Rational Defeasible Belief Change

Giovanni Casini[1,2,5] , **Thomas Meyer**[2,5] , **Ivan Varzinczak**[3,4,5]
[1]ISTI-CNR, Italy
[2]University of Cape Town, South Africa
[3]CRIL, Univ. Artois-CNRS, France
[4]CS Division, Stellenbosch University, South Africa
[5]CAIR, South Africa
giovanni.casini@isti.cnr.it, tmeyer@cs.uct.ac.za, varzinczak@cril.fr

Abstract

We present a formal framework for modelling belief change within a non-monotonic reasoning system. Belief change and non-monotonic reasoning are two areas that are formally closely related, with recent attention being paid towards the analysis of belief change within a non-monotonic environment. In this paper we consider the classical AGM belief change operators, contraction and revision, applied to a defeasible setting in the style of Kraus, Lehmann, and Magidor. The investigation leads us to the formal characterisation of a number of classes of defeasible belief change operators. For the most interesting classes we need to consider the problem of iterated belief change, generalising the classical work of Darwiche and Pearl in the process. Our work involves belief change operators aimed at ensuring logical consistency, as well as the characterisation of analogous operators aimed at obtaining coherence—an important notion within the field of logic-based ontologies.

1 Introduction

Exploring the principles of belief change beyond the standard (propositional) AGM approach has been on the agenda since the early days of the field. The bulk of the effort in this direction has been driven by assuming alternative underlying logical languages that are either extensions of classical propositional logic or fragments thereof (Booth, Meyer, and Varzinczak 2009; Booth et al. 2011; Meyer, Lee, and Booth 2005; Delgrande and Peppas 2015; Zhuang et al. 2019; Zhuang et al. 2016; Delgrande and Wassermann 2013). Nevertheless, for most of the history of AGM-style belief change, the fundamental assumption of an underlying classical (Tarskian) consequence operator has not been challenged, until recently. A case in point is that of revising *defeasible knowledge bases* in which, in addition to classical statements, defeasible statements of the form $\alpha \mid\sim \beta$, read "typically, if α, then β", are also permitted.

Given well-known results in the literature establishing that belief change and non-monotonic reasoning can be defined in terms of each other (Gärdenfors and Makinson 1994), one could argue that an investigation of belief change in non-monotonic settings is superfluous. Nevertheless, a more careful analysis reveals that this is not the case. Indeed, among the consequences of defeasible knowledge bases is a mix of statements that follow non-monotonically and also statements that hold classically. A major challenge in the investigation of belief change in this setting is therefore how to define change for the monotonic part of the formalism while simultaneously ensuring that the non-monotonic part remains well-behaved, and the other way round. The discussion above is better understood in the light of the following example, re-adapted from Casini et al. (2018):

Example 1. *Assume a propositional language built up from the atoms* a, m, n, v, *and* e, *which stand for, respectively, "(being) an avian red-blood cell", "(being) a mammalian red-blood cell", "(having) a nucleus", "(being) a vertebrate red-blood cell", and "(being) extraterrestrial", and an agent having the following beliefs:* $\mathcal{K} = \{\mathsf{v} \rightarrow \mathsf{n}, \mathsf{a} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \neg\mathsf{n}, \mathsf{a} \rightarrow \neg\mathsf{m}\}$.

1. Clearly the agent should conclude $\neg\mathsf{m}$, *which is unintuitive as mammalian red-blood cells do exist. To remedy this, classical belief change would proceed by removing some of the statements in* $\mathcal{K}$, *such as either* $\mathsf{m} \rightarrow \mathsf{v}$, $\mathsf{v} \rightarrow \mathsf{n}$, *or* $\mathsf{m} \rightarrow \neg\mathsf{n}$. *With non-monotonic belief change the goal is to achieve the same result through the introduction of defeasibility, namely by weakening* $\mathsf{v} \rightarrow \mathsf{n}$ *to* $\mathsf{v} \mid\sim \mathsf{n}$ *to obtain* $\mathcal{K}' = \{\mathsf{v} \mid\sim \mathsf{n}, \mathsf{a} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \neg\mathsf{n}, \mathsf{a} \rightarrow \neg\mathsf{m}\}$.

2. Any well-behaved version of defeasible inheritance should ensure that the agent defeasibly concludes $\mathsf{a} \mid\sim \mathsf{n}$ *from* $\mathcal{K}'$. *But upon learning that* $\mathsf{a} \mid\sim \neg\mathsf{n}$, *the agent should not defeasibly conclude* $\mathsf{a} \mid\sim \mathsf{n}$ *from* $\mathcal{K}'' = \mathcal{K}' \cup \{\mathsf{a} \mid\sim \neg\mathsf{n}\}$, *since this was a defeasible conclusion from* $\mathcal{K}'$, *which is in conflict with the explicitly provided new information.*

3. If the agent then learns that $\mathsf{a} \mid\sim \mathsf{n}$ *is indeed the case to arrive at* $\mathcal{K}''' = \mathcal{K}'' \cup \{\mathsf{a} \mid\sim \mathsf{n}\}$, *it should introduce a conflict with* $\mathsf{a} \mid\sim \neg\mathsf{n}$, *forcing the conclusion* $\mathsf{a} \mid\sim \bot$ *(that there are no avian cells). Such a conflict cannot be handled by the non-monotonic machinery alone. In order to get rid of it, one needs to define principled revision operators.*

4. Finally, assume that, even though the agent considers the existence of mammalian, non-vertebrate red-blood cells to be impossible because of $\mathsf{m} \rightarrow \mathsf{v}$, *it is informed that mammalian, non-vertebrate red-blood cells are presumably extraterrestrial* ($\mathsf{m} \land \neg\mathsf{v} \mid\sim \mathsf{e}$). *It can react to this information in two ways: either it considers the defeasible statement* $\mathsf{m} \land \neg\mathsf{v} \mid\sim \mathsf{e}$ *already satisfied in a trivial way, since it is a weaker form of* $\mathsf{m} \land \neg\mathsf{v} \rightarrow \mathsf{e}$ *which itself is implied by* $\mathsf{m} \rightarrow \mathsf{v}$, *or it follows the reasoning principle, popular for example in formal ontologies, that if we speak of some entity we must* *assume that it exists: the statement* $\mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e}$ *requires us to forego* $\mathsf{m} \rightarrow \mathsf{v}$.

In this paper our focus is on belief change (revision and contraction) for sets of defeasible statements. We provide representation results for basic forms of belief change for a non-monotonic setting, but then show that these are not sufficiently restrictive to capture well-behaved forms of belief change. It turns out that an extension to iterated belief change in the style of Darwiche and Pearl (1997) addresses this deficiency. We generalise the classical work of Darwiche and Pearl (1997) and characterise a family of operators for iterated revision, followed by an analogous characterisation of operators for iterated contraction. In particular, we consider belief change operators aimed at obtaining logical consistency for defeasible sets, and analogous operators aimed at the obtaining coherence, an important notion in the specification of formal ontologies and that we have briefly encountered in Item 4 of Example 1 above.

The remainder of the paper is organised as follows. Section 2 introduces the basics of belief change and KLM-style defeasible reasoning. Section 3 introduces basic defeasible revision and basic coherence defeasible revision via AGM-style sets of postulates and construction methods, linked up through appropriate representation results. Section 4 extends the results of section 3 and defines full defeasible revision and full coherence revision by adding Darwiche-Pearl style postulates, and linking these up to extended construction methods via representation results. Sections 5 and 6 follow patterns similar to those of sections 3 and 4, but for defeasible contraction. Section 7 investigates the extent to which defeasible revision and contraction can be defined in terms of each other using versions of the Levi and Harper Identities. Section 8 discusses related and future work.

2 Preliminaries

$\mathcal{L}$ denotes the (finitely generated) language built up from a finite set of atomic propositions $\mathcal{P}$. We use α, β, $\ldots$ to denote the sentences of $\mathcal{L}$. $\mathcal{U}$ denotes the set of all propositional valuations. Given $\alpha \in \mathcal{L}$, $[\![\alpha]\!] \stackrel{\text{def}}{=} \{u \in \mathcal{U} \mid u \Vdash \alpha\}$ denotes the *models* of α, where $\Vdash$ is the classical satisfaction relation. For $X \subseteq \mathcal{L}$, $[\![X]\!] \stackrel{\text{def}}{=} \bigcap_{\alpha\in X}[\![\alpha]\!]$. *Cn*$(\cdot)$ denotes classical (Tarskian) logical consequence in propositional logic. In our examples we denote valuations as sequences of literals, with a negative literal denoted by a barred atom. E.g., for $\mathcal{P} = \{\mathsf{p}, \mathsf{q}\}$, $\mathsf{p}\overline{\mathsf{q}}$ denotes the valuation satisfying p, but not q.

AGM-Style Belief Change. Belief change is concerned with the problem of modifying an agent's beliefs in a principled way. The AGM approach to belief change (Alchourrón, Gärdenfors, and Makinson 1985) has become the gold standard in the area. It assumes an underlying logic with a propositional language and the above consequence operator *Cn*$(\cdot)$.

A *belief set* $\mathcal{B}$ is assumed to be a set of propositional sentences *closed under* *Cn*$(\cdot)$. The AGM approach is concerned with three basic types of operations on belief sets: *expansion*, *contraction*, and *revision*. Expansion is simply defined as adding a sentence and closing under classical entailment: $\mathcal{B} + \alpha \stackrel{\text{def}}{=} Cn(\mathcal{B} \cup \{\alpha\})$. The expected outcome of contraction of $\mathcal{B}$ with α is a belief set $\mathcal{B} - \alpha$ not entailing α. Dually, the revision of $\mathcal{B}$ with α corresponds to a consistent belief set $\mathcal{B} * \alpha$ from which α follows. Revision can be defined in terms of contraction and expansion via the Levi Identity (Levi 1977): $\mathcal{B} * \alpha = \mathcal{B} - (\neg\alpha) + \alpha$, while contraction can be defined in terms of revision and intersection via the Harper Identity (Gärdenfors 1988): $\mathcal{B} - \alpha = (\mathcal{B} * \neg\alpha) \cap \mathcal{B}$.

In the AGM approach, the behaviour of the change operators described above is characterised in terms of postulates, and several construction methods have been investigated over the past 35 years. A popular method is to define the operators on *epistemic states* (Darwiche and Pearl 1997; Booth and Meyer 2006), each with an associated total preorder and belief set. A feature of epistemic states is that they provide an elegant way to define iterated revision.

KLM-Style Defeasible Reasoning. The framework proposed by Kraus et al. (1990), often referred to as the *KLM approach*, employs defeasible statements of the form $\alpha \mid\!\sim \beta$, read as "typically, if α, then β". The semantics of rational defeasible statements is in terms of ordered structures referred to as *ranked interpretations* (Lehmann and Magidor 1992). Here we adopt the following alternative representation thereof. A ranked interpretation $\mathscr{R}$ is a function from $\mathcal{U}$ to $\mathbb{N}\cup\{\infty\}$ s.t. $\mathscr{R}(u) = 0$ for some $u \in \mathcal{U}$, and satisfying the following *convexity* property: for every $v \in \mathcal{U}$, if $\mathscr{R}(v) = i$, then, for every j s.t. $0 \leq j < i$, there is a $u \in \mathcal{U}$ for which $\mathscr{R}(u) = j$. Valuations with a lower rank are deemed more normal (or typical) than those with a higher rank, while those with an infinite rank are regarded as so atypical as to be impossible. A ranked interpretation $\mathscr{R}$ satisfies (is a ranked model of) $\alpha \mid\!\sim \beta$ (denoted $\mathscr{R} \Vdash \alpha \mid\!\sim \beta$), where $\alpha, \beta \in \mathcal{L}$, if the α-valuations u that are minimal w.r.t. their rank $\mathscr{R}(u)$ also satisfy β. Observe that *all* classical propositional sentences can be expressed as defeasible statements: $\{u \in \mathcal{U} \mid \mathscr{R}(u) \in \mathbb{N}\} \subseteq [\![\alpha]\!]$ iff $\mathscr{R} \Vdash \neg\alpha \mid\!\sim \bot$. KLM-style defeasible reasoning can therefore be viewed as an extension of propositional logic.

3 Basic Defeasible Revision

In the approach to belief revision taken by Darwiche and Pearl (1997) and elaborated on by Booth and Meyer (2006), revision is performed on *epistemic states* where "An epistemic state contains, in addition to a knowledge base, all the information needed for coherent reasoning including, in particular, the strategy for belief revision which the agent wishes to employ at a given time." (Booth and Meyer 2006)

Darwiche and Pearl associate a total preorder $\preceq$ over the set of all valuations with an epistemic state, as well as a belief set, closed under classical consequence, with the valuations that are $\preceq$-minimal being the models of the belief set. The belief set associated with an epistemic state represents the beliefs of an agent, whereas the epistemic state as a whole represents, not only the current beliefs of the agent, but also the additional information required for the agent to perform belief revision. That additional information is captured in the ordering over all valuations, including the counter-models of the belief set.

We follow a similar approach and denote the set of all epistemic states by $\mathcal{E}$. In our case, to define revision for defeasible reasoning we associate with each epistemic state an enriched structure that we refer to as an *epistemic interpretation*. Intuitively the idea is to provide an ordering over the set of counter-models of the classical belief set, similar to the way in which Darwiche and Pearl do. But in order to account for defeasible statements we also provide an ordering over the models of the belief set. Models of the belief set are taken to have a *finite rank* whereas counter-models of the belief set are taken to have an *infinite rank*. This is formalised by assigning to each valuation u a tuple of the form $\langle f, i\rangle$ or $\langle \infty, i\rangle$ where $i \in \mathbb{N}$. The f in $\langle f, i\rangle$ is intended to indicate that u has a *finite rank*, while the ∞ in $\langle \infty, i\rangle$ is intended to indicated that u has an *infinite rank*, where finite ranks are viewed as more typical than infinite ranks. To capture this formally, let $R \stackrel{\text{def}}{=} \{\langle f, i\rangle \mid i \in \mathbb{N}\} \cup \{\langle \infty, i\rangle \mid i \in \mathbb{N}\}$. We define the total ordering $\preceq$ over R as follows: $\langle x_1, y_1\rangle \preceq \langle x_2, y_2\rangle$ if and only if $x_1 = x_2$ and $y_1 \leq y_2$, or $x_1 = f$ and $x_2 = \infty$. Before we can define epistemic interpretations, we need to extend the notion of convexity of ranked interpretations (encountered in the preliminaries) to epistemic interpretations. Let e be a function from $\mathcal{U}$ to R. e is said to be *convex* (w.r.t.$\preceq$) if and only the following holds: i) If $e(u) = \langle f, i\rangle$, then, for all j s.t. $0 \leq j < i$, there is a $u_j \in \mathcal{U}$ s.t. $e(u_j) = \langle f, j\rangle$; and ii) If $e(u) = \langle \infty, i\rangle$, then, for all j s.t. $0 \leq j < i$, there is a $u_j \in \mathcal{U}$ s.t. $e(u_j) = \langle \infty, j\rangle$.

Definition 1. *An **epistemic interpretation** e is a total function from $\mathcal{U}$ to R that is convex.*

Viewed from the classical perspective, epistemic interpretations are refined versions of the total preorders over valuations used in classical AGM revision, in which the models of the belief set on which to perform revision are also ordered (or ranked) to provide a semantics for defeasible statements. Viewed from the perspective of defeasible reasoning, epistemic interpretations are refined versions of *ranked interpretations* in which the valuations with infinite rank are also ordered (or ranked) to perform revision.

We refer to the epistemic interpretation e associated with the epistemic state E as e^E, but we abuse notation slightly by referring to e^E simply as E. We let $\mathcal{U}^f(E) \stackrel{\text{def}}{=} \{u \in \mathcal{U} \mid E(u) = \langle f, i\rangle$ for some $i \in \mathbb{N}\}$ and $\mathcal{U}^\infty(E) \stackrel{\text{def}}{=} \{u \in \mathcal{U} \mid E(u) = \langle \infty, i\rangle$ for some $i \in \mathbb{N}\}$. E is *satisfiable* if $\mathcal{U}^f(E) \neq \emptyset$. We let $\min_E[\![\alpha]\!] \stackrel{\text{def}}{=} \{u \in [\![\alpha]\!] \mid E(u) \preceq E(v)$ for all $v \in [\![\alpha]\!]\}$, and $\min^f_E[\![\alpha]\!] \stackrel{\text{def}}{=} \min_E[\![\alpha]\!] \cap \mathcal{U}^f(E)$. Let $r^E_{\max} V = \max_{\preceq}\{E(u) \mid u \in V\}$ and $r^E_{\min} V = \min_{\preceq}\{E(u) \mid u \in V\}$.

The notion of satisfaction carries over to epistemic interpretations in the obvious way: $\alpha \mathrel{|\!\sim} \beta$ is true in E ($E \Vdash \alpha \mathrel{|\!\sim} \beta$) if and only if $\min^f_E[\![\alpha]\!] \subseteq [\![\beta]\!]$. The set of *epistemic models* of a defeasible statement is defined as follows: $[\![\alpha \mathrel{|\!\sim} \beta]\!] \stackrel{\text{def}}{=} \{E \mid E \Vdash \alpha \mathrel{|\!\sim} \beta\}$. Two defeasible statements $\alpha \mathrel{|\!\sim} \beta$ and $\gamma \mathrel{|\!\sim} \delta$ are said to be *rank-equivalent*, denoted by $\alpha \mathrel{|\!\sim} \beta \equiv_{\mathcal{E}} \gamma \mathrel{|\!\sim} \delta$, if and only if $[\![\alpha \mathrel{|\!\sim} \beta]\!] = [\![\gamma \mathrel{|\!\sim} \delta]\!]$.

The defeasible analogue of a belief set w.r.t. an epistemic state is a *defeasible belief set*.

Definition 2. *A **defeasible belief set** $\mathcal{D}$ is a defeasible set for which there is an epistemic state E s.t. $E \Vdash \alpha \mathrel{|\!\sim} \beta$ if and only if $\alpha \mathrel{|\!\sim} \beta \in \mathcal{D}$. The **defeasible belief set associated with an epistemic state** E is defined as follows: $\mathcal{B}(E) \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim} \beta \mid E \Vdash \alpha \mathrel{|\!\sim} \beta\}$. The **classical belief set associated with an epistemic state** E is defined as follows: $\mathcal{B}^c(E) \stackrel{\text{def}}{=} Cn(\{\neg\alpha \mid E \Vdash \alpha \mathrel{|\!\sim} \bot\})$.*

To make sense of the last part of the definition above, observe that, similar to the case for ranked interpretations mentioned in section 2, a defeasible statement $\alpha \mathrel{|\!\sim} \bot$ can be viewed as equivalent to the classical statement $\neg\alpha$ since $E \Vdash \alpha \mathrel{|\!\sim} \bot$ if and only $\mathcal{U}^f(E) \subseteq [\![\neg\alpha]\!]$. Observe that $\mathcal{B}^c(E)$ is, by definition, a belief set (closed under classical consequence). It is easily shown that $[\![\mathcal{B}^c(E)]\!] = \mathcal{U}^f(E)$.

Our initial focus is on *consistency-based* revision in which revision should not result in an inconsistent epistemic state (one from which $\top \mathrel{|\!\sim} \bot$ follows). As a start, we provide some basic AGM-style postulates for revision. Our initial goal is to focus on the (classical) belief set associated with an epistemic state after revision.

E is said to be *α-enforcing* when $\mathcal{U}^f(E) \subseteq [\![\alpha]\!]$. $\alpha \mathrel{|\!\sim} \beta$ is said to be *compatible* when $\mathcal{U} \not\subseteq [\![\neg(\alpha \wedge \beta)]\!]$, *E-compatible* when $\mathcal{U}^f(E) \not\subseteq [\![\neg(\alpha \wedge \beta)]\!]$, and *weakly E-compatible* when it is E-compatible or E is $\neg\alpha$-enforcing. Note that if $\alpha \mathrel{|\!\sim} \beta$ is incompatible, then it is also E-incompatible.

(B*1) $E * \alpha \mathrel{|\!\sim} \beta$ is an epistemic state (Closure)

(B*2) $\mathcal{B}^c(E * \alpha \mathrel{|\!\sim} \beta) \subseteq \mathcal{B}^c(E) + \alpha \to \beta$ (Inclusion)

(B*3) If $\alpha \mathrel{|\!\sim} \beta$ is weakly E-compatible, then $\mathcal{B}^c(E) = \mathcal{B}^c(E * \alpha \mathrel{|\!\sim} \beta)$ (Strict Vacuity)

(B*4) $\alpha \mathrel{|\!\sim} \beta \in \mathcal{B}(E * \alpha \mathrel{|\!\sim} \beta)$ (Success)

(B*5) If $\alpha \mathrel{|\!\sim} \beta \equiv_{\mathcal{E}} \gamma \mathrel{|\!\sim} \delta$, then $E * (\alpha \mathrel{|\!\sim} \beta) = E * (\gamma \mathrel{|\!\sim} \delta)$ (Extensionality)

(B*6) $\top \mathrel{|\!\sim} \bot \in \mathcal{B}(E * \alpha \mathrel{|\!\sim} \beta)$ if and only if $\alpha \mathrel{|\!\sim} \beta \equiv_{\mathcal{E}} \top \mathrel{|\!\sim} \bot$ (Consistency)

(B*7) $\mathcal{B}^c(E * \alpha \wedge \gamma \mathrel{|\!\sim} \beta) \cap \mathcal{B}^c(E) \subseteq (\mathcal{B}^c(E * \alpha \mathrel{|\!\sim} \beta) + \gamma) \cap \mathcal{B}^c(E)$

(B*8) If $\neg\gamma \notin \mathcal{B}^c(E * \alpha \mathrel{|\!\sim} \beta)$, then $(\mathcal{B}^c(E * \alpha \mathrel{|\!\sim} \beta) + \gamma) \cap \mathcal{B}^c(E) \subseteq \mathcal{B}^c(E * \alpha \wedge \gamma \mathrel{|\!\sim} \beta) \cap \mathcal{B}^c(E)$

(B*9) $\mathcal{B}^c(E) \cap Cn(\alpha \wedge \beta) \subseteq \mathcal{B}^c(E * \alpha \mathrel{|\!\sim} \beta)$ (Containment)

(B*1)-(B*8) are close analogues of the classical AGM postulates for revision and we will not discuss all of them in detail here. Note that (B*2), (B*3), and (B7)-(B*9) all refer to the classical belief set associated with the revised state. To see, for example, why (B*3) applies only to the classical belief set associated with the revised epistemic state, and not to the defeasible belief set as a whole, bear in mind that we are dealing with non-monotonic reasoning. As a very simple example, consider an epistemic state E over p and q, with the elements of $[\![\neg\mathsf{p}]\!]$ ranked at $(f, 0)$ and the elements of $[\![\mathsf{p}]\!]$ at $(f, 1)$ (and nothing ranked at (∞, i) for all i). Note that $\top \mathrel{|\!\sim} \neg\mathsf{p} \in \mathcal{B}(E)$, and that $\top \mathrel{|\!\sim} \mathsf{p}$ is weakly E-compatible, but that any reasonable defeasible revision operator will have that $\top \mathrel{|\!\sim} \neg\mathsf{p} \notin \mathcal{B}(E * \top \mathrel{|\!\sim} \mathsf{p})$.

Containment (B*9) is the only postulate without a classical counterpart. It ensures that the belief set associated with E, when restricted to the consequences of $\alpha \wedge \beta$, are contained in the belief set associated with the epistemic state resulting from a revision with $\alpha \mathrel{|\!\sim} \beta$. Semantically speaking this ensures that only elements of $[\![\alpha \wedge \beta]\!]$ may be added to $\mathcal{U}^f(E)$ to obtain $\mathcal{U}^f(E * \alpha \mathrel{|\!\sim} \beta)$.

E :

$\langle\infty,1\rangle$	$\mathcal{U}\setminus([\![\mathcal{B}^c(E)]\!]\cup\{\overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{vn}}\mathsf{e},\ \overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{v}}\mathsf{ne}\})$
$\langle\infty,0\rangle$	$\overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{vn}}\mathsf{e},\ \overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{v}}\mathsf{ne}$
$\langle f,1\rangle$	$\mathsf{a}\overline{\mathsf{m}}\mathsf{v}\overline{\mathsf{n}}\mathsf{e},\ \mathsf{a}\overline{\mathsf{m}}\mathsf{v}\overline{\mathsf{ne}},\ \overline{\mathsf{a}}\mathsf{mv}\overline{\mathsf{n}}\mathsf{e},\ \overline{\mathsf{a}}\mathsf{mv}\overline{\mathsf{ne}},\ \overline{\mathsf{am}}\mathsf{v}\overline{\mathsf{n}}\mathsf{e},\ \overline{\mathsf{am}}\mathsf{v}\overline{\mathsf{ne}}$
$\langle f,0\rangle$	$\mathsf{a}\overline{\mathsf{m}}\mathsf{vne},\ \mathsf{a}\overline{\mathsf{m}}\mathsf{vn}\overline{\mathsf{e}},\ \overline{\mathsf{am}}\mathsf{vne},\ \overline{\mathsf{am}}\mathsf{vn}\overline{\mathsf{e}},\ \overline{\mathsf{amv}}\mathsf{ne},\ \overline{\mathsf{amv}}\mathsf{n}\overline{\mathsf{e}},\ \overline{\mathsf{amvn}}\mathsf{e},\ \overline{\mathsf{amvne}}$

$E * \mathsf{a} \mid\!\sim \neg\mathsf{n}$:

$\langle\infty,1\rangle$	$\mathcal{U}\setminus[\![\mathcal{B}^c(E)]\!]\cup(\{\overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{vn}}\mathsf{e},\ \overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{v}}\mathsf{ne}\})$
$\langle\infty,0\rangle$	$\overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{vn}}\mathsf{e},\ \overline{\mathsf{a}}\mathsf{m}\overline{\mathsf{v}}\mathsf{ne}$
$\langle f,2\rangle$	$\mathsf{a}\overline{\mathsf{m}}\mathsf{vne},\ \mathsf{a}\overline{\mathsf{m}}\mathsf{vn}\overline{\mathsf{e}}$
$\langle f,1\rangle$	$\mathsf{a}\overline{\mathsf{m}}\mathsf{v}\overline{\mathsf{n}}\mathsf{e},\ \mathsf{a}\overline{\mathsf{m}}\mathsf{v}\overline{\mathsf{ne}},\ \overline{\mathsf{a}}\mathsf{mv}\overline{\mathsf{n}}\mathsf{e},\ \overline{\mathsf{a}}\mathsf{mv}\overline{\mathsf{ne}},\ \overline{\mathsf{am}}\mathsf{v}\overline{\mathsf{n}}\mathsf{e},\ \overline{\mathsf{am}}\mathsf{v}\overline{\mathsf{ne}}$
$\langle f,0\rangle$	$\overline{\mathsf{am}}\mathsf{vne},\ \overline{\mathsf{am}}\mathsf{vn}\overline{\mathsf{e}},\ \overline{\mathsf{amv}}\mathsf{ne},\ \overline{\mathsf{amv}}\mathsf{n}\overline{\mathsf{e}},\ \overline{\mathsf{amvn}}\mathsf{e},\ \overline{\mathsf{amvne}}$

Figure 1: Example of a basic defeasible revision operator.

Definition 3. $*$ *is an* ***admissible defeasible revision operator*** *if $E * \alpha \mid\!\sim \beta \Vdash \alpha \mid\!\sim \beta$, and $E * \alpha \mid\!\sim \beta = E * \gamma \mid\!\sim \delta$ if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \gamma \mid\!\sim \delta$. An admissible defeasible revision operator $*$ is a* ***basic defeasible revision operator*** *if it is defined as follows: $E * \alpha \mid\!\sim \beta \stackrel{\text{def}}{=} E'$, where E' is an epistemic state such that $\mathcal{U}^f(E') = \mathcal{U}^f(E)$ if $\alpha \mid\!\sim \beta$ is weakly E-compatible, $\mathcal{U}^f(E') = \emptyset$ if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \top \mid\!\sim \bot$, and $\mathcal{U}^f(E') \stackrel{\text{def}}{=} \mathcal{U}^f(E) \cup \min_E[\![\alpha \wedge \beta]\!]$ otherwise.*

Theorem 1. $*$ *satisfies (B*1) to (B*9) if and only if it is a basic defeasible revision operator.*

Example 2. *Consider Item 2 of Example 1 and suppose that the beliefs of an agent are described as the defeasible belief set associated with the epistemic state E in Fig. 1. Note that E is satisfiable, that $\{\mathsf{v} \mid\!\sim \mathsf{n}, \mathsf{a} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \neg\mathsf{n}, \mathsf{a} \rightarrow \neg\mathsf{m}\} \subseteq \mathcal{B}(E)$, that $\mathcal{B}^c(E) = Cn(\{\mathsf{a} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \neg\mathsf{n}, \mathsf{a} \rightarrow \neg\mathsf{m}\})$, and that $\mathsf{a} \mid\!\sim \mathsf{n} \in \mathcal{B}(E)$. In Fig. 1 we show how one of the possible basic defeasible revision operators $*$ behaves when revising with $\mathsf{a} \mid\!\sim \neg\mathsf{n}$. First off, note that $\mathsf{a} \mid\!\sim \neg\mathsf{n}$ is E-compatible, from which it follows (via (B*3)) that $\mathcal{U}^f(E * \mathsf{a} \mid\!\sim \neg\mathsf{n}) = \mathcal{U}^f(E)$. And then observe that E is modified as follows to obtain $E * \mathsf{a} \mid\!\sim \neg\mathsf{n}$. Except for two valuations, $\mathsf{a}\overline{\mathsf{m}}\mathsf{vne}$ and $\mathsf{a}\overline{\mathsf{m}}\mathsf{vn}\overline{\mathsf{e}}$, all valuations retain the same rank. These two valuations are moved 'up' by changing their rank of $\langle f,0\rangle$ to a (finite) rank of $\langle f,2\rangle$. This has the effect of ensuring that $\mathsf{a} \mid\!\sim \neg\mathsf{n} \in \mathcal{B}(E * \mathsf{a} \mid\!\sim \neg\mathsf{n})$ since $\min_{E*\mathsf{a}\mid\sim\neg\mathsf{n}}[\![\mathsf{a}]\!] \subseteq [\![\neg\mathsf{n}]\!]$. It is important to note that this is one of many ways in which basic defeasible revision operators could have dealt with a revision with $\mathsf{a} \mid\!\sim \neg\mathsf{n}$.*

Observe that basic defeasible revision (via Strict Vacuity) will keep the belief set associated with the epistemic state E unchanged through a revision by $\alpha \mid\!\sim \beta$ whenever $\alpha \mid\!\sim \bot$ is true in E, since $\alpha \mid\!\sim \beta$ is then also true in E. This may seem counter-intuitive, but it is essentially because basic defeasible revision is consistency-based. If we aim to obtain coherence we need to consider *coherence-based* revision where a revision by $\alpha \mid\!\sim \beta$ should not result in an epistemic state from which $\alpha \mid\!\sim \bot$ follows. This can be done by replacing (B*3) with (B*3') and adding (B*10) below.

(B*3') If $\alpha \mid\!\sim \beta$ is E-compatible, then $\mathcal{B}^c(E) = \mathcal{B}^c(E * \alpha \mid\!\sim \beta)$ (Weak Vacuity)

(B*10) $\alpha \mid\!\sim \bot \in \mathcal{B}(E * \alpha \mid\!\sim \beta)$ if and only if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \alpha \mid\!\sim \bot$ (Coherence)

Weak Vacuity requires E to remain unchanged after a revision by $\alpha \mid\!\sim \beta$ only when the latter is E-compatible. Coherence ensures that the only reason for an $\alpha \mid\!\sim \beta$-revision to contain $\alpha \mid\!\sim \bot$ is if the former is, in fact, rank-equivalent to $\alpha \mid\!\sim \bot$.

Definition 4. *An admissible revision operator $*$ is a* ***basic coherence defeasible revision operator*** *if it is defined as follows: $E * \alpha \mid\!\sim \beta \stackrel{\text{def}}{=} E'$ where E' is an epistemic state such that $\mathcal{U}^f(E') = \mathcal{U}^f(E)$ if $\alpha \mid\!\sim \beta$ is E-compatible, $\mathcal{U}^f(E') \stackrel{\text{def}}{=} \emptyset$ if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \top \mid\!\sim \bot$, and $\mathcal{U}^f(E') \stackrel{\text{def}}{=} \mathcal{U}^f(E) \cup \min_E[\![\alpha \wedge \beta]\!]$ otherwise.*

Theorem 2. $*$ *satisfies (B*1) to (B*10), with (B*3) replaced by (B*3'), if and only if it is a basic coherence defeasible revision operator.*

Example 3. *Consider Item 4 of Example 1. E' in Fig. 2 corresponds to the epistemic state $E * \mathsf{a} \mid\!\sim \neg\mathsf{n}$ from Example 2. We want to revise with $\mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e}$. Note that $\mathcal{B}^c(E') = Cn(\{\mathsf{a} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \mathsf{v}, \mathsf{m} \rightarrow \neg\mathsf{n}\})$, that $\mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e} \in \mathcal{B}(E')$ since $\neg(\mathsf{m} \wedge \neg\mathsf{v}) \in \mathcal{B}^c(E')$, and $\mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e}$ is weakly E'-compatible, but not E'-compatible. We can implement two kinds of revision policies. If we are interested in obtaining consistency, we would implement a* basic defeasible revision operator $*$. *We already have that $\mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e} \in \mathcal{B}(E')$ and, because of (B*3), we are not allowed to change the belief set associated with the epistemic state. Hence a viable revision policy is simply doing nothing: $E' * \mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e} = E'$.*

*If we are interested in obtaining coherence then, with the introduction of $\mathsf{m} \wedge \neg\mathsf{v} \mid\!\sim \mathsf{e}$, we are forced by (B*10) and allowed by (B*3') to change the belief set associated with the epistemic state, allowing for the existence of non-vertebrate mammalian red-blood cells ($\mathsf{m} \wedge \neg\mathsf{v}$). We can do it, for example, by moving elements of $[\![\mathsf{m} \wedge \neg\mathsf{v} \wedge \mathsf{e}]\!]$ 'down' from the (infinite) rank of $\langle\infty,0\rangle$ to the (finite) rank of $\langle f,3\rangle$, as in Fig. 2. In this way we make extraterrestrial non-vertebrate mammalians conceivable, but we consider them as a very implausible option, since elements of $[\![\mathsf{m} \wedge \neg\mathsf{v} \wedge \mathsf{e}]\!]$ occur only in the highest finite rank. As required, the defeasible statement $\mathsf{m} \mid\!\sim \mathsf{v}$ now becomes purely defeasible ($\mathsf{m} \mid\!\sim \mathsf{v}$ holds in the new epistemic state, but $\mathsf{m} \rightarrow \mathsf{v}$ doesn't).*

From Theorems 1 and 2 we have that both forms of revision defined above generalise classical basic AGM revision.

Corollary 1. *1. Let $*$ be a basic defeasible revision operator, E an epistemic state, and $\alpha \in \mathcal{L}$. Then there is a basic AGM revision operator $*'$ s.t. $\mathcal{B}^c(E) *' \alpha = \mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot)$. Conversely, suppose that $*'$ is a basic AGM revision operator, let B be a belief set, and let E be such that $\mathcal{B}^c(E) = K$. Then there is a basic defeasible revision operator $*$ s.t. $\mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot) = B *' \alpha$.*

2. Let $*$ be a basic coherence defeasible revision operator, E an epistemic state, and $\alpha \in \mathcal{L}$. Then there is a basic AGM

E' :

$\langle\infty, 1\rangle$	$\mathcal{U} \setminus [\![\mathcal{B}^c(E)]\!] \cup (\{\overline{a}m\overline{v}\overline{n}e, \overline{a}m\overline{v}ne\})$
$\langle\infty, 0\rangle$	$\overline{a}m\overline{v}\overline{n}e, \overline{a}m\overline{v}ne$
$\langle f, 2\rangle$	$a\overline{m}vne, a\overline{m}vn\overline{e}$
$\langle f, 1\rangle$	$a\overline{m}v\overline{n}e, a\overline{m}v\overline{n}\overline{e}, \overline{a}mv\overline{n}e, \overline{a}mv\overline{n}\overline{e}, \overline{a}\overline{m}v\overline{n}e, \overline{a}\overline{m}v\overline{n}\overline{e}$
$\langle f, 0\rangle$	$\overline{a}\overline{m}vne, \overline{a}\overline{m}vn\overline{e}, \overline{a}\overline{m}\overline{v}ne, \overline{a}\overline{m}\overline{v}n\overline{e}, \overline{a}\overline{m}\overline{v}\overline{n}e, \overline{a}\overline{m}\overline{v}\overline{n}\overline{e}$

$E' * \mathsf{m} \wedge \neg\mathsf{v} \mathrel{|\!\!\sim} \mathsf{e}$:

$\langle\infty, 1\rangle$	$\mathcal{U} \setminus [\![\mathcal{B}^c(E)]\!] \cup (\{\overline{a}m\overline{v}\overline{n}e, \overline{a}m\overline{v}ne\})$
$\langle f, 3\rangle$	$\overline{a}m\overline{v}\overline{n}e, \overline{a}m\overline{v}ne$
$\langle f, 2\rangle$	$a\overline{m}vne, a\overline{m}vn\overline{e}$
$\langle f, 1\rangle$	$a\overline{m}v\overline{n}e, a\overline{m}v\overline{n}\overline{e}, \overline{a}mv\overline{n}e, \overline{a}mv\overline{n}\overline{e}, \overline{a}\overline{m}v\overline{n}e, \overline{a}\overline{m}v\overline{n}\overline{e}$
$\langle f, 0\rangle$	$\overline{a}\overline{m}vne, \overline{a}\overline{m}vn\overline{e}, \overline{a}\overline{m}\overline{v}ne, \overline{a}\overline{m}\overline{v}n\overline{e}, \overline{a}\overline{m}\overline{v}\overline{n}e, \overline{a}\overline{m}\overline{v}\overline{n}\overline{e}$

Figure 2: Basic coherence defeasible revision.

revision operator $*'$ s.t. $\mathcal{B}^c(E) *' \alpha = \mathcal{B}^c(E * \neg\alpha \mathrel{|\!\!\sim} \bot)$. Conversely, suppose that $*'$ is a basic AGM revision operator, let B be a belief set, and let E be such that $\mathcal{B}^c(E) = K$. Then there is a basic coherence defeasible revision operator $*$ s.t. $\mathcal{B}^c(E * \neg\alpha \mathrel{|\!\!\sim} \bot) = B *' \alpha$.

It follows that basic defeasible revision and basic coherence defeasible revision coincide when revising with statements of the form $\alpha \mathrel{|\!\!\sim} \bot$.

Basic (coherence) defeasible revision uses the ranking of elements in $\mathcal{U}^\infty(E)$ to perform revision in a systematic way, but largely ignores the ranking of elements in $\mathcal{U}^f(E)$. The extended forms of revision we present in the next section takes the ranking of elements in $\mathcal{U}^f(E)$ into account as well.

The following two properties require the epistemic state to remain unchanged following a revision with $\alpha \mathrel{|\!\!\sim} \beta$ when $\alpha \mathrel{|\!\!\sim} \beta$ already follows from it (and in the case of obtaining coherence, $\alpha \mathrel{|\!\!\sim} \bot$ doesn't).

If $\alpha \mathrel{|\!\!\sim} \beta \in \mathcal{B}(E)$, then $E * \alpha \mathrel{|\!\!\sim} \beta = E$ (Inertia)

If $\alpha \mathrel{|\!\!\sim} \beta \in \mathcal{B}(E)$ and $\alpha \mathrel{|\!\!\sim} \bot \notin \mathcal{B}(E)$, then
$E * \alpha \mathrel{|\!\!\sim} \beta = E$ (Coherence Inertia)

The following example shows that Inertia and Coherence Inertia are too strong.

Example 4. *For $\mathcal{P} = \{\mathsf{p}, \mathsf{q}\}$ let E be such that $E(\mathsf{pq}) = \langle \mathsf{f}, 0\rangle$, $E(\overline{\mathsf{p}}\overline{\mathsf{q}}) = E(\overline{\mathsf{p}}\mathsf{q}) = \langle\infty, 0\rangle$, and $E(\mathsf{p}\overline{\mathsf{q}}) = \langle\infty, 1\rangle$. Note that $\mathsf{p} \vee \mathsf{q} \mathrel{|\!\!\sim} \mathsf{p} \in \mathcal{B}(E)$. It seems reasonable to require that under some circumstances, strengthening $\mathsf{p} \vee \mathsf{q}$ by adding $\neg(\mathsf{p} \wedge \mathsf{q})$ should still yield p as the consequent of a defeasible statement. And indeed, basic (coherence) defeasible revision allows for that. Yet, Inertia and Coherence Inertia will have as a result that $(\mathsf{p} \vee \mathsf{q}) \wedge \neg(\mathsf{p} \wedge \mathsf{q}) \not\mathrel{|\!\!\sim} \mathsf{p} \in \mathcal{B}(E * \mathsf{p} \vee \mathsf{q} \mathrel{|\!\!\sim} \mathsf{p})$.*

The following weaker versions of Inertia seem more reasonable. They require the defeasible belief set associated with an epistemic state to remain unchanged when $\alpha \mathrel{|\!\!\sim} \beta$ already follows from it (and in the case of obtaining coherence, $\alpha \mathrel{|\!\!\sim} \bot$ doesn't).

(B*11) If $\alpha \mathrel{|\!\!\sim} \beta \in \mathcal{B}(E)$, then
$\mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \beta) = \mathcal{B}(E)$ (Weak Inertia)

(B*12) If $\alpha \mathrel{|\!\!\sim} \beta \in \mathcal{B}(E)$ and $\alpha \mathrel{|\!\!\sim} \bot \notin \mathcal{B}(E)$, then
$\mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \beta) = \mathcal{B}(E)$ (Coherence Weak Inertia)

The following example shows that basic defeasible revision does not satisfy Weak Inertia, and basic coherence defeasible revision does not satisfy Coherence Weak Inertia.

Example 5. *For $\mathcal{P} = \{\mathsf{p}, \mathsf{q}\}$ let E be such that $E(\mathsf{pq}) = \langle \mathsf{f}, 0\rangle$ and $E(\mathsf{p}\overline{\mathsf{q}}) = E(\overline{\mathsf{p}}\overline{\mathsf{q}}) = E(\overline{\mathsf{p}}\mathsf{q}) = \langle \mathsf{f}, 1\rangle$. Basic defeasible revision and basic coherence defeasible revision allow for a revision with $\mathsf{p} \mathrel{|\!\!\sim} \mathsf{q}$ to result in E' where $E(\mathsf{pq}) = \langle \mathsf{f}, 0\rangle$ and $E(\mathsf{p}\overline{\mathsf{q}}) = \langle \mathsf{f}, 1\rangle$, and $E(\overline{\mathsf{p}}\overline{\mathsf{q}}) = E(\overline{\mathsf{p}}\mathsf{q}) = \langle \mathsf{f}, 2\rangle$. Note that $\mathsf{p} \mathrel{|\!\!\sim} \mathsf{q} \in \mathcal{B}(E)$ and $\mathsf{p} \mathrel{|\!\!\sim} \bot \notin \mathcal{B}(E)$, that $\neg\mathsf{q} \mathrel{|\!\!\sim} \mathsf{p} \notin \mathcal{B}(E)$, but that $\neg\mathsf{q} \mathrel{|\!\!\sim} \mathsf{p} \in \mathcal{B}(E')$.*

We conclude this section with a semantic description of these versions of Inertia.

(SB*11) If $\alpha \mathrel{|\!\!\sim} \beta \in \mathcal{B}(E)$, then $\mathcal{U}^f(E) = \mathcal{U}^f(E * \alpha \mathrel{|\!\!\sim} \beta)$ and $E(u) = (E * \alpha \mathrel{|\!\!\sim} \beta)(u)$
for all $u \in \mathcal{U}^f(E)$ (Semantic Weak Inertia)

(SB*12) If $\alpha \mathrel{|\!\!\sim} \beta \in \mathcal{B}(E)$ and $\alpha \mathrel{|\!\!\sim} \bot \notin \mathcal{B}(E)$, then $\mathcal{U}^f(E) = \mathcal{U}^f(E * \alpha \mathrel{|\!\!\sim} \beta)$ and $E(u) = (E * \alpha \mathrel{|\!\!\sim} \beta)(u)$ for all $u \in \mathcal{U}^f(E)$ (Semantic Coherence Weak Inertia)

Proposition 1. *A basic defeasible revision operator $*$ (coherence defeasible revision operator) satisfies (B*11) (satisfies (B*12)) iff it satisfies (SB*11) (satisfies (SB*12)).*

Example 6. *Looking at Example 2, the basic defeasible revision operator described there satisfies Inertia (and therefore Weak Inertia) for a revision with $\mathsf{m} \wedge \neg\mathsf{v} \mathrel{|\!\!\sim} \mathsf{e}$ since E' remains unchanged following such a revision. The coherence defeasible revision operator described in the example vacuously satisfies Coherence Inertia (and therefore Coherence Weak Inertia) for a revision with $\mathsf{m} \wedge \neg\mathsf{v} \mathrel{|\!\!\sim} \mathsf{e}$. Observe that (B*10) forces us to modify E' by introducing elements of $[\![\mathsf{m} \wedge \neg\mathsf{v} \wedge \mathsf{e}]\!]$ into $\mathcal{U}^f(E * \mathsf{m} \wedge \neg\mathsf{v} \mathrel{|\!\!\sim} \mathsf{e})$ in order to obtain coherence w.r.t. $\mathsf{m} \wedge \neg\mathsf{v}$.*

4 Full Defeasible Revision

The focus in the previous section was on specifying the classical belief set associated with a revised epistemic state. In this section we focus on determining the full defeasible belief set associated with a revised epistemic state. It turns out that this can be done by specifying a number of postulates that constrain the way *iterated revision* is performed.

(BDP*1) If $\alpha \models \gamma$, then for all δ,
$\alpha \mathrel{|\!\!\sim} \delta \in \mathcal{B}(E{*}\alpha \mathrel{|\!\!\sim} \beta)$ iff $\alpha \mathrel{|\!\!\sim} \delta \in \mathcal{B}(E{*}\gamma \mathrel{|\!\!\sim} \beta{*}\alpha \mathrel{|\!\!\sim} \beta)$

(BDP*2) If $\alpha \models \gamma$, then for all δ, $\alpha \mathrel{|\!\!\sim} \delta \in \mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \neg\beta)$ iff $\alpha \mathrel{|\!\!\sim} \delta \in \mathcal{B}(E * \gamma \mathrel{|\!\!\sim} \beta * \alpha \mathrel{|\!\!\sim} \neg\beta)$

(BDP*3) If $\alpha \mathrel{|\!\!\sim} \gamma \in \mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \beta)$, then $\alpha \mathrel{|\!\!\sim} \gamma \in \mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \gamma * \alpha \mathrel{|\!\!\sim} \beta)$

(BDP*4) If $\alpha \mathrel{|\!\!\sim} \neg\gamma \notin \mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \beta)$, then $\alpha \mathrel{|\!\!\sim} \neg\gamma \notin \mathcal{B}(E * \alpha \mathrel{|\!\!\sim} \gamma * \alpha \mathrel{|\!\!\sim} \beta)$

(BDP*5) If $\alpha \models \neg\gamma$, then for all δ,
$\alpha \mathrel{|\!\!\sim} \delta \in \mathcal{B}(E{*}\alpha \mathrel{|\!\!\sim} \beta)$ iff $\alpha \mathrel{|\!\!\sim} \delta \in \mathcal{B}(E{*}\gamma \mathrel{|\!\!\sim} \beta{*}\alpha \mathrel{|\!\!\sim} \beta)$

The first four postulates above are analogues of the postulates for iterated revision proposed by Darwiche and Pearl (1997). Motivations for the latter were discussed in detail by Darwiche and Pearl (1997), and also by Booth and Meyer (2006). These motivations carry over to our versions of the postulates, and we don't discuss them here. The fifth postulate does not have an analogue in the classical case. It states that revision with a defeasible statement with α as the antecedent overrides a prior revision with a defeasible statement with an antecedent that is disjoint with α when it comes to defeasible consequences with α as antecedent. Its motivation is similar to the motivation for (BDP*2). Kern-Isberner (2018) proposed semantic versions of the five postulates above in the context identifying a a special case of a more general operator for change by sets of conditionals. She refers to them as (CP1a-CP1c), (CP2*), and (CP3*). We restate those postulates below.

(SDP*1) If $u, v \in [\![\alpha \wedge \beta]\!]$, then $E(u) \preceq E(v)$ iff $(E * \alpha \mid\!\sim \beta)(u) \preceq (E * \alpha \mid\!\sim \beta)(v)$

(SDP*2) If $u, v \in [\![\alpha \wedge \neg\beta]\!]$, then $E(u) \preceq E(v)$ iff $(E * \alpha \mid\!\sim \beta)(u) \preceq (E * \alpha \mid\!\sim \beta)(v)$

(SDP*3) If $u \in [\![\alpha \wedge \beta]\!]$ and $v \in [\![\alpha \wedge \neg\beta]\!]$, then $E(u) \prec E(v)$ implies
$(E * \alpha \mid\!\sim \beta)(u) \prec (E * \alpha \mid\!\sim \beta)(v)$

(SDP*4) If $u \in [\![\alpha \wedge \beta]\!]$ and $v \in [\![\alpha \wedge \neg\beta]\!]$, then $E(u) \preceq E(v)$ implies
$(E * \alpha \mid\!\sim \beta)(u) \preceq (E * \alpha \mid\!\sim \beta)(v)$

(SDP*5) If $u, v \in [\![\neg\alpha]\!]$, then $E(u) \preceq E(v)$ iff $(E * \alpha \mid\!\sim \beta)(u) \preceq (E * \alpha \mid\!\sim \beta)(v)$

(SDP*1) states that the elements of $[\![\alpha \wedge \beta]\!]$ keep their relative ranking when revising by $\alpha \mid\!\sim \beta$. (SDP*2) states that the elements of $[\![\alpha \wedge \neg\beta]\!]$ keep their relative ranking when revising by $\alpha \mid\!\sim \beta$. (SDP*3) and (SDP*4) together state that when revising by $\alpha \mid\!\sim \beta$, the elements of $[\![\alpha \wedge \beta]\!]$ can be slid "down" the ranking w.r.t. the elements of $[\![\alpha \wedge \neg\beta]\!]$. (SDP*5) states that when revising by $\alpha \mid\!\sim \beta$, the elements of $[\![\neg\alpha]\!]$ keep their relative ranking, but there is no constraint on how the elements of $[\![\neg\alpha]\!]$ are ranked w.r.t. the elements of $[\![\alpha]\!]$. The motivation for (SDP*5) is that, since defeasible statements are conditional on the antecedent, a revision with a defeasible statement having α as the antecedent should be independent of the ranking of elements of $[\![\neg\alpha]\!]$.

Lemma 1. *Given postulates (B*1)-(B*9), for $i \in \{1, 2, 3, 4, 5\}$, (BDP*i) and (SDP*i) are equivalent.*

This strengthens basic (coherence) defeasible revision

Definition 5. *A **full (coherence) defeasible revision operator** is a basic (coherence) defeasible revision operator satisfying (SDP*1)-(SDP*5).*

Theorem 3. *$*$ satisfies (B*1)-(B*9) and (BDP*1)-(BDP*5) if and only if it is a full defeasible revision operator. $*$ satisfies (B*1)-(B*10), with (B*3) replaced by (B*3'), and (BDP*1)-(BDP*5) if and only if it is a full coherence defeasible revision operator.*

As mentioned in the previous section, unlike basic (coherence) defeasible revision, full (coherence) defeasible revision takes the ranking of the elements of $\mathcal{U}^f(E)$ into account. It does so by restricting the relative ranking of elements $\mathcal{U}$ as specified by (SDP*1)-(SDP*5) and, as a consequence, it restricts the relative ranking of elements of $\mathcal{U}(E)$ as well. Consider the following simple example.

Example 7. *For $\mathcal{P} = \{\mathsf{p}, \mathsf{q}\,\mathsf{r}\}$ let E be such that $E(\mathsf{p}\overline{\mathsf{q}}\mathsf{r}) = \langle \mathsf{f}, 0\rangle$, $E(\mathsf{pqr}) = \langle \mathsf{f}, 1\rangle$, $E(\mathsf{p}\overline{\mathsf{qr}}) = \langle \mathsf{f}, 2\rangle$, and $u = \langle(\infty, 0\rangle$ for all $u \in \mathcal{U} \setminus \{\mathsf{p}\overline{\mathsf{q}}\mathsf{r}, \mathsf{pqr}, \mathsf{p}\overline{\mathsf{qr}}\}$. Note that $\mathsf{p} \mid\!\sim \neg\mathsf{q} \in \mathcal{B}(E)$ and $\neg\mathsf{q} \mid\!\sim \mathsf{r} \in \mathcal{B}(E)$, but $\mathsf{p} \mid\!\sim \mathsf{q} \notin \mathcal{B}(E)$. It seems reasonable to require that $\neg\mathsf{q} \mid\!\sim \mathsf{r} \in \mathcal{B}(E * \mathsf{p} \mid\!\sim \mathsf{q})$. And while full (coherence) defeasible revision requires us to draw this conclusion, basic (coherence) defeasible revision does not guarantee it. The difference in this case is that full (coherence) defeasible revision (via (SDP*2) requires of us to retain the relative ranking of $\mathsf{p}\overline{\mathsf{q}}\mathsf{r}$ and $\mathsf{p}\overline{\mathsf{qr}}$ from E to $E * \mathsf{p} \mid\!\sim \mathsf{q}$ since both are elements of $[\![\mathsf{p} \wedge \neg\mathsf{q}]\!]$.*

We conclude this section by presenting a concrete (coherence) defeasible revision operator. It is styled on Boutilier's Natural Revision operator (1996) for classical iterated revision. When performing a revision with $\alpha \mid\!\sim \beta$, it changes the relative ranking of only the minimal elements of $[\![\alpha \wedge \beta]\!]$, and then it moves them down as little as possible to ensure that $\alpha \mid\!\sim \beta$ holds in the revised epistemic state.

Definition 6. *(Natural Defeasible Revisions)*

Natural Defeasible Revision: $E * \alpha \mid\!\sim \beta \stackrel{\text{def}}{=} E'$ *where*

1. *if $E \Vdash \alpha \mid\!\sim \beta$, then $E' = E$; else*
2. *if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \top \mid\!\sim \bot$, then $E'(u) = \langle \infty, i\rangle$ where $E(u) = \langle f, i\rangle$ for all $u \in \mathcal{U}^f(E)$, and $E'(u) = \langle \infty, j + i\rangle$ where $r^E_{\max}\mathcal{U}^f(E) = \langle f, j\rangle$; else*
3. *$E'(u) = E(u)$ for all $(u \in \mathcal{U}^\infty(E) \setminus \min_E[\![\alpha \wedge \beta]\!]) \cup \{v \mid v \prec r^E_{\min}[\![\alpha \wedge \neg\beta]\!]$, $E'(u) = r^E_{\min}[\![\alpha \wedge \neg\beta]\!]$ for all $u \in \min_E[\![\alpha \wedge \beta]\!]$, and $E'(u) = E(u) + 1$ for all $u \in \{v \in \mathcal{U}^f(E) \mid r^E_{\min}[\![\alpha \wedge \neg\beta]\!] \preceq v\}$.*

Natural Coherence Defeasible Revision: $E * \alpha \mid\!\sim \beta \stackrel{\text{def}}{=} E'$ *where*

1. *if $E \Vdash \alpha \mid\!\sim \beta$ and $E \nVdash \alpha \mid\!\sim \bot$, then $E' = E$; else*
2. *if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \top \mid\!\sim \bot$, then $E'(u) = \langle \infty, i\rangle$ where $E(u) = \langle f, i\rangle$ for all $u \in \mathcal{U}^f(E)$, and $E'(u) = \langle \infty, j + i\rangle$ where $r^E_{\max}\mathcal{U}^f(E) = \langle f, j\rangle$; else*
3. *$E'(u) = E(u)$ for all $(u \in \mathcal{U}^\infty(E) \setminus \min_E[\![\alpha \wedge \beta]\!]) \cup \{v \mid v \prec r^E_{\min}[\![\alpha \wedge \neg\beta]\!]$, $E'(u) = r^E_{\min}[\![\alpha \wedge \neg\beta]\!]$ for all $u \in \min_E[\![\alpha \wedge \beta]\!]$, and $E'(u) = E(u) + 1$ for all $u \in \{v \in \mathcal{U}^f(E) \mid r^E_{\min}[\![\alpha \wedge \neg\beta]\!] \preceq v\}$.*

Proposition 2. *Natural (Coherence) Defeasible Revision is a full (coherence) defeasible revision operator satisfying (Coherence) Inertia (and (Coherence) Weak Inertia).*

To get a sense of what Natural (Coherence) Defeasible Revision is about, consider the following simple example.

Example 8. *For $\mathcal{P} = \{\mathsf{p}, \mathsf{q}\}$ let E be such that $E(\overline{\mathsf{pq}}) = \langle f, 0\rangle$, $E(\overline{\mathsf{p}}\mathsf{q}) = \langle \mathsf{f}, 1\rangle$, $E(\mathsf{pq}) = \langle \mathsf{f}, 2\rangle$, and $E(\mathsf{p}\overline{\mathsf{q}}) = \langle \mathsf{f}, 3\rangle$. A (coherence) defeasible revision with p will result in E' where $E'(\mathsf{pq}) = \langle \mathsf{f}, 0\rangle$, $E'(\overline{\mathsf{pq}}) = \langle f, 1\rangle$, $E'(\overline{\mathsf{p}}\mathsf{q}) = \langle \mathsf{f}, 2\rangle$, and $E'(\mathsf{p}\overline{\mathsf{q}}) = \langle \mathsf{f}, 3\rangle$. This ensures, for example, that $\neg\mathsf{q} \mid\!\sim \neg\mathsf{p}$, which is in $\mathcal{B}(E)$, is retained in $\mathcal{B}(E * \mathsf{p})$. But note that full (coherence) defeasible revision with p also allows for E'' where $E'(\mathsf{pq}) = \langle \mathsf{f}, 0\rangle$, $E'(\mathsf{p}\overline{\mathsf{q}}) = \langle \mathsf{f}, 1\rangle$, $E'(\overline{\mathsf{pq}}) = \langle f, 2\rangle$, and $E'(\overline{\mathsf{p}}\mathsf{q}) = \langle \mathsf{f}, 3\rangle$, and that $\neg\mathsf{q} \mid\!\sim \neg\mathsf{p} \notin \mathcal{B}(E'')$.*

5 Basic Defeasible Contraction

We now turn to defeasible *contraction*. Observe that there is no distinction between obtaining consistency and obtaining coherence in the case of contraction operators. Or rather, obtaining coherence is folded into contraction in the following sense. If a contraction by $\alpha \mid\!\sim \beta$ is performed, the Success postulate below guarantees that $\alpha \mid\!\sim \beta$ will not be in the defeasible belief set associated with the contracted epistemic state (unless it is rank-equivalent to $\top \mid\!\sim \top$).

(B-1) $E - (\alpha \mid\!\sim \beta))$ is an epistemic state (Closure)

(B-2) $\mathcal{B}^c(E - \alpha \mid\!\sim \beta) \subseteq \mathcal{B}^c(E)$ (Inclusion)

(B-3) If $\alpha \rightarrow \beta \notin \mathcal{B}(E)$, then $\mathcal{B}^c(E) \subseteq \mathcal{B}^c(E - \alpha \mid\!\sim \beta)$ (Vacuity)

(B-4) If $\alpha \mid\!\sim \beta \not\equiv_{\mathcal{E}} \top \mid\!\sim \top$, then $\alpha \mid\!\sim \beta \notin \mathcal{B}(E - \alpha \mid\!\sim \beta)$ (Success)

(B-5) If $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \gamma \mid\!\sim \delta$, then $E - (\alpha \mid\!\sim \beta) = E - (\gamma \mid\!\sim \delta)$ (Extensionality)

(B-6) $\mathcal{B}^c(E) \cap Cn(\alpha \wedge \neg\beta) \subseteq \mathcal{B}^c(E - \alpha \mid\!\sim \beta)$ (Containment)

(B-7) $(\mathcal{B}^c(E) - \alpha \mid\!\sim \beta) \cap (\mathcal{B}^c(E) - \gamma \mid\!\sim \beta) \subseteq \mathcal{B}^c(E) - \alpha \wedge \gamma \mid\!\sim \beta$ (Conjunction Overlap)

(B-8) If $\gamma \mid\!\sim \beta \notin \mathcal{B}^c(E) - \alpha \wedge \gamma \mid\!\sim \beta$, then $\mathcal{B}^c(E) - \alpha \wedge \gamma \mid\!\sim \beta \subseteq (\mathcal{B}^c(E) - \alpha \mid\!\sim \beta) \cap (\mathcal{B}^c(E) - \gamma \mid\!\sim \beta)$ (Conjunction Inclusion)

Except for (B-6), these postulates are all close analogues of the classical AGM postulates for contraction, ans we will not discuss them here in detail. Note that (B-2), (B-3), and (B-6)-(B-8) refer to the classical belief set associated with an epistemic state, and not to the full defeasible belief set. Containment ensures that when contracting by $\alpha \mid\!\sim \beta$, only elements of $[\![\alpha \wedge \neg\beta]\!]$ can change from an infinite rank to a finite rank. It is a replacement for the following postulate which is too strong in this context.

If $\alpha \mid\!\sim \beta \in \mathcal{B}(E)$, then $\mathcal{B}^c(E) = \mathcal{B}^c(E - \alpha \mid\!\sim \beta) + \alpha \rightarrow \beta$ (Semi-Classic Recovery)

Semi-Classic Recovery insists that if $\alpha \mid\!\sim \beta$ holds in an epistemic state *E*, then after a contraction by $\alpha \mid\!\sim \beta$, the belief set associated with the contracted *E* should be such that an expansion by $\alpha \rightarrow \beta$ will give us back the original belief set. But this does not take into account that it is possible that $\alpha \mid\!\sim \beta$ can follow from *E* without $\alpha \rightarrow \beta$ following from the belief set associated with *E*. Note that the following version of Recovery is a consequence of Containment.

(B-10) If $\alpha \rightarrow \beta \in \mathcal{B}^c(E)$, then $\mathcal{B}^c(E) = \mathcal{B}^c(E - \alpha \mid\!\sim \beta) + \alpha \rightarrow \beta$ (Classic Recovery)

We shall see below that it makes sense to consider a defeasible version of Recovery. For this we need an appropriate definition of defeasible *expansion*, which will be touched on in next section. In the meantime, note that Classical Recovery follows from Containment.

Definition 7. *$-$ is an **admissible defeasible contraction operator** if $E - \alpha \mid\!\sim \beta \not\Vdash \alpha \mid\!\sim \beta$ and $E - \alpha \mid\!\sim \beta = E - \gamma \mid\!\sim \delta$ if $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \gamma \mid\!\sim \delta$. An admissible contraction operator $-$ is a **basic defeasible contraction operator** if it is defined as follows: $E - \alpha \mid\!\sim \beta \stackrel{\text{def}}{=} E'$, where E' is an epistemic state such that $\mathcal{U}^f(E') = \mathcal{U}^f(E)$ if $\alpha \mid\!\sim \beta \notin \mathcal{B}(E)$, and $\mathcal{U}^f(E') = \mathcal{U}^f(E) \cup \min_E[\![\alpha \wedge \neg\beta]\!]$ otherwise.*

Theorem 4. *$-$ satisfies (B-1) to (B-8) if and only if it is a basic defeasible contraction operator.*

The Inertia postulate for contraction requires that an epistemic state remains unchanged after a contraction by $\alpha \mid\!\sim \beta$ when $\alpha \mid\!\sim \beta$ is not in the defeasible belief set associated with the epistemic satate.

If $\alpha \mid\!\sim \beta \notin \mathcal{B}(E)$, then $E - (\alpha \mid\!\sim \beta) = E$ (Inertia)

Inertia should not hold, in general, for reasons similar to the case for Inertia for revision. Instead, it is more reasonable to consider a weaker version of Inertia where the defeasible belief set associated with an epistemic state remains unchanged after a contraction by $\alpha \mid\!\sim \beta$ when $\alpha \mid\!\sim \beta$ is not in the defeasible belief set associated with the epistemic satate.

If $\alpha \mid\!\sim \beta \notin \mathcal{B}(E)$, then $\mathcal{B}(E - \alpha \mid\!\sim \beta) = \mathcal{B}(E)$ (Weak Inertia)

Basic defeasible contraction does not satisfy Weak Inertia.

Example 9. *For $\mathcal{P} = \{p, q\}$ let E be s.t. $E(\mathsf{pq}) = E(\mathsf{p\overline{q}}) = \langle \mathsf{f}, 0\rangle$ and $E(\mathsf{\overline{p}q}) = E(\mathsf{\overline{pq}}) = \langle \mathsf{f}, 1\rangle$. Observe that $\mathsf{p} \mid\!\sim \mathsf{q} \notin \mathcal{B}(E)$. Basic defeasible contraction allows for contraction by $\mathsf{p} \mid\!\sim \mathsf{q}$ to result in E' where $E'(\mathsf{p\overline{q}}) = \langle \mathsf{f}, 0\rangle$ and $E'(\mathsf{pq}) = E'(\mathsf{\overline{p}q}) = E'(\mathsf{\overline{pq}}) = \langle \mathsf{f}, 1\rangle$. Note that $\mathsf{q} \mid\!\sim \mathsf{p} \in \mathcal{B}(E)$, but $\mathsf{q} \mid\!\sim \mathsf{p} \notin \mathcal{B}(E')$.*

The postulate below is a semantic version of Weak Inertia.

If $\alpha \mid\!\sim \beta \notin \mathcal{B}(E)$, then $\mathcal{U}^f(E) = \mathcal{U}^f(E - \alpha \mid\!\sim \beta)$ and $E(u) = (E - \alpha \mid\!\sim \beta)(u)$ for all $u \in \mathcal{U}^f(E)$ (Semantic Weak Inertia)

Proposition 3. *A basic defeasible contraction operator $-$ satisfies Weak Inertia iff it satisfies Semantic Weak Inertia.*

The main difference between classsical propositional logic and the version of defeasibility logic presented in this paper is the ability to express defeasible implications, with statements of the the form $\alpha \mid\!\sim \beta$ being viewed as defeasible (and therefore weaker) versions of classical implications of the form $\alpha \rightarrow \beta$. Given this context, it makes sense to consider the following postulate.

If $\gamma \rightarrow \delta \in \mathcal{B}^c(E) \setminus \mathcal{B}^c(E - \alpha \mid\!\sim \beta)$, then $\gamma \mid\!\sim \delta \in \mathcal{B}(E - \alpha \mid\!\sim \beta)$. (Defeasibility)

Defeasibility is an expression of minimal change. It states that when a classical implication does not hold anymore, it should be weakened to a defeasible implication. This postulate can be expressed semantically as follows.

(S-0) If $v \notin \mathcal{U}^f(E)$, then $(E - \alpha \mid\!\sim \beta)(u) \prec (E - \alpha \mid\!\sim \beta)(v)$ for all $u \in \mathcal{U}^f(E)$ (Defeasible Weakening)

Proposition 4. *A basic defeasible contraction operator satisfies Defeasibility if and only if it satisfies (S0).*

6 Full Defeasible Contraction

In this section we strengthen basic defeasible contraction to obtain a version of full defeasible contraction that is the contraction counterpart of full (coherence) defeasible revision. The starting point is the specification of five postulates for iterated contraction.

(BDP-1) If $\alpha \models \gamma$, then $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha \vee \beta))$ iff $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha \vee \beta))$

(BDP-2) If $\gamma \models \alpha$, then $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha \vee \beta))$ iff $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha \vee \beta))$

(BDP-3) If $\neg\beta \mid\!\sim \gamma$, then if $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha \vee \beta))$, then $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha \vee \beta))$

(BDP-4) If $\gamma \mid\!\sim \beta$, then if $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma - \alpha \vee \beta))$, then $((\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha)$ implies $\neg\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha \vee \beta))$

(BDP-5) If $\alpha \models \neg\gamma$, then for all δ, $\alpha \mid\!\sim \delta \in \mathcal{B}(E - \alpha \mid\!\sim \beta)$ iff $\alpha \mid\!\sim \delta \in \mathcal{B}(E - \gamma \mid\!\sim \beta - \alpha \mid\!\sim \beta)$

The fifth postulate is the contraction analogue of (BDP*5) in the section on Full Defeasible Revision. The first four postulates are analogues of the iterated contraction postulates for classical contraction presented by Koniezny and Pino-Pérez (). The latter can be viewed as the contraction analogues of the Darwiche-Pearl postulates for classical revision. This is perhaps made more clear when considering the semantic versions of the five postulates above.

(SDP-1) If $u, v \in [\![\alpha \wedge \neg\beta]\!]$, then $E(u) \preceq E(v)$ iff $(E - \alpha \mid\!\sim \beta)(u) \preceq (E - \alpha \mid\!\sim \beta)(v)$

(SDP-2) If $u, v \in [\![\alpha \wedge \beta]\!]$, then $E(u) \preceq E(v)$ iff $(E - \alpha \mid\!\sim \beta)(u) \preceq (E - \alpha \mid\!\sim \beta)(v)$

(SDP-3) If $u \in [\![\alpha \wedge \neg\beta]\!]$ and $v \in [\![\alpha \wedge \beta]\!]$, then $E(u) \prec E(v)$ implies $(E - \alpha \mid\!\sim \beta)(u) \prec (E - \alpha \mid\!\sim \beta)(v)$

(SDP-4) If $u \in [\![\alpha \wedge \neg\beta]\!]$ and $v \in [\![\alpha \wedge \beta]\!]$, then $E(u) \preceq E(v)$ implies $(E - \alpha \mid\!\sim \beta)(u) \preceq (E - \alpha \mid\!\sim \beta)(v)$

(SDP-5) If $u, v \in [\![\neg\alpha]\!]$, then $E(u) \preceq E(v)$ iff $(E - \alpha \mid\!\sim \beta)(u) \preceq (E - \alpha \mid\!\sim \beta)(v)$

(SDP-1) states that elements of $[\![\alpha \wedge \neg\beta]\!]$ keep their relative ordering when contracting by $\alpha \mid\!\sim \beta$. (SDP-2) states that elements of $[\![\alpha \wedge \beta]\!]$ keep their relative ordering when contracting by $\alpha \mid\!\sim \beta$. (SDP-3) and (SDP-4) together state that when contracting by $\alpha \mid\!\sim \beta$, elements of $[\![\alpha \wedge \neg\beta]\!]$ can be slid "down" w.r.t. the elements of $[\![\alpha \wedge \beta]\!]$. (SDP-5) does not have a classical Darwiche-Pearl counterpart but is the contraction analogue of (SDP*5). It says that when contracting by $\alpha \mid\!\sim \beta$, the elements of $[\![\neg\alpha]\!]$ keep their relative ordering, but there is no constraint on how the elements of $[\![\neg\alpha]\!]$ move w.r.t. the elements of $[\![\alpha]\!]$.

Lemma 2. *Given postulates (B-1)-(B-8), for* $i \in \{1, 2, 3, 4, 5\}$, *(BDP-i) and (SDP-i) are equivalent.*

This allows us to define full defeasible contraction.

Definition 8. *A **full defeasible contraction operator** is defined to be a basic defeasible contraction operator satisfying (SDP-1)-(SDP-5).*

Theorem 5. $-$ *satisfies (B-1)-(B-8) and (BDP-1)-(BDP-5) if and only if it is a full defeasible contraction operator.*

We conclude this section with a concrete defeasible contraction operator. It is the contraction analogue of Natural (Coherence) Defeasible Revision. It moves the minimal elements of $[\![\alpha \wedge \neg\beta]\!]$ down as little as possible to ensure that $\alpha \mid\!\sim \beta$ does not hold in the contracted epistemic state.

Definition 9. *(Natural Defeasible Contraction)*

Natural Defeasible Contraction: $E - \alpha \mid\!\sim \beta \stackrel{\text{def}}{=} E'$ *where*

1. *if* $E \not\Vdash \alpha \mid\!\sim \beta$ *or* $\alpha \mid\!\sim \beta \equiv_{\mathcal{E}} \top \mid\!\sim \top$, *then* $E' = E$; *else*
2. $E'(u) = E(u)$ *for all* $u \in (\mathcal{U}^{\infty}(E) \setminus \min_E[\![\alpha \wedge \neg\beta]\!]) \cup \{v \in \mathcal{U}^f(E) \mid E(v) \prec r^E_{\min}[\![\alpha \wedge \beta]\!] \cap \mathcal{U}^f(E)\}$, $E'(u) = E(u) + 1$ *for all* $u \in \{v \in \mathcal{U}^f(E) \mid r^E_{\min}[\![\alpha \wedge \beta]\!] \preceq v\}$, *and* $E'(u) = r^E_{\min}[\![\alpha \wedge \beta]\!]$ *for all* $u \in \min_E[\![\alpha \wedge \neg\beta]\!]$.

Natural Defeasible Contraction is a full defeasible contraction operator that satisfies Inertia (and therefore Weak Inertia), as well as Defeasibility.

7 Interdefinability

In this section we investigate the extent to which defeasible revision and contraction can be defined in terms of each other. Our first results show that the Classical Levi and Harper Identities hold when restricting revision and contraction to classical statements.

Proposition 5. *1. For every basic (full) defeasible revision operator* $*$ *there is a basic (full) defeasible contraction operator* $-$ *s.t.* $\mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot) = \mathcal{B}^c(E - \alpha \mid\!\sim \bot) + \alpha$.

2. For every basic (full) coherence defeasible revision operator $*$ *there is a basic (full) defeasible contraction operator* $-$ *s.t.* $\mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot) = \mathcal{B}^c(E - \alpha \mid\!\sim \bot) + \alpha$.

3. For every basic (full) defeasible contraction operator $-$ *there is a basic (full) defeasible revision operator* $*$ *s.t.* $\mathcal{B}^c(E - \alpha \mid\!\sim \bot = \mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot) \cap \mathcal{B}^c(E)$.

4. For every basic (full) defeasible contraction operator $-$ *there is a basic (full) coherence defeasible revision operator* $*$ *s.t.* $\mathcal{B}^c(E - \alpha \mid\!\sim \bot = \mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot) \cap \mathcal{B}^c(E)$.

The next results show that the Levi Identity does not hold on the level of the belief sets associated with epistemic states, and that the Harper Identity holds only for coherence defeasible revision.

Proposition 6. *1. It is not always the case that for every basic (full) defeasible revision operator* $*$ *there is a basic (full) defeasible contraction operator* $-$ *s.t.* $\mathcal{B}^c(E * \alpha \mid\!\sim \beta) = \mathcal{B}^c(E - \alpha \mid\!\sim \neg\beta) + \alpha \to \beta$.

2. It is not always the case that for every basic (full) coherence defeasible revision operator $*$ *there is a basic (full) defeasible contraction operator* $-$ *s.t.* $\mathcal{B}^c(E * \alpha \mid\!\sim \beta) = \mathcal{B}^c(E - \alpha \mid\!\sim \neg\beta) + \alpha \to \beta$.

3. It is not always the case that for every basic (full) defeasible revision operator $*$ *there is a basic (full) defeasible contraction operator* $-$ *s.t.* $\mathcal{B}^c(E - \alpha \mid\!\sim \beta) = \mathcal{B}^c(E * \alpha \mid\!\sim \neg\beta) \cap \mathcal{B}^c(E)$.

4. For every basic (full) coherence defeasible revision operator $*$ *there is a basic (full) defeasible contraction operator* $-$ *s.t.* $\mathcal{B}^c(E - \alpha \mid\!\sim \beta) = \mathcal{B}^c(E * \alpha \mid\!\sim \neg\beta) \cap \mathcal{B}^c(E)$.

In order to investigate interdefinability on the level of defeasible belief sets we need appropriate versions of the Levi and Harper Identities on this level. For a version of the Levi Identity for defeasible belief sets, we need an appropriate version of *expansion* for defeasible belief sets, which means that we need an appropriate version of defeasible consequence. But defeasible consequence for a defeasible set $\mathcal{D}$ is not unique: it can be obtained from any epistemic state whose associated defeasible belief set contains $\mathcal{D}$.

Definition 10. *An operator $\mathcal{C}(\cdot)$ on the set of defeasible sets is a* ***defeasible consequence operator*** *if and only if for every defeasible set $\mathcal{D}$ there is an epistemic state E s.t. $\mathcal{D} \subseteq \mathcal{C}(\mathcal{D}) = \mathcal{B}(E)$. $+$ is a* ***defeasible expansion operator*** *if and only if there is a defeasible consequence operator $\mathcal{C}(\cdot)$ s.t. $\mathcal{B} + \alpha \mid\!\sim \beta = \mathcal{C}(\mathcal{B} \cup \{\alpha \mid\!\sim \beta\})$.*

While there is a case to be made for restricting this broad definition of defeasible consequence, such an investigation is beyond the scope of this paper.

The next results show that, on the level of defeasible belief sets associated with epistemic states, the Levi Identity does not hold for basic (full) defeasible revision, but it holds for basic (full) coherence revision.

Proposition 7. *1. It is not the case that for every basic (full) defeasible revision operator $*$ there is a basic (full) defeasible contraction operator $-$ and an expansion operator $+$ s.t. $\mathcal{B}(E * \alpha \mid\!\sim \beta) = \mathcal{B}(E - \alpha \mid\!\sim \neg\beta) + \alpha \mid\!\sim \beta$.*

2. For every basic (full) coherence defeasible revision operator $*$ there is a basic (full) defeasible contraction operator $-$ and a defeasible expansion operator $+$ s.t. $\mathcal{B}(E * \alpha \mid\!\sim \beta) = \mathcal{B}(E - \alpha \mid\!\sim \neg\beta) + \alpha \mid\!\sim \beta$.

This places us in a position to introduce a version of the Recovery postulate for defeasible belief sets.

If $\alpha \mid\!\sim \beta \in \mathcal{B}(E)$ and $\alpha \mid\!\sim \bot \notin \mathcal{B}(E)$ then
$\mathcal{B}(E - \alpha \mid\!\sim \beta + \alpha \mid\!\sim \beta) = \mathcal{B}(E)$ (Recovery)

Proposition 8. *For every basic defeasible contraction operator there is a defeasible expansion operator s.t. Recovery holds.*

We conclude this section with an investigation into a version of the Harper Identity for defeasible belief sets.

$\mathcal{B}(E - \alpha \mid\!\sim \beta) =$
$\mathcal{B}(E * \alpha \mid\!\sim \neg\beta) \cap \mathcal{B}(E)$ (Harper Identity)

Our next results show that it does not hold for basic (full) defeasible revision or for basic (full) coherence revision. Not even if Inertia holds.

Proposition 9. *1. It is not the case that for every basic (full) defeasible contraction operator satisfying Inertia there is a basic (full) defeasible revision operator s.t. the Harper Identity holds.*

2. It is not the case that for every basic (full) defeasible contraction operator satisfying Inertia there is a basic (full) coherence defeasible revision operator s.t. the Harper Identity holds.

8 Related Work and Conclusion

In this paper we characterised belief change (revision and contraction) for defeasible reasoning in terms of postulates and semantic constructions, with a focus on the defeasible belief sets associated with epistemic states.

The transformation of material implications ($\alpha \rightarrow \beta$) into defeasible implications ($\alpha \mid\!\sim \beta$), as set out in Example 1, was originally proposed by Falappa et al (2002; 2013).

The connection between belief change and non-monotonic reasoning is well-known (Gärdenfors and Makinson 1994). Initial attempts to combine these two areas resulted in the Gardënfors' paradox (Gärdenfors 1988; Rott 1989) which deals with the revision of subjunctive conditionals representing the revision policies themselves. Most of the work dealing with the revision of conditionals, such as those by Kern-Isberner (1999; 2018) and Wobcke (1995), assume such a perspective. They model the revision of subjunctive conditionals representing the revision operator in the object language. In contrast, our versions of conditionals (defeasible statements) do not encode the revision policies, but rather follow from such policies which are encoded in the epistemic states. In our case there is a clear separation between object level and meta level. Note that only one direction of the Ramsey Test (Gärdenfors 1988) holds for our proposal. While it is the case that $\beta \in \mathcal{B}^c(E * \neg\alpha \mid\!\sim \bot)$ implies $\alpha \mid\!\sim \beta \in \mathcal{B}(E)$, the converse does not hold (recall that α can be represented as $\neg\alpha \mid\!\sim \bot$).

In the original work of Kern-Isberner (1999) her epistemic states have less structure than ours. The two approaches are compatible, but not equivalent. E.g., our approach doesn't satisfy her CR2 postulate. Our motivation for CR2 not holding is related to the additional structure of our epistemic states. Also, while her CR3 postulate can be made to hold in our framework, it carries with it an assumption that is at odds with our proposal: that $\top \mid\!\sim \alpha$ is an appropriate representation of the propositional statement α. In our approach such an α is expressible as $\neg\alpha \mid\!\sim \bot$.

Strongly connected to the present paper is previous work by Casini et al. (2016; 2017; 2018), formulating a general AGM-like approach for defeasible reasoning. While the main overall focus is the same, there are important differences. (*i*) We focus on defeasible belief sets that are *rational* (Lehmann and Magidor 1992), which allows for a simpler semantic representation. (*ii*) The simpler semantic representation used here has some other advantages, such as the ability to perform iterated change. (*iii*) The previous framework is premised on the ability to distinguish, within a defeasible belief set, between those that are truly defeasible, and those that are not, the latter forming the *monotonic core* of the defeasible belief set (Casini and Meyer 2017). Ware unable to draw such a distinction. Or rather, our assumption is the boundary case where the monotonic core of a defeasible belief set associated with an epistemic state *E* coincides with the classical belief set associated with *E*. Other recent contributions to revision in a non-monotonic framework include work from Hunter (2016), dealing with highly implausible conditionals, and Delgrande et al. (2013), analysing belief change in an ASP framework.

The next step is to enrich the representation of epistemic states by including the notion of the monotonic core (Casini and Meyer 2017) of a defeasible belief set. Beyond generalising the present proposal, this will allow us to link the current paper more directly to our previous work. Other future work includes an extension to other well-known forms of iterated change, such as lexicographic revision (Nayak 1994), as well as restrained revision and the more general class of admissible revision operators (Booth and Meyer 2006).

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 Meyer, T. A. 2006. Admissible and restrained revision. *J. Artif. Intell. Res.* 26:127–151.

Booth, R.; Meyer, T.; Varzinczak, I.; and Wassermann, R. 2011. On the link between partial meet, kernel, and infra contraction and its application to horn logic. *Journal of Artificial Intelligence Research (JAIR)* 42:31–53.

Booth, R.; Meyer, T.; and Varzinczak, I. 2009. Next steps in propositional Horn contraction. In Boutilier, C., ed., *IJCAI*, 702–707. AAAI Press.

Boutilier, C. 1996. Iterated revision and minimal change of conditional beliefs. *J. Philosophical Logic*.

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.; Fermé, E.; Meyer, T.; and Varzinczak, I. 2018. A semantic perspective on belief change in a preferential non-monotonic framework. In Thielscher, M.; Toni, F.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, 220–229. AAAI Press.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artif. Intell.* 89(1–2):1–29.

Delgrande, J. P., and Peppas, P. 2015. Belief revision in horn theories. *Artif. Intell.* 218:1–22.

Delgrande, J. P., and Wassermann, R. 2013. Horn clause contraction functions. *J. Artif. Intell. Res.* 48:475–511.

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.

Falappa, M. A.; García, A. J.; Kern-Isberner, G.; and Simari, G. R. 2013. Stratified belief bases revision with argumentative inference. *Journal of Philosophical Logic* 42(1):161–193.

Falappa, M. A.; Kern-Isberner, G.; and Simari, G. R. 2002. Explanations, belief revision and defeasible reasoning. *Artificial Intelligence* 141(1/2):1–28.

Gärdenfors, P., and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artificial Intelligence* 65(2):197–245.

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

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.

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 Publishers Inc.

Kern-Isberner, G. 2018. Axiomatizing a qualitative principle of conditional preservation for iterated belief change. In Thielscher, M.; Toni, F.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, 248–256. AAAI Press.

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

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

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

Meyer, T.; Lee, K.; and Booth, R. 2005. Knowledge integration for description logics. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 645–650. AAAI Press/MIT Press.

Nayak, A. C. 1994. Foundational belief change. *J. Philosophical Logic* 23(5):495–533.

Rott, H. 1989. Conditionals and theory change: Revisions, expansions, and additions. *Synthese* 81(1):91–113.

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

Zhuang, Z.; Wang, Z.; Wang, K.; and Qi, G. 2016. Dl-lite contraction and revision. *J. Artif. Intell. Res.* 56:329–378.

Zhuang, Z.; Wang, Z.; Wang, K.; and Delgrande, J. P. 2019. A generalisation of AGM contraction and revision to fragments of first-order logic. *J. Artif. Intell. Res.* 64:147–179.