(\min_{r}(]A \mathrel{|\!\sim} \neg B[)))]\!] \cap [\![A \mathrel{|\!\sim} B]\!])$$

Since the contraction is to be defined on $\mathcal{K}^{\mathcal{PC}}$, the only constraint is for the ranking r to be built with $\mathfrak{L}_0 = [\![\mathcal{K}^{\mathcal{PC}}]\!]$.

5 Conclusions

This paper, in combination with previous work by Casini and Meyer (2017), provide the basis for the definition of belief change in a (preferential) non-monotonic framework. Building on the characterisation in terms of entrenchment and the ranked semantics of the AGM approach, we define and characterise preferential expansion, contraction and revision (for both consistency and coherence preservation) on the monotonic core of the class of the *spc*-operators.

Most of the previous works dealing with the revision of conditionals, such as those by Kern-Isberner (1999) and Wobcke (1995), is connected to the revision of subjunctive conditionals, which are used to represent the revision policies themselves. Here, instead, we do not consider the conditionals as linked to the revision policies, and leave their interpretation quite open (see Section 1). Other recent contributions to revision in a non-monotonic framework have come from Hunter (2016), dealing with highly implausible conditionals, and Delgrande and others (2013), that analysed belief change in the non-monotonic framework of ASP.

Regarding future developments, there are some obvious steps, such as modelling iterated revision and some forms of base revisions. An initial step in this direction has been made by Casini and Meyer (2016). In recent years there has been particular interest in preferential non-monotonic reasoning in the framework of Description Logics (Britz, Heidema, and Meyer 2008; Britz, Meyer, and Varzinczak 2011; Casini and Straccia 2010; 2013; Casini et al. 2014; Giordano et al. 2013; 2015; Lukasiewicz 2008), and in such a framework the notion of coherence, as we use it here, is very relevant. Hence a major goal is to reformulate the present approach for Description logics.

An important aspect of belief change in non-monotonic frameworks is contraction. The present paper and the previous work by Casini and Meyer (2017) are focused on revision for non-monotonic theories, and we have reduced it to revision of its monotonic core, that in turn can be analysed in terms of contraction and expansion of the monotonic core. Hence, the only form of contraction that we have analysed up to this point is the one related to the monotonic core, while we have still to develop an analysis of the contraction of a conditional from a non-monotonic closure. The interesting aspect is that if we want to eliminate some piece of defeasible information, we have two possibilities: either we contract information from the monotonic core (as our contraction operations do), or we can actually *add* information to it. Consider Example 1, Step 2, where $a \mathrel{|\!\sim} n \in \mathcal{C}(\mathcal{B}')$. Assume we are just interested in eliminating $a \mathrel{|\!\sim} n$ from our conclusions. In a non-monotonic framework we can proceed in two ways. Either we eliminate some pieces of information from $\mathcal{B}'$, preventing the derivation of $a \mathrel{|\!\sim} n$ (e.g., eliminating $v \mathrel{|\!\sim} n$ or $a \vdash v$), or we can add some information to $\mathcal{B}$ that conflicts with $a \mathrel{|\!\sim} n$ (e.g., adding $a \mathrel{|\!\sim} \neg n$, as it happens in the example). To the best of our knowledge, the only attempt to at modelling this kind of contraction (in the KLM framework) is by Booth and Paris (1998), in which they introduce negated conditionals $A \mathrel{|\!\not\sim} B$ into the language. However, a proper analysis of this kind of contraction operation still needs to be done.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

Booth, R., and Paris, J. 1998. A note on the rational closure of knowledge bases with both positive and negative knowledge. *Journal of Logic, Lang. and Inform.* 7(2):165–190.

Britz, K.; Heidema, J.; and Meyer, T. 2008. Semantic preferential subsumption. In Lang, J., and Brewka, G., eds., *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning, KR 2008*, 476–484. AAAI Press/MIT Press.

Britz, K.; Meyer, T.; and Varzinczak, I. 2011. Semantic foundation for preferential Description Logics. In Wang, D., and Reynolds, M., eds., *Proceedings of the 24th Australasian Joint Conference on Artificial Intelligence*, number 7106 in LNAI, 491–500. Springer.

Casini, G., and Meyer, T. 2016. Using defeasible information to obtain coherence. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning, KR 2016*, 537–540.

Casini, G., and Meyer, T. 2017. Belief change in a preferential non-monotonic framework. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 929–935.

Casini, G., and Straccia, U. 2010. Rational closure for defeasible Description Logics. In Janhunen, T., and Niemelä, I., eds., *Proceedings of the 12th European Conference on Logics in Artificial Intelligence (JELIA)*, number 6341 in LNCS, 77–90. Springer-Verlag.

Casini, G., and Straccia, U. 2013. Defeasible inheritance-based Description Logics. *J. Artif. Intell. Res. (JAIR)* 48:415–473.

Casini, G.; Meyer, T.; Moodley, K.; and Nortjé, R. 2014. Relevant closure: A new form of defeasible reasoning for description logics. In Fermé, E., and Leite, J., eds., *Logics in Artificial Intelligence*, 92–106. Springer International Publishing.

Delgrande, J.; Schaub, T.; Tompits, H.; and Woltran, S. 2013. A model-theoretic approach to belief change in answer set programming. *ACM Trans. Comput. Logic* 14(2):1–46.

Fermé, E. 1999. *Revising the AGM postulates*. Ph.D. Dissertation, University of Buenos Aires.

Gärdenfors, P., and Makinson, D. 1988. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning About Knowledge*, TARK '88, 83–95. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

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

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. L. 2013. A non-monotonic Description Logic for reasoning about typicality. *Artif. Intell.* 195:165–202.

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. L. 2015. Semantic characterization of rational closure: From propositional logic to Description Logics. *Artif. Intell. J.* 226:1–33.

Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17:157–170.

Hansson, S. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Publishers.

Hunter, A. 2016. Nearly counterfactual revision. In Khoury, R., and Drummond, C., eds., *Advances in Artificial Intelligence: 29th Canadian Conference on Artificial Intelligence, Canadian AI 2016, Victoria, BC, Canada, May 31 - June 3, 2016. Proceedings*. Cham: Springer International Publishing. 263–269.

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

Kern-Isberner, G. 1999. Postulates for conditional belief revision. In *Proceedings of the 16th International Joint Conference on Artifical Intelligence - Volume 1*, IJCAI'99, 186–191. San Francisco, CA, USA: Morgan Kaufmann Publ.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* 44:167–207.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artif. Intell.* 55:1–60.

Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15(1):61–82.

Levi, I. 1977. Subjunctives, dispositions and chances. *Synthese* 34:423–455.

Lukasiewicz, T. 2008. Expressive probabilistic Description Logics. *Artif. Intell.* 172(6-7):852–883.

Makinson, D. 1993. Five faces of minimality. *Studia Logica* 52:339–379.

Qi, G., and Hunter, A. 2007. Measuring incoherence in Description Logic-based ontologies. In *The Semantic Web*, volume 4825 of *LNCS*. Springer. 381–394.

Rott, H. 1991. Two methods of constructing contractions and revisions of knowledge systems. *J. of Phil. Logic* 20:149–173.

Wobcke, W. 1995. Belief revision, conditional logic and nonmonotonic reasoning. *Notre Dame J. Formal Logic* 36(1):55–102.