Annals of Mathematics and Artificial Intelligence (2023) 91:489–515
https://doi.org/10.1007/s10472-023-09835-4

Deductive belief change

Theofanis Aravanis[1]

Accepted: 29 January 2023 / Published online: 13 February 2023

Abstract

In a 2003-article, Sven Ove Hansson discusses the justificatory structure of a belief base, by highlighting that some beliefs of the belief base are held only because they are (deductively) justified by some other beliefs. He concludes that the relation between the justificatory structure of a belief base and the vulnerability of its beliefs (which in turn reflects their resistance to change) remains an open issue, both on a conceptual and on a technical level. Motivated by Hanssons' remarks, we introduce in this article a new interesting type of change-operation, called deductive belief change (contraction and revision), and abbreviated as DBC. DBC associates in a natural manner the deductive justification that the logical sentences of the language have, in the context of a belief base B, with their vulnerability relative to B. According to DBC, the more explicit B-beliefs imply a sentence φ, the more resistant to change φ is, with respect to B. We characterize DBC both axiomatically, in terms of natural postulates, and constructively, in terms of kernel belief change, illustrating its simple and intuitive structure. Interestingly enough, as we prove, kernel belief change (and its central specialization partial-meet belief change) already encodes a strong coupling between justificatory structure and vulnerability, as it implements DBC. Furthermore, we show that deductive belief revision, properly adapted to the belief-sets realm, is indistinguishable from Parikh's relevance-sensitive revision, a fundamental type of revision which, due to its favourable properties, constitutes a promising candidate for a variety of real-world applications. As a last contribution, we study relevance in the context of belief bases, and prove that kernel belief change respects Parikh's notion of relevance.

Keywords Deductive justification · Vulnerability · Kernel belief change · Relevance · Knowledge representation

Mathematics Subject Classification (2010) 03B42 · 68T27 · 68T30

✉ Theofanis Aravanis
taravanis@upatras.gr

[1] Department of Mechanical Engineering, School of Engineering, University of the Peloponnese, Patras 263 34, Greece

1 Introduction

An intelligent agent should be capable of gathering information about the world, and modifying their state of belief in response to new evidential input. As a consequence, the agent should be able to perform *belief change* —namely, *belief contraction* (simply contraction) and *belief revision* (simply revision)— an operation which is heavily studied in the realm of Artificial Intelligence [16].

The process of belief change was formalized by Alchourrón, Gärdenfors and Makinson in the cornerstone work [1], in which a versatile framework was introduced, now called the *AGM paradigm*, after the initials of its three founders. Within the AGM paradigm, the belief corpus of an agent is typically represented by a logical theory K (which is an *infinite* set of sentences closed under logical consequence, also referred to as a *belief set*), the *epistemic input* (i.e., new information) is represented by a logical formula φ, and the contraction/revision of K in response to φ is constrained by well-accepted *rationality-postulates*.

One of the most controversial assumptions of the AGM paradigm is the modelling of an agent's belief corpus as a belief set [32]. The main concerns regarding belief sets can be summarized in the following points: Belief sets make no distinction between *explicit* beliefs and *implicit* beliefs (with the latter ones being "merely derived" from the former ones), they are computationally intractable, and they cannot distinguish between different inconsistent states of belief. In response to these weaknesses of belief sets, several authors have proposed the use of *belief bases* as a model for representing states of belief. A belief base is a set of sentences of the language, which, contrary to a belief set, is *not* closed under logical consequence, and for all practical purposes it is *finite*; as such, a belief base does not suffer from the aforementioned shortcomings of belief sets.[1]

In a 2003-article [34], Sven Ove Hansson introduces ten principal open problems of the belief-change theory, related to the representation of states of belief, to different notions of degrees of belief, and to the nature of change-operations. In that article, Hansson discusses the *justificatory structure* of a belief base, by highlighting that some beliefs of the belief base are held only because they are (deductively) *justified* by some other beliefs. He, thereafter, wonders to what degree that justificatory structure relates to the *vulnerability* of the beliefs, which in turn reflects their resistance to change. Finally, Hansson concludes to the following statement, which is borrowed directly from [34, Section 5].

> *"The relation between vulnerability and justificatory structure remains an open issue. It is not clear, either on a conceptual or a technical level, to what degree the justificatory structure can be expressed in terms of vulnerability, or vice versa."*

Motivated by the aforementioned Hanssons' remarks, we introduce in this article a new interesting type of change-operation, called *deductive belief change* (abbreviated as DBC). DBC comes to associate in a natural manner the *deductive justification* that the logical sentences of the language have, in the context of a belief base B, with their vulnerability relative to B. The core principle of DBC is that, if a sentence ψ is implied by *set-theoretically more explicit* B-beliefs than a sentence φ, then ψ is more resistant to change than φ, with respect

[1] The assumption that the belief-base approach corresponds to a *foundationalist* epistemology, whereas, a belief-set approach represents a *coherentist* epistemology constitutes a common ground in the belief-change literature [25].

to B. For demonstrating its simple and intuitive structure, DBC is defined both *axiomatically*, in terms of natural postulates, and *constructively*, in terms of *kernel belief change* [31] (in the context of which, belief bases are modified via a mechanism that selects sentences to be retracted).

As we prove, kernel belief change —and thus all its specializations, including *partial-meet* belief change originally developed by the AGM trio in [1]— implements DBC, hence, it provides a proof of concept for the introduced type of belief change. This is an important result pointing out that the fundamental kernel belief change *already* encodes a strong coupling between the justificatory structure of belief bases and the vulnerability of their beliefs.

Beyond belief bases, we study DBC in the realm of belief sets. Accordingly, we show that an appropriate refinement of deductive belief revision for belief sets is *indistinguishable* from Parikh's *relevance-sensitive* revision, a principal type of revision that is more well-behaved than the one identified by the AGM paradigm, which has been criticized as overly liberal towards relevance [47, 48].[2] As a last contribution, we study relevance in the context of belief bases, and prove that kernel belief change (on belief bases) respects Parikh's notion of relevance.

The remainder of this article is organized as follows: The next section establishes the formal background for our discussion. Thereafter, Section 3 presents a brief overview of the AGM paradigm, followed by Section 4 which discusses kernel belief change on belief bases. Section 5 explores representative works on change-operations for belief bases. In Section 6, DBC is introduced, whereas, in Section 7, it is shown that the introduced type of belief change is in fact implemented by kernel belief change. Section 8 investigates DBC in the realm of belief sets, as well as the relevance-sensitivity of change-operations on belief bases. The article closes with some concluding remarks and promising avenues for feature research.

2 Formal prelude

In the present article, we shall work with a propositional language $\mathcal{L}$, built over a *finite*, non-empty set $\mathcal{P}$ of atoms (propositional variables), using the standard Boolean connectives $\wedge$ (conjunction), $\vee$ (disjunction), $\rightarrow$ (implication), $\leftrightarrow$ (equivalence), $\neg$ (negation), and governed by *classical propositional logic*. The classical consequence relation is denoted by $\models$.

A sentence φ of $\mathcal{L}$ is *contingent* iff it is neither a tautology (i.e., $\not\models \varphi$) nor a contradiction (i.e., $\not\models \neg\varphi$). For a set of sentences Γ, $Cn(\Gamma)$ denotes the set of all logical consequences of Γ; i.e., $Cn(\Gamma) = \{\varphi \in \mathcal{L} : \Gamma \models \varphi\}$. For a set of sentences Γ and a sentence $\varphi \in \mathcal{L}$, $\Gamma \models \varphi$ stands for an abbreviation of $\varphi \in Cn(\Gamma)$. For sentences $\varphi_1, \ldots, \varphi_n$ of $\mathcal{L}$, we shall write $Cn(\varphi_1, \ldots, \varphi_n)$ instead of $Cn(\{\varphi_1, \ldots, \varphi_n\})$. Since the language $\mathcal{L}$ is governed by classical propositional logic, it follows that the operator Cn satisfies the standard Tarskian properties, hence, among others, it satisfies *deduction*; that is, for any set of sentences Γ and any two sentences $\varphi, \psi \in \mathcal{L}$, $\psi \in Cn(\Gamma \cup \{\varphi\})$ iff $(\varphi \rightarrow \psi) \in Cn(\Gamma)$.

[2]Relevance is a vital principle, not only in the realm of belief change, but also in the context of Artificial Intelligence in general [27].

A set of sentences of $\mathcal{L}$ that is closed under logical consequence is called a *belief set* (or theory); otherwise, it is called a *belief base*. A sentence φ of $\mathcal{L}$ is a belief of a belief base (or belief set) B iff $B \models \varphi$, or equivalently $\varphi \in Cn(B)$. Although belief bases need not be finite, most work on them focuses on the finite case, and this tactic shall be followed herein as well. For a belief set K and a sentence φ, we define the *expansion* of K by φ, denoted by $K + \varphi$, as the logical closure of the set $K \cup \{\varphi\}$; in symbols, $K + \varphi = Cn\big(K \cup \{\varphi\}\big)$.

A *literal* is an atom $p \in \mathcal{P}$ or its negation. For a set of literals Q, $\bigvee Q$ denotes the sentence of $\mathcal{L}$ resulting from the disjunction of all the literals in Q, and $\overline{Q}$ denotes the set of all the negated literals in Q. A *possible world* (simply *world*) r is a consistent set of literals, such that, for every atom $p \in \mathcal{P}$, either $p \in r$ or $\neg p \in r$. The set of all possible worlds is denoted by $\mathbb{M}$. For a sentence or set of sentences φ of $\mathcal{L}$, $[\varphi]$ denotes the set of worlds at which φ is true.

A *preorder* over a set of possible worlds M is a reflexive, transitive binary relation in M. A preorder $\preceq$ is *total* iff, for all $r, r' \in M$, $r \preceq r'$ or $r' \preceq r$. The strict part of $\preceq$ is denoted by $\prec$; i.e., $r \prec r'$ iff $r \preceq r'$ and $r' \npreceq r$. The indifference part of $\preceq$ is denoted by $\approx$; i.e., $r \approx r'$ iff $r \preceq r'$ and $r' \preceq r$. Furthermore, $\min(M, \preceq)$ denotes the set of all $\preceq$-minimal worlds of M; i.e.,

$$\min(M, \preceq) = \Big\{ r \in M : \text{ for all } r' \in M, \text{ if } r' \preceq r, \text{ then } r \preceq r' \Big\}.$$

3 The AGM paradigm

As already stated, within the AGM paradigm, the belief corpus of an agent is represented by a belief set K and the epistemic input is represented by a sentence φ of $\mathcal{L}$. Between K and φ, the AGM paradigm considers two fundamental change-operations, that is, *belief contraction* (or simply contraction) and *belief revision* (or simply revision). In their seminal article [1], Alchourrón, Gärdenfors and Makinson characterized both these types of change-operations *axiomatically*, in terms of well-accepted *rationality-postulates*, as well as *constructively*, in terms of the so-called *partial-meet constructive model*, which is based on a mechanism that selects sets of sentences of a belief corpus that are eligible to be retained. In a subsequent work [39], Katsuno and Mendelzon developed another popular constructive model for the change-operations of the AGM paradigm, which is based on a special type of total preorders over possible worlds, called *faithful preorders*. In this section, we present the axiomatic side of the AGM paradigm, as well as the faithful-preorders constructive model (given that it shall be used in Section 8). Furthermore, we highlight a well-known connection between the operations of contraction and revision.

3.1 Axiomatic characterization

Alchourrón, Gärdenfors and Makinson model the process of contraction by a *contraction function* $\dot{-}$, which is a binary function that maps a belief set K and a sentence φ to a belief set $K \dot{-} \varphi$, representing the result of contracting φ from K. We shall say that a contraction function $\dot{-}$ is an *AGM contraction function* iff it respects the following postulates, known as the *AGM contraction postulates* [1].

$(K \dot{-} 1)$ $K \dot{-} \varphi$ is a theory.

$(K \dot{-} 2)$ $K \dot{-} \varphi \subseteq K$.

$(K \dot{-} 3)$ If $\varphi \notin K$, then $K \dot{-} \varphi = K$.

$(K \dot{-} 4)$ If φ is not tautological, then $\varphi \notin K \dot{-} \varphi$.

$(K \dot{-} 5)$ If $\varphi \in K$, then $K \subseteq (K \dot{-} \varphi) + \varphi$.

$(K \dot{-} 6)$ If $Cn(\varphi) = Cn(\psi)$, then $K \dot{-} \varphi = K \dot{-} \psi$.

$(K \dot{-} 7)$ $(K \dot{-} \varphi) \cap (K \dot{-} \psi) \subseteq K \dot{-} (\varphi \wedge \psi)$.

$(K \dot{-} 8)$ If $\varphi \notin K \dot{-} (\varphi \wedge \psi)$, then $K \dot{-} (\varphi \wedge \psi) \subseteq K \dot{-} \varphi$.

Likewise, the process of revision is encoded into a *revision function* $*$, which is a binary function that maps a belief set K and a sentence φ to a belief set $K * \varphi$, representing the result of revising K by φ. We shall say that a revision function $*$ is an *AGM revision function* iff it respects the following postulates, known as the *AGM revision postulates* [1].

$(K * 1)$ $K * \varphi$ is a theory.

$(K * 2)$ $\varphi \in K * \varphi$.

$(K * 3)$ $K * \varphi \subseteq K + \varphi$.

$(K * 4)$ If $\neg\varphi \notin K$, then $K + \varphi \subseteq K * \varphi$.

$(K * 5)$ $K * \varphi$ is inconsistent iff φ is inconsistent.

$(K * 6)$ If $Cn(\varphi) = Cn(\psi)$, then $K * \varphi = K * \psi$.

$(K * 7)$ $K * (\varphi \wedge \psi) \subseteq (K * \varphi) + \psi$.

$(K * 8)$ If $\neg\psi \notin K * \varphi$, then $(K * \varphi) + \psi \subseteq K * (\varphi \wedge \psi)$.

A discussion on the AGM contraction/revision postulates can be found in [23, Chapter 3] and [49, Section 8.3]. Herein, we will suffice to state that the guiding principle behind $(K \dot{-} 1)$–$(K \dot{-} 8)$ and $(K * 1)$–$(K * 8)$ is the *economy of information*; that is, the modification of K, in response to the new information φ, should be done so that K is changed as little as possible.

Each one of the two aforementioned sets of postulates does *not uniquely* define particular change-operations; it only identifies the class of all different *rational* change-strategies. The definition of specific change-operations that respect the AGM contraction/revision postulates is implemented through the so-called *constructive models* for belief change, which constitute formalizations that provide "recipes" for specifying particular change-strategies. Two such constructive models are the partial-meet model, which was already proposed by the AGM trio in their seminal article [1], and the faithful-preorders model by Katsuno and Mendelzon [39]. Herein, we shall focus on the latter constructive model, which we discuss subsequently.

3.2 Constructive characterization: the faithful-preorders model

Katsuno and Mendelzon built their constructive model for the case of revision, based on Grove's system of spheres [28]. At the heart of Katsuno and Mendelzon's approach lies a special kind of total preorder over all possible worlds, called *faithful preorder* [39].

Definition 1 (Faithful Preorder, [39]) A preorder $\preceq_K$ over $\mathbb{M}$ is faithful to a belief set K iff it is total, and such that $[K] \neq \varnothing$ entails $[K] = \min(\mathbb{M}, \preceq_K)$.

Intuitively, a total preorder $\preceq_K$ over $\mathbb{M}$ encodes the *comparative plausibility* of the possible worlds of $\mathbb{M}$, relative to the belief set K. Hence, $r \preceq_K r'$ states that the world r is at least as plausible as the world r', with respect to K.

Relying on the notion of faithful preorder, Katsuno and Mendelzon proved in [39] that a revision function $*$ satisfies the AGM revision postulates iff, for *each* belief set K, there exists a faithful preorder $\preceq_K$ over $\mathbb{M}$, such that, for any $\varphi \in \mathcal{L}$:

$$\textbf{(F}*\textbf{)} \quad [K * \varphi] = \min([\varphi], \preceq_K).$$

Therefore, an AGM revision function $*$ can be constructed by means of a family $\{\preceq_K\}_{\forall K}$ of faithful preorders (one for each belief set K), with the aid of condition (F$*$); that is, by specifying the revised belief set $K * \varphi$ as the theory corresponding to the most plausible φ-worlds, with respect to K.

3.3 From contraction to revision

The change-operation identified by the AGM contraction postulates and that identified by the AGM revision postulates are not independent to each other; on the contrary, they are *closely related*. That connection was suggested by Isaac Levi in [43], before Alchourrón, Gärdenfors and Makinson formulated their rationality-postulates. According to Levi, one should in principle be able to define revision in terms of contraction by means of the following procedure: For revising a belief set K by an epistemic input φ, first, retract $\neg\varphi$ from K (thus anything that may contradict the epistemic input is removed), and then expand the resulting belief set with φ. This method connecting revisions and contractions is encoded into the following condition (LV), which is known as the *Levi Identity*.

$$\textbf{(LV)} \quad K * \varphi = \big(K \dot{-} \neg\varphi\big) + \varphi.$$

Interestingly, Alchourrón, Gärdenfors and Makinson proved that the revision function $*$ generated from an AGM contraction function $\dot{-}$, through the Levi Identity, is an AGM revision function, and, conversely, any AGM revision function $*$ can be generated from an AGM contraction function $\dot{-}$, by means of the Levi Identity [1].

We conclude our discussion on the AGM paradigm noting that several AGM-compliant types of change-operations (particularly, revision-operations) have been proposed in the literature, each one with its own favourable properties that make it suitable for certain applications — the interested reader is, indicatively, referred to the Hamming-based change-method by Dalal [9], its generalization parametrized-difference belief revision [50], as well as uniform belief revision [3, 8] and theory-relational belief revision [5]. Following this line of research, the main aim of this article is the introduction of a well-behaved type of change-operation applicable (primarily) to *belief bases*. In order to facilitate our presentation, we shall consider in what follows only the principal case of the modification of *consistent* belief corpora by *contingent* epistemic inputs.

4 Kernel belief change on belief bases

As noted in the Introduction section, one of the most controversial assumptions of the AGM paradigm is the modelling of an agent's belief corpus as a belief set [32]. A common concern regarding belief sets is that they are *infinite* objects, and, as such, cannot be incorporated directly into real-world Artificial-Intelligence applications.[3] In response to this concern, *belief bases* have been proposed as an alternative to belief sets.

Beyond computational considerations, one may also argue in favour of belief bases due to their ability of distinguishing the *unique syntax/structure* of beliefs. For example, given that $\mathcal{P} = \{a, b\}$, the belief bases $B_1 = \{a, b\}$ and $B_2 = \{a, a \leftrightarrow b\}$ represent identical beliefs, as they have the same logical closure, yet, they represent genuinely *distinct* states of belief. This can be evident if we consider the *dynamics* of B_1 and B_2. Suppose, for instance, that we revise B_1 and B_2 so that the new belief $\neg a$ be included into both of them. A plausible outcome for the revision of B_1 by $\neg a$ is the belief base $B_1' = \{\neg a, b\}$, whereas, a plausible outcome for the revision of B_2 by $\neg a$ is the belief base $B_2' = \{\neg a, a \leftrightarrow b\}$. Notice, now, that a belief of B_1' is the sentence b, whereas, a belief of B_2' is the sentence $\neg b$. The aforementioned scenario indicates that two *statically equivalent* belief bases (i.e., two belief bases that represent identical beliefs) are *not* in general *dynamically equivalent*, as their modification (with respect to new epistemic input) may lead to different outcomes. Thus, in a sense, belief bases are *more expressive* than belief sets.[4]

One other favourable feature of belief bases —which is not present in belief sets— is their ability of distinguishing the *explicit* from the *implicit* beliefs. The explicit beliefs of a belief base B are all the sentences contained in B, whereas, the implicit beliefs of B are all the sentences that logically follow from B, but are *not* in B. In this sense, the explicit beliefs are *independent* of any other beliefs, whereas, the implicit beliefs have no independent standing and are "merely derived" from the explicit beliefs; hence, if an implicit belief loses its (deductive) support from the explicit beliefs, then it will be automatically discarded.

In this article, we shall focus on a fundamental type of change-operation on belief bases, called *kernel contraction*. Kernel contraction was proposed by Hansson in [31], and originates from Alchourrón and Makinson's *safe contraction* [2]. Before discussing Hansson's proposal, we note that, in the realm of belief bases, a contraction function $\dot{-}$ is a binary function that maps a belief base B and a sentence φ to a belief base $B \dot{-} \varphi$, representing the result of contracting φ from B. Likewise, a revision function $*$ is a binary function that maps a belief base B and a sentence φ to a belief base $B * \varphi$, representing the result of revising B by φ. On that premises, let us introduce the notions of *kernel* and *incision function*, upon which Hansson defines his contraction functions.

Definition 2 (Kernel, [31]) Let B be a belief base and let φ be a sentence of $\mathcal{L}$. A φ-kernel of B is a $\subseteq$-minimal subset of B that entails φ; i.e., B' is a φ-kernel of B iff $B' \subseteq B$, $B' \models \varphi$, and no proper subset of B' entails φ.

The set of all φ-kernels of B shall be denoted by $B \perp\!\!\!\perp \varphi$.

Now, in order to retract a belief φ from B, *at least one* element of *each* φ-kernel of B must be removed from B, since otherwise φ would still be implied. This retraction is

---

[3] Although, for a propositional language $\mathcal{L}$ built from *finitely* many atoms (as it is the case herein), any belief set is *finitely axiomatizable* (i.e., it can be represented as the logical closure of a single sentence of $\mathcal{L}$).

[4] A consequence of that feature is that belief bases allow us to express different types of inconsistent states of belief, contrary to belief sets, in the context of which there is *only one* inconsistent belief set, which coincides with the whole language $\mathcal{L}$.

implemented through an *incision function*, which makes an incision (cut) into *each* φ-kernel of B.

Definition 3 (Incision Function, [31]) An incision function σ is a function such that, for every belief base B and all sentences $\varphi \in \mathcal{L}$, $\sigma(B \perp\!\!\!\perp \varphi) \subseteq \bigcup (B \perp\!\!\!\perp \varphi)$, and if $\varnothing \neq X \in B \perp\!\!\!\perp \varphi$, then $X \cap \sigma(B \perp\!\!\!\perp \varphi) \neq \varnothing$.

Every incision function σ gives rise to a contraction function $\dot{-}_\sigma$, defined by means of condition (K$\dot{-}$), for any belief base B and any sentence $\varphi \in \mathcal{L}$.

$$(K\dot{-}) \quad B \dot{-}_\sigma \varphi = B - \sigma\big(B \perp\!\!\!\perp \varphi\big).$$

Thus, the $\dot{-}_\sigma$-contraction of φ from B produces a set that contains all the elements of B that have *not been cut off* by the incision function σ.

Any contraction function $\dot{-}_\sigma$ induced from an incision function σ, by means of condition (K$\dot{-}$), shall be called *kernel contraction function*. Hansson establishes the following representation result, which *axiomatically* characterizes the class of kernel contraction functions, in terms of four postulates.

Theorem 1 ([33]) *A contraction function* $\dot{-}$ *is a kernel contraction function iff it satisfies the following postulates* $(B \dot{-} 1)$–$(B \dot{-} 4)$*:*

$(B \dot{-} 1)$ *If* φ *is not tautological, then* $B \dot{-} \varphi \nvDash \varphi$.

$(B \dot{-} 2)$ $B \dot{-} \varphi \subseteq B$.

$(B \dot{-} 3)$ *If, for all subsets* B' *of* B, $B' \models \varphi$ *iff* $B' \models \psi$, *then* $B \dot{-} \varphi = B \dot{-} \psi$.

$(B \dot{-} 4)$ *If* $\psi \in B$ *and* $\psi \notin B \dot{-} \varphi$, *then there is a subset* B' *of* B, *such that* $B' \nvDash \varphi$ *and* $B' \cup \{\psi\} \models \varphi$.

The rationale behind postulates $(B \dot{-} 1)$–$(B \dot{-} 4)$ that follows is borrowed from [31] and [16]. Postulate $(B \dot{-} 1)$, named *Success*, states that, if the epistemic input is not tautological, then it is not a belief of the belief base $B \dot{-} \varphi$. Postulate $(B \dot{-} 2)$, named *Inclusion*, requires that the contracted belief base $B \dot{-} \varphi$ is a subset of the initial belief base B.[5] Postulate $(B \dot{-} 3)$, called *Uniformity*, asserts that, if two sentences φ and ψ have the same behaviour relative to a belief base B (that is, if they are implied by the same subsets of B), then the contraction of φ from B is identical to the contraction of ψ from B. Lastly, postulate $(B \dot{-} 4)$, also referred to as *Core-Retainment*, prevents unnecessary losses of beliefs, by requiring that, if a belief ψ is removed from a belief base B due to the contraction of φ from B, then ψ contributes in some way to make B imply φ.

The following example concretely illustrates the operation of kernel contraction.

Example 1 (Kernel Contraction) Let $\mathcal{P} = \{a, b, c\}$ and let $B = \big\{a, a \leftrightarrow b, a \vee b\big\}$ be a belief base. Clearly then, the sentence $\varphi = a \wedge b$ is an implicit belief of B. The set of

[5] Postulates $(B \dot{-} 1)$ and $(B \dot{-} 2)$ encode the intuition of the AGM contraction postulates $(K \dot{-} 2)$ and $(K \dot{-} 4)$, respectively.

all φ-kernels of B is $B \perp\!\!\!\perp \varphi = \big\{\{a, a \leftrightarrow b\}, \{a \vee b, a \leftrightarrow b\}\big\}$. Now, let σ_1 be an incision function such that $\sigma_1\big(B \perp\!\!\!\perp \varphi\big) = \{a \leftrightarrow b\}$, and let σ_2 be an incision function such that $\sigma_2\big(B \perp\!\!\!\perp \varphi\big) = \{a, a \leftrightarrow b\}$. The incision functions σ_1, σ_2 induce in turn, via condition (K$\dot{-}$), two kernel contraction functions $\dot{-}_{\sigma_1}$, $\dot{-}_{\sigma_2}$, respectively, such that $B \dot{-}_{\sigma_1} \varphi = \big\{a, a \leftrightarrow b, a \vee b\big\} - \{a \leftrightarrow b\} = \{a, a \vee b\}$ and $B \dot{-}_{\sigma_2} \varphi = \big\{a, a \leftrightarrow b, a \vee b\big\} - \{a, a \leftrightarrow b\} = \{a \vee b\}$. Observe that the contraction-strategies of the operators $\dot{-}_{\sigma_1}$ and $\dot{-}_{\sigma_2}$ are distinct, which is to be expected since $\sigma_1 \neq \sigma_2$.

It is noteworthy that kernel contraction circumscribes a *very general* category of contraction-strategies, which has close connections with a plethora of important frameworks, such as *Truth Maintenance Systems* [10–12] and *Argumentation Systems* [52]. By imposing restrictions on the operation of incision functions, several well-studied sub-classes of kernel contraction functions can be identified. For example, by imposing *smoothness* on incision functions, we get the so-called *smooth kernel contraction* [31]. By *further* restricting incision functions, we get *partial-meet contraction*, which basically constitutes the belief-base counterpart of the partial-meet constructive model for belief sets, discussed in Section 3 [13].[6] Both the aforementioned indicative types of kernel contraction play a cornerstone role in the belief-change literature; the interested reader can find details on them in [33]. We only mention here the following result by Hansson, which concerns the relation between (smooth or unsmooth) kernel contraction and partial-meet contraction.

Proposition 1 [31] *The family of smooth kernel contraction functions forms a proper sub-class of the family of kernel contraction functions, and the family of partial-meet contraction functions forms a proper sub-class of the family of smooth kernel contraction functions.*

We conclude our discussion on kernel contraction noting that, although Theorem 1 refers to contraction functions, a variant of the Levi Identity, presented below as condition (LVB), allows us to construct a revision function $*$ on belief bases from a contraction function $\dot{-}$ on belief bases; notice that condition (LVB), contrary to condition (LV) of Section 3, yields an outcome that is *not* closed under logical consequence.

$$\textbf{(LVB)} \qquad B * \varphi = \big(B \dot{-} \neg\varphi\big) \cup \{\varphi\}.$$

Hence, with the aid of condition (LVB), we can specify (smooth or unsmooth) kernel revision functions from (smooth or unsmooth) kernel contraction functions, as well as partial-meet revision functions from partial-meet contraction functions. Note, lastly, that a result totally symmetric to Proposition 1 can also be formulated for the relation among kernel revision, smooth kernel revision and partial-meet revision.[7]

[6] Let us briefly outline the formal mechanism of partial-meet contraction. Let Γ be a set of sentences of $\mathcal{L}$, representing a belief corpus (either belief base or belief set), and let φ be a sentence of $\mathcal{L}$ that we would like to contract from Γ. Denote by $\Gamma \perp \varphi$ the set of $\subseteq$-maximal subsets of Γ that do *not* entail φ. A *selection function* γ is a function such that, for every set Γ of sentences and all sentences $\varphi \in \mathcal{L}$, if $\Gamma \perp \varphi \neq \varnothing$, then $\gamma(\Gamma \perp \varphi)$ is a non-empty subset of $\Gamma \perp \varphi$, whereas, if $\Gamma \perp \varphi = \varnothing$, then $\gamma(\Gamma \perp \varphi) = \{\Gamma\}$. On that basis, the selection function γ gives rise to an operator $\dot{-}_\gamma$ of partial-meet contraction, such that, for any set Γ of sentences and any sentence $\varphi \in \mathcal{L}$, $\Gamma \dot{-}_\gamma \varphi = \bigcap \gamma(\Gamma \perp \varphi)$. Thus, contrary to kernel contraction which is based on incision functions that select sentences that are relevant to derive the information to be retracted, partial-meet contraction is based on selection functions that select sets of sentences of a belief corpus that are eligible to be retained.

[7] The processes of kernel revision, smooth kernel revision and partial-meet revision are axiomatically characterized in [21, 54] and [29], respectively.

5 Related work

The favourable properties of belief bases, mentioned in the previous section, have made them a proper finite representation of belief corpora, in the context of Artificial Intelligence. It is, therefore, expected that several aspects of belief bases have extensively been discussed in the literature. In this section, and before proceeding to our contribution, we briefly overview some cornerstone works on belief bases.

We begin by stating that several important types of change-operations for belief bases, along with notable properties of them, are presented in Hansson's book [33], as well as in the surveys [16, Chapter 6] and [49, Section 8.4]. Now, an interesting type of change-operations on belief bases, which are implemented by means of a total preorder over a belief base, that essentially *prioritizes* its elements, is presented by Nebel [45]. Nebel's approach was generalized by Weydert [57], who also associated it to the AGM revision and contraction postulates. In a subsequent work [46], Nebel explored the *computational complexity* of change-operations on belief bases.

Non-prioritized change-operations, namely, operations that do *not* always accept the epistemic input, have also been adapted to the belief-base context. For example, [17] and [22] recast the model of *credibility-limited* revision, initially proposed for belief sets [35], in the realm of belief bases. In a more recent work [20], Garapa adapted *selective revision*, proposed for belief sets in [15], to the belief-base context, obtaining a model for revising belief bases that allows the acceptance of only part of the new information.

In a different vein, Fuhrmann in [18] pointed out that change-operations on a belief base B induce change-operations on its logical closure $K = Cn(B)$, which is of course a belief set. Hence, if $\dot{-}$ is a contraction function for B, a *base-generated* contraction function $\div$ can be defined, such that, for any sentence $\varphi \in \mathcal{L}$,

$$K \div \varphi = Cn\big(B \dot{-} \varphi\big).$$

Hansson in [30] and [31] obtained an axiomatic characterization for operations on belief sets, generated by partial-meet and kernel contraction of belief bases, respectively. In Section 8.5, we shall highlight some interesting connections between deductive belief change and base-generated change-operations. In his subsequent book [19], Fuhrmann explored *multiple* change-operations on belief bases; i.e., operations on belief bases in the context of which the epistemic input is a *set* of sentences, instead of a single sentence. Falappa et al. also covered this issue in [14], by studying a prioritized and a non-prioritized approach to multiple change applicable to belief bases.

We close our overview with the approach of Di Giusto and Governatori [26], in the realm of which the sentences of a belief base are partitioned into two classes; *facts*, which can be removed if new facts is necessary to be accommodated, and *rules*, which cannot be removed, but can instead be changed. Clearly, Di Giusto and Governatori's proposal prescribe guidelines on the vulnerability of sentences of belief bases. Deductive belief change, to be introduced in the next section, also imposes directives on the vulnerability of the sentences of a belief base, through the deductive justification that these sentences have within the belief base.

6 Deductive belief change

Deductive belief change (DBC) is a type of contraction and/or revision that is implemented by *particular* change-operations on *belief bases*, namely, change-operations that *respect*

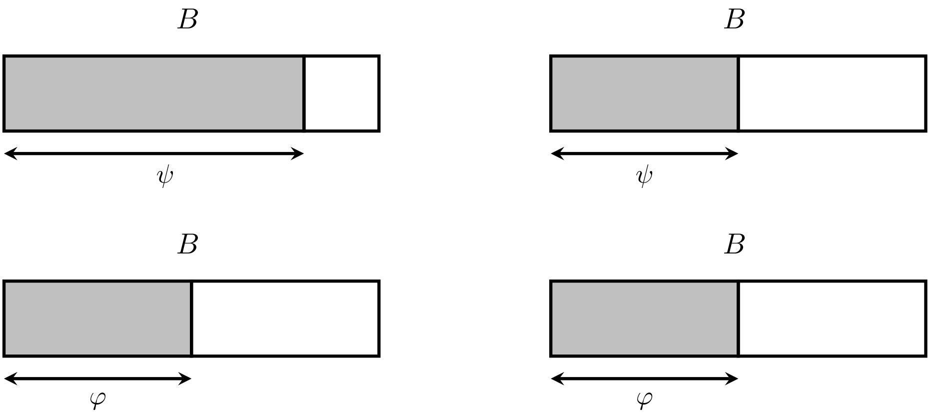

Fig. 1 Abstract representation of all the explicit beliefs (grey areas) of a belief base B that imply φ and ψ, where $\varphi \lhd_B \psi$ (left) and $\varphi \sim_B \psi$ (right)

specific constraints. This section is devoted to the presentation of the *axiomatic* and *constructive* side of these constraints, which, as we shall see, highlight the simple and intuitive structure of the introduced type of belief change. To that end, we first introduce the notion of *degree of support* of sentences, in the context of a belief base.[8]

Definition 4 (Unequal Degree of Support) Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$. We shall say that ψ is better supported than φ in B, which we shall denote by $\varphi \lhd_B \psi$, iff, for any sentence $\mu \in B$, $\mu \models \varphi$ entails $\mu \models \psi$, and there is a sentence $\nu \in B$ such that $\nu \not\models \varphi$ and $\nu \models \psi$.

Intuitively, ψ is better supported than φ in a belief base B whenever the explicit beliefs of B that (deductively) justify ψ are *set-theoretically more* than the explicit beliefs of B that (deductively) justify φ. Thus, if one remove enough "links" (deductions) to disconnect ψ from B, then φ gets also disconnected (cf. Figure 1, left).

Definition 5 (Equal Degree of Support) Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$. We shall say that φ and ψ are equally supported in B, which we shall denote by $\varphi \sim_B \psi$, iff, for any sentence $\mu \in B$, $\mu \models \varphi$ iff $\mu \models \psi$.

In a similar vein, φ and ψ are equally supported in a belief base B whenever the explicit beliefs of B that (deductively) justify φ are identical to the explicit beliefs of B that (deductively) justify ψ. Thus, φ gets disconnected from B iff ψ gets disconnected from B (cf. Figure 1, right). Evidently then, the degree of support that a sentence φ has in a belief base B is specified by the *explicit* beliefs of B that (deductively) make φ a belief of B.

Remark 1 Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$. It follows from Definition 4 that, if $B \not\models \varphi$ and $B \models \psi$, then $\varphi \lhd_B \psi$. Furthermore, in view of Definition 5, we have that, if $B \not\models \varphi$ and $B \not\models \psi$, then $\varphi \sim_B \psi$. Therefore, all the non-beliefs of B are

[8] The notion of degree of support presented herein is inspired by [7].

equally supported in B, and, at the same time, they are the least supported sentences of $\mathcal{L}$, with respect to B.

The following concrete example illustrates the notion of degree of support.

Example 2 (Degree of Support) Let $\mathcal{P} = \{a, b, c, d, e, f\}$ and let $B = \{a \leftrightarrow b, c \leftrightarrow d, e, e \rightarrow f\}$ be a belief base of a rational agent. Moreover, let $\varphi = \neg a \vee b \vee c \vee d$, $\chi = \neg a \vee b \vee \neg c \vee \neg d$ and $\psi = \neg a \vee b \vee \neg c \vee d$ be three sentences of $\mathcal{L}$, which are all implicit beliefs of B. Firstly, observe that, for any sentence $\mu \in B$, $\mu \models \varphi$ entails $\mu \models \psi$, and there is a sentence $\nu \in B$ (i.e., $c \leftrightarrow d$) such that $\nu \not\models \varphi$ and $\nu \models \psi$. Furthermore, for any sentence $\mu \in B$, $\mu \models \chi$ entails $\mu \models \psi$, and there is a sentence $\nu \in B$ (i.e., $c \leftrightarrow d$) such that $\nu \not\models \chi$ and $\nu \models \psi$. Therefore, in view of Definition 4, we derive that the belief ψ is better supported than the belief φ and χ in B; that is, $\varphi \lhd_B \psi$ and $\chi \lhd_B \psi$. Thereafter, observe that, for any sentence $\mu \in B$, $\mu \models \varphi$ iff $\mu \models \chi$. Thus, in view of Definition 5, it follows that the beliefs φ and χ are equally supported in B; that is, $\varphi \sim_B \chi$.

6.1 Axiomatic characterization

In view of the notion of degree of support, let $\dot{-}$ be an arbitrary contraction function, and consider the following two postulates (DC1) & (DC2), which reflect *specific properties* of $\dot{-}$.

(DC1) If $\varphi \lhd_B \psi$, then $B \dot{-} (\varphi \wedge \psi) \models \psi$.

(DC2) If $\varphi \sim_B \psi$, then $B \dot{-} (\varphi \wedge \psi) \not\models \varphi$.

Both postulates (DC1) & (DC2) relate the degree of support of two arbitrary sentences in a belief base B —which reflects the justificatory structure of B— with their resistance to change (i.e., vulnerability). The idea underlying postulates (DC1) & (DC2) is that, when we contract the belief $\varphi \wedge \psi$ from a belief base B, we are forced to give up *at least one* of the beliefs φ and ψ. Bearing that in mind, postulate (DC1) states that, if ψ is better supported than φ in B, then the belief ψ should be retained (and of course the belief φ should be withdrawn). Postulate (DC2) deals with the limiting case of two sentences φ and ψ that are equally supported in B. For such sentences, (DC2) requires that they should be *both contracted* from the belief base B.[9] We shall say that any contraction function that satisfies postulates (DC1) & (DC2) *implements deductive belief contraction*.

Obviously, a contraction function $\dot{-}$ that implements deductive belief contraction induces, through condition (LVB), a corresponding revision function $*$. Theorem 2, presented subsequently, proves that $\dot{-}$ satisfies postulates (DC1) & (DC2) iff $*$ satisfies the following two postulates (DR1) & (DR2), respectively.

(DR1) If $\varphi \lhd_B \psi$, then $B * (\neg\varphi \vee \neg\psi) \models \psi$.

(DR2) If $\varphi \sim_B \psi$, then $B * (\neg\varphi \vee \neg\psi) \not\models \varphi$.

Theorem 2 *Let $\dot{-}$ be a contraction function, and let $*$ be the corresponding revision function induced from $\dot{-}$, by means of (LVB). Then, $\dot{-}$ satisfies postulates (DC1) & (DC2) iff $*$ satisfies postulates (DR1) & (DR2), respectively.*

[9] Given that we confine our study to the principal case of contingent epistemic inputs, postulate (DC2) implies postulate ($B \dot{-} 1$), which any kernel contraction function respects (cf. Theorem 1); to see that, simply replace ψ with φ in (DC2).

Proof Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$. It suffices to prove that $B \dot{-} (\varphi \wedge \psi) \models \psi$ iff $B * (\neg\varphi \vee \neg\psi) \models \psi$, as well as that $B \dot{-} (\varphi \wedge \psi) \not\models \varphi$ iff $B * (\neg\varphi \vee \neg\psi) \not\models \varphi$.

The left-to-right implication of $B \dot{-} (\varphi \wedge \psi) \models \psi$ iff $B * (\neg\varphi \vee \neg\psi) \models \psi$ follows directly from condition (LVB). The right-to-left implication of $B \dot{-} (\varphi \wedge \psi) \models \psi$ iff $B * (\neg\varphi \vee \neg\psi) \models \psi$ follows from the deduction-property of the operator Cn (cf. Section 2). In particular, assume that $B * (\neg\varphi \vee \neg\psi) \models \psi$. Then, from the deduction-property of Cn and condition (LVB), it follows that $B \dot{-} (\varphi \wedge \psi) \models (\neg\varphi \vee \neg\psi) \rightarrow \psi$. Observe, however, that the sentence $(\neg\varphi \vee \neg\psi) \rightarrow \psi$ is logically equivalent to $(\varphi \wedge \psi) \vee \psi$, which is in turn logically equivalent to ψ. Therefore, we obtain that $B \dot{-} (\varphi \wedge \psi) \models \psi$, as desired.

The left-to-right implication of $B \dot{-} (\varphi \wedge \psi) \not\models \varphi$ iff $B * (\neg\varphi \vee \neg\psi) \not\models \varphi$ follows, like above, from the deduction-property of the operator Cn. In particular, assume that $B \dot{-} (\varphi \wedge \psi) \not\models \varphi$, and suppose towards contradiction that $B * (\neg\varphi \vee \neg\psi) \models \varphi$. Then, from the deduction-property of Cn and condition (LVB), it follows that $B \dot{-} (\varphi \wedge \psi) \models (\neg\varphi \vee \neg\psi) \rightarrow \varphi$. Hence, $B \dot{-} (\varphi \wedge \psi) \models \varphi$, a conclusion that contradicts our initial assumption. For the right-to-left implication of $B \dot{-} (\varphi \wedge \psi) \not\models \varphi$ iff $B * (\neg\varphi \vee \neg\psi) \not\models \varphi$, assume that $B * (\neg\varphi \vee \neg\psi) \not\models \varphi$. Suppose towards contradiction that $B \dot{-} (\varphi \wedge \psi) \models \varphi$. Then, we derive from condition (LVB) that $B * (\neg\varphi \vee \neg\psi) \models \varphi$, which is a contradiction once again. □

Like in the case of contraction, we shall say that any revision function that satisfies postulates (DR1) & (DR2) *implements deductive belief revision*.

6.2 Constructive characterization

We now turn to the *constructive* characterization of DBC in terms of *kernel belief change*. To that end, consider the following two conditions (DK1) & (DK2), which reflect specific properties of an incision function σ.

(DK1) If $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$, then there is a $X \in B \perp\!\!\!\perp \psi$ such that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X = \varnothing$.

(DK2) If $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$, then, for every $X \in B \perp\!\!\!\perp \varphi$, $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X \neq \varnothing$.

Condition (DK1) states that, if the ψ-kernels of B are *set-theoretically more* than the φ-kernels of B, then there is a ψ-kernel of B which is *disjoint* from the set $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big)$. Condition (DK2), on the other hand, deals with the case where the φ-kernels of B are identical to the ψ-kernels of B, in which circumstance (DK2) requires that the set $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big)$ intersects every φ-kernel of B (and of course every ψ-kernel of B as well).

Subsequently, we shall establish a correspondence between the contraction-postulates (DC1) & (DC2) and conditions (DK1) & (DK2); thus, in view of Theorem 2, a correspondence between the revision-postulates (DR1) & (DR2) and (DK1) & (DK2) shall be also established. Let us first, however, present the following lemma, which points out a natural connection between the degree of support of two sentences φ, ψ, in the context of a belief base B, and the φ-kernels and ψ-kernels of B.

Lemma 1 *Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$. Then, $\varphi \lhd_B \psi$ iff $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$, and $\varphi \sim_B \psi$ iff $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$.*

Proof Obvious from Definition 2 of Section 4, and Definitions 4 and 5 of the present section. □

Against this background, Theorem 3 is a representation result that establishes the alluded correspondence between the contraction-postulates (DC1) & (DC2) and conditions (DK1) & (DK2), and thus, it constructively characterizes DBC in terms of kernel belief change.

Theorem 3 *Let σ be an incision function, and let $\dot{-}_\sigma$ be a kernel contraction function induced from σ, by means of (K$\dot{-}$). Then, $\dot{-}_\sigma$ satisfies postulates (DC1) & (DC2) iff σ satisfies conditions (DK1) & (DK2), respectively.*

Proof For the left-to-right implication, assume first that $\dot{-}_\sigma$ satisfies postulate (DC1). We show that σ satisfies condition (DK1). Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$ such that $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$ (thus, $B \models \psi$). Then, it follows from Lemma 1 that $\varphi \lhd_B \psi$. Therefore, from postulate (DC1), we have that $B \dot{-}_\sigma (\varphi \wedge \psi) \models \psi$. Hence, we conclude from condition (K$\dot{-}$) that the set $B - \sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big)$ contains a ψ-kernel. Consequently, there is a $X \in B \perp\!\!\!\perp \psi$ such that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X = \varnothing$, as desired.

Next, assume that $\dot{-}_\sigma$ satisfies postulate (DC2). We show that σ satisfies condition (DK2). Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$ such that $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$. If $B \not\models \varphi$, then $B \perp\!\!\!\perp \varphi$ contains only the empty set, and thus (DK2) trivially holds. Assume, therefore, that $B \models \varphi$. Then, from $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$ and Lemma 1, we have that $\varphi \sim_B \psi$. Therefore, from postulate (DC2), we obtain that $B \dot{-}_\sigma (\varphi \wedge \psi) \not\models \varphi$. Hence, we conclude from condition (K$\dot{-}$) that the set $B - \sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big)$ does *not* contain a φ-kernel. Consequently, it follows from condition (K$\dot{-}$) that, for every $X \in B \perp\!\!\!\perp \varphi$, $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X \neq \varnothing$, as desired.

For the right-to-left implication, assume first that σ satisfies condition (DK1). We show that $\dot{-}_\sigma$ satisfies postulate (DC1). Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$ such that $\varphi \lhd_B \psi$. Then, it follows from Lemma 1 that $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$. Therefore, from condition (DK1), we have that there is a $X \in B \perp\!\!\!\perp \psi$ such that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X = \varnothing$. This again entails from condition (K$\dot{-}$) that $B \dot{-}_\sigma (\varphi \wedge \psi) \models \psi$, as desired.

Next, assume that σ satisfies condition (DK2). We show that $\dot{-}_\sigma$ satisfies postulate (DC2). Let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$ such that $\varphi \sim_B \psi$. Then, it follows from Lemma 1 that $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$. Therefore, from condition (DK2), we have that, for every $X \in B \perp\!\!\!\perp \varphi$, $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X \neq \varnothing$. This again entails from condition (K$\dot{-}$) that $B \dot{-}_\sigma (\varphi \wedge \psi) \not\models \varphi$, as desired. □

Figure 2 presents a visual representation of the conjunction of Theorems 2 and 3.

7 Kernel belief change implements deductive belief change

Having axiomatically and constructively introduced DBC in terms of natural constraints on the behaviour of change-operations on belief bases, we shall prove in this section that the well-established kernel belief change implements DBC and, thus, it provides a *proof of concept* for the introduced type of belief change. This result is established by Theorem 4 below, which proves that every kernel-based change-operation implements DBC.

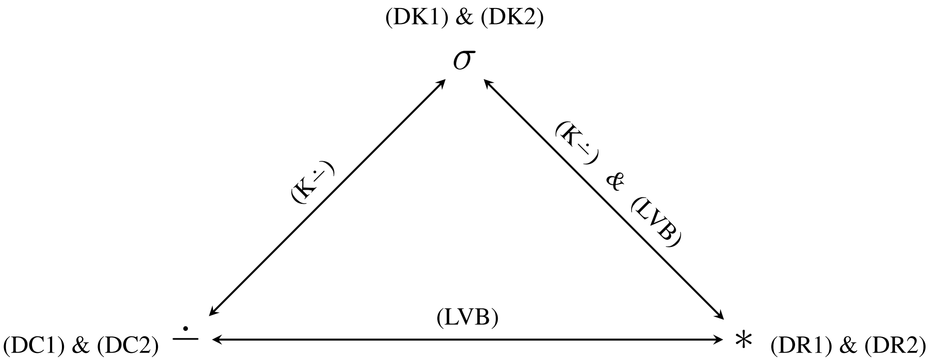

Fig. 2 An incision function σ satisfying conditions (DK1) & (DK2) induces, via condition (K$\dot{-}$), a kernel contraction function $\dot{-}$ that satisfies (DC1) & (DC2). The kernel contraction function $\dot{-}$ induces in turn, via condition (LVB), a kernel revision function $*$ that satisfies (DR1) & (DR2)

Theorem 4 *Let σ be an incision function, and let $\dot{-}_\sigma$ be a kernel contraction function induced from σ, by means of (K$\dot{-}$). Moreover, let $*_\sigma$ be the kernel revision function corresponding to $\dot{-}$, by means of (LVB). Then, $\dot{-}_\sigma$ satisfies postulates (DC1) & (DC2), and $*_\sigma$ satisfies postulates (DR1) & (DR2).*

Proof In view of Theorems 2 and 3, it suffices to prove that the incision function σ satisfies conditions (DK1) & (DK2). To that end, let B be a belief base, and let φ, ψ be two sentences of $\mathcal{L}$.

For (DK1), assume that $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$. It follows then that $B \perp\!\!\!\perp (\varphi \wedge \psi) = B \perp\!\!\!\perp \varphi$. Since, by definition, it is true that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \subseteq \bigcup\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big)$, we derive that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \subseteq \bigcup\big(B \perp\!\!\!\perp \varphi\big)$. This, combined with $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$, entails that there is a $X \in B \perp\!\!\!\perp \psi$ such that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X = \varnothing$, as desired.

For (DK2), assume that $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$. It follows then that $B \perp\!\!\!\perp (\varphi \wedge \psi) = B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \psi$. Since, by definition, the set $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big)$ intersects every element of $B \perp\!\!\!\perp (\varphi \wedge \psi)$, we derive that it also intersects every element of $B \perp\!\!\!\perp \varphi$. Consequently, we conclude that, for every $X \in B \perp\!\!\!\perp \varphi$, $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) \cap X \neq \varnothing$, as desired. □

Theorem 4 attributes an important property on kernel belief change, as it shows that this fundamental type of change-operation *already* encodes a strong connection between the justificatory structure of a belief base and the vulnerability of its beliefs. Furthermore, Theorem 4, in conjunction with Proposition 1 of Section 4, implies the subsequent corollary, which indicates that smooth kernel belief change and partial-meet belief change implement DBC as well.

Corollary 1 *Smooth kernel contraction functions and partial-meet contraction functions satisfy postulates (DC1) & (DC2). Moreover, smooth kernel revision functions and partial-meet revision functions, obtained through condition (LVB), satisfy postulates (DR1) & (DR2).*

Figure 3 summarizes the results established so far, by illustrating the classes of the principal change-operations on belief bases and their relation to DBC.

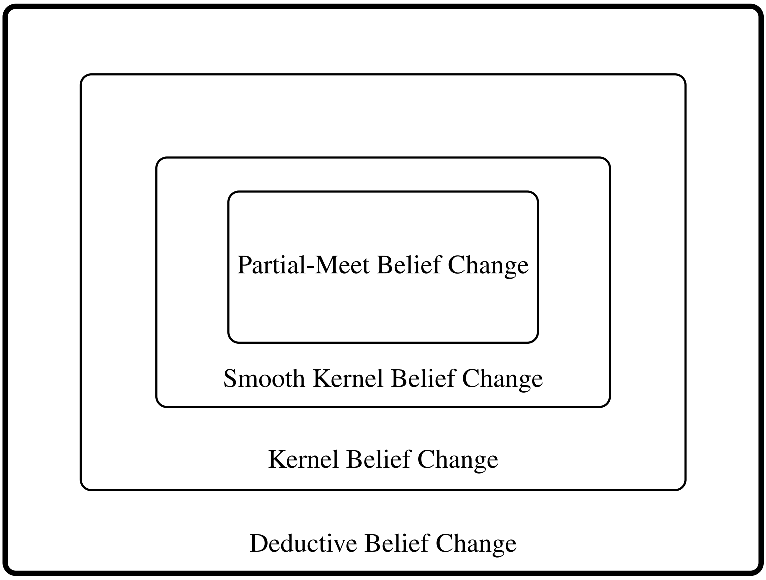

Fig. 3 The classes of the principal change-operations on belief bases and their relation to DBC

We conclude this section with Example 3, which demonstrates a concrete application of DBC (contraction and revision) implemented through kernel-based change-operations.

Example 3 (DBC via Kernel Belief Change, Cont'd Example 2) Recall that $B = \{a \leftrightarrow b, c \leftrightarrow d, e, e \rightarrow f\}$ is a belief base, and $\varphi = \neg a \vee b \vee c \vee d$, $\chi = \neg a \vee b \vee \neg c \vee \neg d$ and $\psi = \neg a \vee b \vee \neg c \vee d$ are three implicit beliefs of B. Observe that $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \chi = \{\{a \leftrightarrow b\}\}$, $B \perp\!\!\!\perp \psi = \{\{a \leftrightarrow b\}, \{c \leftrightarrow d\}\}$ and $B \perp\!\!\!\perp (\varphi \wedge \psi) = B \perp\!\!\!\perp (\varphi \wedge \chi) = \{\{a \leftrightarrow b\}\}$. Clearly, $B \perp\!\!\!\perp \varphi \subset B \perp\!\!\!\perp \psi$ and $B \perp\!\!\!\perp \varphi = B \perp\!\!\!\perp \chi$, which, in view of Lemma 1, gives that $\varphi \lhd_B \psi$ and $\varphi \sim_B \chi$, respectively. Now, let σ be an incision function. It follows then that $\sigma\big(B \perp\!\!\!\perp (\varphi \wedge \psi)\big) = \sigma\big(B \perp\!\!\!\perp (\varphi \wedge \chi)\big) = \{a \leftrightarrow b\}$, and obviously σ satisfies conditions (DK1) & (DK2), as expected due to Theorems 3 and 4. The incision function σ induces, through condition (K$\dot{-}$), a kernel contraction function $\dot{-}_\sigma$, such that $B \dot{-}_\sigma (\varphi \wedge \psi) = B \dot{-}_\sigma (\varphi \wedge \chi) = \{a \leftrightarrow b, c \leftrightarrow d, e, e \rightarrow f\} - \{a \leftrightarrow b\} = \{c \leftrightarrow d, e, e \rightarrow f\}$. The kernel contraction function $\dot{-}_\sigma$ induces in turn, through condition (LVB), a corresponding kernel revision function $*_\sigma$, such that $B *_\sigma (\neg\varphi \vee \neg\psi) = \Big\{c \leftrightarrow d, e, e \rightarrow f, \big(a \wedge \neg b \wedge \neg c \wedge \neg d\big) \vee \big(a \wedge \neg b \wedge c \wedge \neg d\big)\Big\}$ and $B *_\sigma (\neg\varphi \vee \neg\chi) = \Big\{c \leftrightarrow d, e, e \rightarrow f, \big(a \wedge \neg b \wedge \neg c \wedge \neg d\big) \vee \big(a \wedge \neg b \wedge c \wedge d\big)\Big\}$. Observe that $B \dot{-}_\sigma (\varphi \wedge \psi) \models \psi$ and $B \dot{-}_\sigma (\varphi \wedge \chi) \not\models \varphi$, according to the dictates of postulates (DC1) & (DC2), as well as that $B *_\sigma (\neg\varphi \vee \neg\psi) \models \psi$ and $B *_\sigma (\neg\varphi \vee \neg\chi) \not\models \varphi$, according to the dictates of postulates (DR1) & (DR2), as expected due to Theorem 4.

8 Deductive belief change in the realm of belief sets

So far, we have studied DBC on belief bases. This section is devoted to the translation of DBC in the realm of *belief sets*, which, contrary to belief bases, are logically closed (thus, infinite) sets of sentences. We first proceed to a *naïve* translation of DBC in the context of belief sets, and discuss its relation to AGM belief change (Section 8.1). Thereafter, we turn to a more interesting circumstance. We introduce a *relevance-sensitive* variant of the notion of degree of support (developed in Section 6), and, based on this variant, we reformulate postulates (DR1) & (DR2) in terms of belief sets. As we prove, in that case, deductive belief revision and Parikh's *relevance-sensitive* revision, [47, 48], are two different manifestations of the *same* operation (Section 8.3). Furthermore, we explore the relevance-sensitivity of kernel belief change on *belief bases* (Section 8.4). We shall close this section with a discussion on the relation between DBC and *base-generated* change-operations (Section 8.5).

8.1 A naïve translation

For a naïve translation of DBC in the realm of belief sets, one has to note that the definition of degree of support, as well as postulates (DC1) & (DC2) and (DR1) & (DR2) of Section 6, can all *straightforwardly* be formulated in terms of belief sets, simply be replacing B (denoting belief bases) with K (denoting belief sets).

On that premises, it follows immediately that *every* AGM contraction function *satisfies* the reformulated postulate (DC2). To see this, let $\dot{-}$ be an AGM contraction function, let K be a belief set, and let φ, ψ be two sentences of $\mathcal{L}$, such that $\varphi \sim_K \psi$. Firstly, observe that the case where one of φ, ψ is contained in K, and the other one is not contained in K, is excluded by the reformulated definition of $\sim_K$. Hence, either $\varphi, \psi \in K$ or $\varphi, \psi \notin K$. In the former case where $\varphi, \psi \in K$, it follows, by the reformulated definition of $\sim_K$, that φ and ψ are logically equivalent; that is, $Cn(\varphi) = Cn(\psi) = Cn(\varphi \wedge \psi)$. Therefore, postulates $(K \dot{-} 3)$ and $(K \dot{-} 6)$ entail that $\varphi, \psi \notin K \dot{-} (\varphi \wedge \psi)$. In the latter case where $\varphi, \psi \notin K$ (thus, $\varphi \wedge \psi \notin K$), postulate $(K \dot{-} 3)$ entails again that $\varphi, \psi \notin K \dot{-} (\varphi \wedge \psi)$. Consequently, in any case, the naïve recast of postulate (DC2) holds. Yet, contrary to (DC2), the naïve recast of postulate (DC1) is *not* respected by all AGM contraction functions, as one can easily find AGM contraction functions that violate it.

It is not hard to verify that a result analogous to Theorem 2 can be established, connecting the reformulated contraction-postulates (DC1) & (DC2) with the reformulated revision-postulates (DR1) & (DR2), via the Levi Identity (LV). By means of this result, and with a line of reasoning symmetric to the aforementioned one concerning contraction, we can obtain that, contrary to the naïve recast of (DR2), the naïve recast of (DR1) is respected by *every* AGM revision function.

Let us now turn to a more interesting case. As we shall prove in Section 8.3, a relevance-sensitive variant of the notion of degree of support suffices to make deductive belief revision *indistinguishable* from Parikh's relevance-sensitive revision. Parikh's relevance-sensitive revision is a fundamental type of revision that is more well-behaved than the one identified by the AGM revision postulates, which have been criticized as overly liberal towards relevance [47, 48]. Before further discussing Parikh's notion of relevance in the following subsection, it should be emphasized that relevance is a particularly important principle, both *conceptually* and *computationally*, that plays a cornerstone role, not only in the realm of belief change [4, 6, 24, 36, 40, 42, 44, 48, 51, 55, 56], but also in the context of Artificial Intelligence in general [27].

8.2 Parikh's notion of relevance

Parikh highlighted in [47] that there exist AGM revision functions —thus, AGM contraction functions as well generated by the Levi Identity— that are overly *liberal* towards relevance.

He showed, for example, that there is an AGM revision function $*$, such that, for any belief set K and any epistemic input φ that contradicts K, $K * \varphi = Cn(\varphi)$.[10] Observe that, during the $*$-revision of K by φ, *all* the K-beliefs that are not logical consequences of the epistemic input φ are thrown away, regardless of whether these beliefs are related to φ or not. In response to this unsatisfactory background, Parikh proposed an additional axiom that *supplements* the AGM revision postulates in handling relevance [47].

The main intuition of Parikh's axiom is that, if a belief set K is *splittable* —namely, it can be expressed in two *syntax-disjoint* compartments representing distinct subject matters—, then the revision of K by an epistemic input φ affects *only* the part of K that is syntactically relevant to φ. In a later work [51], Peppas et al. pointed out two different interpretations of Parikh's axiom, namely, is *weak* and its *strong* version, which are both plausible depending on the context. Herein, we shall focus on the weak version of Parikh's axiom, as it is more general and intuitive.

Before presenting the weak version of Parikh's axiom, let us firstly introduce the required notation and terminology. For a (proper) subset Q of $\mathcal{P}$, $\mathcal{L}^Q$ denotes the *sublanguage* of $\mathcal{L}$ defined over Q, using the standard Boolean connectives; if Q is empty, then $\mathcal{L}^Q$ is empty as well. For a sentence (resp., set of sentences) x of $\mathcal{L}$, $\mathcal{L}_x$ denotes the unique *minimal* language within which (resp., the sentences of) x can be expressed. The *complement* language of $\mathcal{L}_x$ is denoted by $\overline{\mathcal{L}_x}$; that is, $\overline{\mathcal{L}_x}$ is the language built from the atoms that do *not* appear in $\mathcal{L}_x$ (if there are no atoms that do not appear in $\mathcal{L}_x$, then $\overline{\mathcal{L}_x}$ is empty). Furthermore, consider the definitions of the *finest theory-splitting* by Parikh [47], and the *difference* between a theory and a possible world by Peppas et al. [51].

Definition 6 (Finest Theory-Splitting, [47]) Let K be a theory, and let $Q = \{Q_1, \ldots, Q_n\}$ be a partition of $\mathcal{P}$; i.e., $\bigcup Q = \mathcal{P}$, $Q_i \neq \varnothing$, and $Q_i \cap Q_j = \varnothing$, for all $1 \leqslant i \neq j \leqslant n$. The set Q is a K-splitting iff there exist sentences $\varphi_1 \in \mathcal{L}^{Q_1}, \ldots, \varphi_n \in \mathcal{L}^{Q_n}$, such that $K = Cn(\varphi_1, \ldots, \varphi_n)$. For each theory K, there exists a unique finest K-splitting (i.e., one which refines every other K-splitting), denoted by $\mathcal{F}_K$.[11]

Definition 7 (Difference between Theories and Possible Worlds, [51]) Let K be a theory with a finest K-splitting $\mathcal{F}_K$, and let r be a possible world of $\mathbb{M}$. The difference between K and r, denoted by Diff (K, r), is the union of the elements F of $\mathcal{F}_K$, for which there exists a sentence μ that can be expressed in the sublanguage $\mathcal{L}^F$, on that K and r disagree. In symbols:

$$\mathit{Diff}\,(K, r) = \bigcup \Big\{ F \in \mathcal{F}_K : \text{ for some } \mu \in \mathcal{L}^F, K \models \mu \text{ and } r \models \neg\mu \Big\}.$$

As it has been shown in [51], it is true that *Diff* $(K, r) = \varnothing$ iff $r \in [K]$. Lastly, in the spirit of Definition 7, we define the difference between two possible worlds r, r' of $\mathbb{M}$, denoted by *Diff* (r, r'), as the set of atoms over which r and r' disagree; in symbols,

$$\mathit{Diff}\,(r, r') = \Big((r - r') \cup (r' - r) \Big) \cap \mathcal{P}.$$

We are now ready to present the weak version of Parikh's axiom, which is encoded into the next postulate (RW). Postulate (RW) states that, if a belief set K splits between two

[10] The type of revision encoded into the AGM revision function $*$ is called *maxichoice* revision [1].

[11] A partition Q' refines another partition Q iff, for every $Q'_i \in Q'$, there exists a $Q_j \in Q$, such that $Q'_i \subseteq Q_j$.

syntax-disjoint compartments $Cn(x)$ and $Cn(y)$, then the revision of K by an epistemic input that can be asserted within the sublanguage $\mathcal{L}_x$ should *not* affect anything outside $\mathcal{L}_x$.

(RW) If $K = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$ and $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, then $(K * \varphi) \cap \overline{\mathcal{L}_x} = K \cap \overline{\mathcal{L}_x}$.

Peppas et al. in [51] formulated a semantic characterization of postulate (RW), which turns out to be the following two constraints (Q1) & (Q2) on a faithful preorder $\preceq_K$, associated with a belief set K.[12]

(Q1) If $\mathit{Diff}\ (K, r) \subset \mathit{Diff}\ (K, r')$ and $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$, then $r \prec_K r'$.

(Q2) If $\mathit{Diff}\ (K, r) = \mathit{Diff}\ (K, r')$ and $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$, then $r \approx_K r'$.

Both conditions (Q1) & (Q2) relate the plausibility of a possible world r to its difference $\mathit{Diff}\ (K, r)$ from the belief set K. In particular, condition (Q1) states that, if a world r differs from K in strictly fewer atoms than r', and moreover r and r' agree on the atoms in $\mathit{Diff}\ (K, r)$, then r should be strictly more plausible than r', with respect to K. Condition (Q2), on the other hand, asserts that, if two worlds r and r' differ from K on exactly the same atoms, and moreover both worlds agree on these atoms, then r and r' should be equally plausible, with respect to K.

Theorem 5 is a representation result by Peppas et al. that establishes the correspondence between postulate (PW) and the semantic conditions (Q1) & (Q2).

Theorem 5 [51] *Let $*$ be an AGM revision function, and let $\{\preceq_T\}_{\forall T}$ be the family of faithful preorders that corresponds to $*$, by means of (F$*$). Then, $*$ satisfies postulate (PW) iff $\{\preceq_T\}_{\forall T}$ satisfies conditions (Q1) & (Q2).*

8.3 Refined deductive belief revision and Parikh's relevance-sensitive revision: two sides of the same coin

Having discussed relevance-sensitive revision, we present the following *relevance-sensitive* variants of Definitions 4 and 5 in the realm of *belief sets*.

Definition 8 (Unequal Degree of Support for Belief Sets) Let K be a belief set with a finest K-splitting $\mathcal{F}_K$, and let φ, ψ be two sentences of $\mathcal{L}$. Denote by $\mathcal{S}_K^\varphi$ the set containing any element $F \in \mathcal{F}_K$, such that, for any sentence $\mu \in K \cap \mathcal{L}^F$, $\mu \models \varphi$ iff $\mu \models \psi$. We shall say that ψ is better supported than φ in K, denoted by $\varphi \lhd_K \psi$, iff there is an element $F' \in \mathcal{F}_K - \mathcal{S}_K^\varphi$, such that, for any sentence $\zeta \in K \cap \mathcal{L}^{F'}$, $\zeta \not\models \varphi$, and for a sentence $\nu \in K \cap \mathcal{L}^{F'}$, $\nu \models \psi$.

Definition 9 (Equal Degree of Support for Belief Sets) Let K be a belief set with a finest K-splitting $\mathcal{F}_K$, and let φ, ψ be two sentences of $\mathcal{L}$. We shall say that φ and ψ are equally supported in K, denoted by $\varphi \sim_K \psi$, iff, for each $F \in \mathcal{F}_K$ and any sentence $\mu \in K \cap \mathcal{L}^F$, $\mu \models \varphi$ iff $\mu \models \psi$.

Intuitively, ψ is better supported than φ in a belief set K whenever, every belief of every finest compartment of K, identified by $\mathcal{S}_K^\varphi$ (which is a *proper* subset of $\mathcal{F}_K$), entails φ iff

[12] For a concrete example of a faithful preorder that satisfies conditions (Q1) & (Q2), which also demonstrates the concepts of Definitions 6 and 7, the reader is referred to [51, Section 7]. We note, moreover, that a semantic characterization of Parikh's axiom in terms of a collection of popular constructive models for belief revision has been developed in [7].

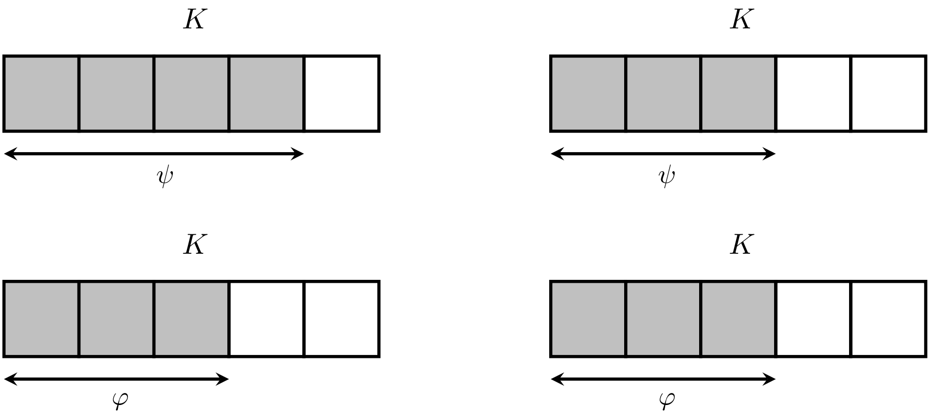

Fig. 4 Abstract representation of all the beliefs of the finest compartments (grey areas) of a belief set K that imply φ and ψ, where $\varphi \lhd_K \psi$ (left) and $\varphi \sim_K \psi$ (right)

it entails ψ, and, moreover, there is a finest compartment of K, which is *not* identified by $\mathcal{S}_K^{\varphi}$, whose beliefs do not entail φ, yet a belief of this compartment entails ψ. Thus, if one remove enough finest compartments of K to disconnect ψ from K, then φ gets also disconnected (cf. Figure 4, left). Analogously, φ and ψ are equally supported in a belief set K whenever, every belief of every finest compartment of K entails ψ iff it entails φ. Thus, the removal of enough finest compartments of K disconnects φ iff it disconnects ψ (cf. Figure 4, right). Hence, according to Definitions 8 and 9, the degree of support that a sentence φ has in a belief set K is specified by the finest compartments of K that (deductively) justify φ in K — in a sense, in the context of the degree of support for belief sets, (the beliefs of) the finest compartments of K take the role that the explicit beliefs have for belief bases.

Remark 2 As in the case of belief bases (cf. Remark 1), for a belief set K and two sentences φ, ψ of $\mathcal{L}$, it is true that, if $\varphi \notin K$ and $\psi \in K$, then $\varphi \lhd_K \psi$, and, moreover, if $\varphi, \psi \notin K$, then $\varphi \sim_K \psi$.

Based on Definitions 8 and 9, we can restate postulates (DR1) & (DR2) of Section 6 in terms of belief sets.[13] Accordingly, we end up to the following postulates (DRS1) & (DRS2), which essentially identify a *refined* version of deductive belief revision; in (DRS1) & (DRS2), K is a belief set, φ, ψ are arbitrary sentences of $\mathcal{L}$, and $*$ is a revision function for belief sets.

(DRS1) If $\varphi \lhd_K \psi$, then $\psi \in K * (\neg\varphi \vee \neg\psi)$.

(DRS2) If $\varphi \sim_K \psi$, then $\varphi \notin K * (\neg\varphi \vee \neg\psi)$.

Assuming that $*$ is an AGM revision function, Theorem 6 proves that the conjunction of postulates (DRS1) & (DRS2) is *equivalent* to postulate (PW), namely, to the weak version of Parikh's axiom.

[13]The focus here is on the revision-postulates (DR1) & (DR2), and not on the contraction postulates (DC1) & (DC2), since Parikh formulated his axiom in terms of revision.

Theorem 6 *Let $*$ be an AGM revision function. Then, $*$ satisfies postulates (DRS1) & (DRS2) iff $*$ satisfies postulate (PW).*

Proof Let $\{\preceq_T\}_{\forall T}$ be the family of faithful preorders that corresponds to $*$, by means of (F$*$). In the presence of Theorem 5, it suffices to show that $*$ satisfies postulates (DRS1) & (DRS2) iff $\{\preceq_T\}_{\forall T}$ satisfies conditions (Q1) & (Q2).

For the left-to-right implication, assume first that $*$ satisfies (DRS1). We show that $\{\preceq_T\}_{\forall T}$ satisfies (Q1). Let K be a belief set, and let r, r' be two possible worlds of $\mathbb{M}$ such that $\mathit{Diff}\ (K, r) \subset \mathit{Diff}\ (K, r')$ and $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$. If $\mathit{Diff}\ (K, r) = \varnothing$, then $r \in [K]$, and of course $r' \notin [K]$. Then, the faithfulness of $\preceq_K$ entails that $r \prec_K r'$, as desired. Assume, therefore, that $\mathit{Diff}\ (K, r) \neq \varnothing$.

Firstly, note that Definition 7 entails that the set $\big\{\mathit{Diff}\ (K, r)\ ,\ \mathcal{P}{-}\mathit{Diff}\ (K, r)\big\}$ is a K-splitting. Consequently, there exist sentences $x, y \in \mathcal{L}$ such that $K = Cn(x, y)$, $\mathcal{L}_x = \mathcal{L}^{Diff(K,r)}$, and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$. Now, construct the contingent sentences φ and ψ, such that $\varphi = \bigvee \overline{r}$ and $\psi = \bigvee \overline{r'}$.[14] Clearly by construction, it is true that $[\neg\varphi] = \{r\}$ and $[\neg\psi] = \{r'\}$. From $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$, it follows that, for any sentence $\mu \in K \cap \mathcal{L}^{Diff(K,r)} = Cn(x)$, $\mu \models \varphi$ iff $\mu \models \psi$. Furthermore, from $\mathit{Diff}\ (K, r) \subset \mathit{Diff}\ (K, r')$, it follows that there is a sentence $\nu \in K \cap \overline{\mathcal{L}^{Diff(K,r)}} = Cn(y)$ such that $\nu \models \psi$, and that there is *no* sentence $\zeta \in K \cap \overline{\mathcal{L}^{Diff(K,r)}} = Cn(y)$ such that $\zeta \models \varphi$. Combining the above, we derive that $\varphi \lhd_K \psi$, which, in view of postulates (DRS1) and $(K * 2)$, entails that $K * (\neg\varphi \vee \neg\psi) \models \psi$ and $K * (\neg\varphi \vee \neg\psi) \not\models \varphi$. It follows, then, that all the worlds in $\big[K * (\neg\varphi \vee \neg\psi)\big]$ are ψ-worlds and that there is a $\neg\varphi$-world in $\big[K * (\neg\varphi \vee \neg\psi)\big]$. Given that $[\neg\varphi] = \{r\}$ and $[\neg\psi] = \{r'\}$ (thus, $[\neg\varphi \vee \neg\psi] = \{r, r'\}$), we obtain that $\big[K * (\neg\varphi \vee \neg\psi)\big] = \{r\}$. This again entails from condition (F$*$) that $r \prec_K r'$, as desired.

Next, assume that $*$ satisfies (DRS2). We show that $\{\preceq_T\}_{\forall T}$ satisfies (Q2). Let K be a belief set, and let r, r' be two possible worlds of $\mathbb{M}$ such that $\mathit{Diff}\ (K, r) = \mathit{Diff}\ (K, r')$ and $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$. If $\mathit{Diff}\ (K, r) = \mathcal{P}$, then $r = r'$, and therefore (Q2) trivially holds. Moreover, if $\mathit{Diff}\ (K, r) = \varnothing$, then $r, r' \in [K]$, and the faithfulness of $\preceq_K$ entails that $r \approx_K r'$, as desired. Assume, therefore, that $\varnothing \subset \mathit{Diff}\ (K, r) \subset \mathcal{P}$.

Firstly, note again that the set $\big\{\mathit{Diff}\ (K, r)\ ,\ \mathcal{P}{-}\mathit{Diff}\ (K, r)\big\}$ is a K-splitting. Consequently, there exist sentences $x, y \in \mathcal{L}$ such that $K = Cn(x, y)$, $\mathcal{L}_x = \mathcal{L}^{Diff(K,r)}$, and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$. Now, construct the contingent sentences φ and ψ, such that $\varphi = \bigvee \overline{r}$ and $\psi = \bigvee \overline{r'}$. Clearly again, it is true that $[\neg\varphi] = \{r\}$ and $[\neg\psi] = \{r'\}$. From $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$, it follows that, for any sentence $\mu \in K \cap \mathcal{L}^{Diff(K,r)} = Cn(x)$, $\mu \models \varphi$ iff $\mu \models \psi$. Furthermore, from $\mathit{Diff}\ (K, r) = \mathit{Diff}\ (K, r')$, it follows that there is *no* sentence $\zeta \in K \cap \overline{\mathcal{L}^{Diff(K,r)}} = Cn(y)$ such that $\zeta \models \varphi$, and that there is *no* sentence $\xi \in K \cap \overline{\mathcal{L}^{Diff(K,r)}} = Cn(y)$ such that $\xi \models \psi$. Combining the above, we derive that $\varphi \sim_K \psi$, which, in view of postulate (DRS2), entails that $K * (\neg\varphi \vee \neg\psi) \not\models \varphi$ and $K * (\neg\varphi \vee \neg\psi) \not\models \psi$. It follows, then, that there is a $\neg\varphi$-world and a $\neg\psi$-world in $\big[K * (\neg\varphi \vee \neg\psi)\big]$. Given that $[\neg\varphi] = \{r\}$ and $[\neg\psi] = \{r'\}$ (thus, $[\neg\varphi \vee \neg\psi] = \{r, r'\}$), we obtain that $\big[K * (\neg\varphi \vee \neg\psi)\big] = \{r, r'\}$. This again entails from condition (F$*$) that $r \approx_K r'$, as desired.

We have shown so far that, if $*$ satisfies (DRS1) & (DRS2), then $\{\preceq_T\}_{\forall T}$ satisfies (Q1) & (Q2), respectively. We shall now prove the converse, namely that, if $\{\preceq_T\}_{\forall T}$ satisfies

[14] Recall that, for a set of literals Q, $\bigvee Q$ denotes the sentence of $\mathcal{L}$ resulting from the disjunction of all the literals in Q, and $\overline{Q}$ denotes the set of all the negated literals in Q.

(Q1) & (Q2), then $*$ satisfies (DRS1) & (DRS2), respectively. To that end, assume first that $\{\preceq_T\}_{\forall T}$ satisfies (Q1). We show that $*$ satisfies (DRS1). Let K be a belief set, and let φ, ψ be two sentences of $\mathcal{L}$, such that $\varphi \lhd_K \psi$.

Let $\mathcal{F}_K$ be the finest K-splitting. Since $\varphi \lhd_K \psi$, we have by definition that, for each $F \in \mathcal{S}_K^{\varphi}$ and any sentence $\mu \in K \cap \mathcal{L}^F$, $\mu \models \varphi$ iff $\mu \models \psi$. Hence, for each $F \in \mathcal{S}_K^{\varphi}$ and any sentence $\mu \in K \cap \mathcal{L}^F$, $\neg\varphi \models \neg\mu$ iff $\neg\psi \models \neg\mu$. Therefore, for each $F \in \mathcal{S}_K^{\varphi}$, any sentence $\mu \in K \cap \mathcal{L}^F$, any world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $r \models \neg\mu$ iff $r' \models \neg\mu$. Consequently, for any world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $\mathit{Diff}\ (K, r) \cap \left(\bigcup \mathcal{S}_K^{\varphi}\right) = \mathit{Diff}\ (K, r') \cap \left(\bigcup \mathcal{S}_K^{\varphi}\right)$ and $\mathit{Diff}\ (r, r') \cap \left(\mathit{Diff}\ (K, r) \cap \left(\bigcup \mathcal{S}_K^{\varphi}\right)\right) = \varnothing$. Furthermore, from the fact that there is an element $F' \in \mathcal{F}_K - \mathcal{S}_K^{\varphi}$, such that, for any sentence $\zeta \in K \cap \mathcal{L}^{F'}$, $\zeta \not\models \varphi$, and for a sentence $\nu \in K \cap \mathcal{L}^{F'}$, $\nu \models \psi$ (or $\neg\psi \models \neg\nu$), we derive that there is an element $F' \in \mathcal{F}_K - \mathcal{S}_K^{\varphi}$, such that, for any sentence $\zeta \in K \cap \mathcal{L}^{F'}$, some world $r \in [\neg\varphi]$, some sentence $\nu \in K \cap \mathcal{L}^{F'}$, and any world $r' \in [\neg\psi]$, $r \not\models \neg\zeta$ and $r' \models \neg\nu$. Combining the above, we conclude that, for some world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $\mathit{Diff}\ (K, r) \subset \mathit{Diff}\ (K, r')$ and $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$. Consequently, it follows from condition (Q1) that, for some world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $r \prec_K r'$. This, in view of condition (F$*$), entails that there is *no* $\neg\psi$-world in $\left[K * (\neg\varphi \vee \neg\psi)\right]$, which means that $K * (\neg\varphi \vee \neg\psi) \models \psi$, as desired.

Next, assume that $\{\preceq_T\}_{\forall T}$ satisfies (Q2). We show that $*$ satisfies (DRS2). Let K be a belief set, and let φ, ψ be two sentences of $\mathcal{L}$, such that $\varphi \sim_K \psi$.

Let $\mathcal{F}_K$ be the finest K-splitting. Since $\varphi \sim_K \psi$, we have by definition that, for each $F \in \mathcal{F}_K$ and any sentence $\mu \in K \cap \mathcal{L}^F$, $\mu \models \varphi$ iff $\mu \models \psi$. Hence, for each $F \in \mathcal{F}_K$ and any sentence $\mu \in K \cap \mathcal{L}^F$, $\neg\varphi \models \neg\mu$ iff $\neg\psi \models \neg\mu$. Therefore, for each $F \in \mathcal{F}_K$, any sentence $\mu \in K \cap \mathcal{L}^F$, any world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $r \models \neg\mu$ iff $r' \models \neg\mu$. Consequently, for any world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $\mathit{Diff}\ (K, r) = \mathit{Diff}\ (K, r')$ and $\mathit{Diff}\ (r, r') \cap \mathit{Diff}\ (K, r) = \varnothing$. This again entails from condition (Q2) that, for any world $r \in [\neg\varphi]$ and any world $r' \in [\neg\psi]$, $r \approx_K r'$. This, in view of condition (F$*$), entails that there is a $\neg\varphi$-world and a $\neg\psi$-world in $\left[K * (\neg\varphi \vee \neg\psi)\right]$, which means that $K * (\neg\varphi \vee \neg\psi) \not\models \varphi$, as desired. □

Theorem 6 shows that the class of AGM revision functions implementing the refined deductive belief revision (thus, satisfying postulates (DRS1) & (DRS2)) *coincides* with the class of AGM revision functions satisfying postulate (PW), namely, the weak version of Parikh's axiom. Therefore, the refined deductive belief revision and Parikh's relevance-sensitive revision are, as a matter of fact, *indistinguishable*. Theorem 6 provides, also, an alternative (equivalent) interpretation/characterization of Parikh's notion of relevance in terms of deductive-based structures; conversely, it provides an alternative (equivalent) interpretation/characterization of the refined deductive belief revision in terms of relevance-sensitive structures. Thus, responding to Hanssons' observations on the relation between justificatory structure and vulnerability, we show herein that the well-behaved Parikh's relevance-sensitive revision (for belief sets) is a type of belief change that *indeed relates* the aforementioned two concepts.

Example 4 demonstrates an instance of the refined version of deductive belief revision, which, as just shown, is equivalent to Parikh's relevance-sensitive revision.

Example 4 (Refined Deductive Belief Revision) Let $\mathcal{P} = \{a, b, c, d, e, f\}$, $x = a \leftrightarrow b$, $y = c \leftrightarrow d$, $z = e \wedge (e \rightarrow f)$, and let $K = Cn(x, y, z)$ be a belief set of a rational agent.[15] Clearly, the belief set K is splittable, and its finest K-splitting is $\mathcal{F}_K = \{\{a, b\}, \{c, d\}, \{e, f\}\}$. Moreover, let $*$ be an AGM revision function that satisfies postulates (DRS1) & (DRS2) (thus, postulate (PW) as well), and let $\varphi = \neg a \vee b \vee c \vee d$ and $\psi = \neg a \vee b \vee \neg c \vee d$ be two sentences of $\mathcal{L}$, which are all beliefs of K; i.e., $\varphi, \psi \in K$. Then, it follows from Definition 8 that $\varphi \lhd_K \psi$. Consider, now, the epistemic input $\omega = \neg\varphi \vee \neg\psi$ which contradicts K, and for which it is true that $\mathcal{L}_\omega \subseteq \mathcal{L}_{x \wedge y}$. On that premises, we obtain from postulate (DRS1) that $\psi \in K * \omega$. Therefore, from postulate $(K * 2)$, we have that $\neg\varphi \in K * \omega$. Since $*$ also satisfies postulate (PW), it holds that $(K * \omega) \cap \overline{\mathcal{L}_{x \wedge y}} = K \cap \overline{\mathcal{L}_{x \wedge y}}$. Combining the above, we derive that $K * \omega = Cn\Big(a \wedge \neg b \wedge \neg c \wedge \neg d\ ,\ e \wedge (e \rightarrow f)\Big)$.

8.4 Parikh's notion of relevance in the realm of belief bases

As stated at the beginning of Section 8.2, the well-established change-operations on belief sets —that is, those identified by the AGM trio— are *not* necessarily relevance-sensitive. But, what about the relevance-sensitivity of change-operations on *belief bases*? In this final subsection, we study this issue and prove that, contrary to the case of belief sets, the well-established operations on belief bases —that is, kernel contraction and revision— *respect syntax-relevance*. To that end, consider the following postulate (PB), which is a *straightforward* recast of postulate (PW) in terms of belief-bases contraction, and, as such, encodes Parikh's notion of relevance.

(PB) If $B = B_1 \cup B_2$, $\mathcal{L}_{B_1} \cap \mathcal{L}_{B_2} = \varnothing$ and $\mathcal{L}_\varphi \subseteq \mathcal{L}_{B_1}$, then $(B \dot{-} \varphi) \cap \overline{\mathcal{L}_{B_1}} = B \cap \overline{\mathcal{L}_{B_1}}$.

On that premise, Theorem 7 proves that kernel contraction respects postulate (PB), thus, it is relevance-sensitive.

Theorem 7 *Every kernel contraction function satisfies postulate (PB).*

Proof Let σ be an incision function, and let $\dot{-}_\sigma$ be the kernel contraction function induced from σ, by means of condition (K$\dot{-}$). Moreover, let B be a belief base such that $B = B_1 \cup B_2$ and $\mathcal{L}_{B_1} \cap \mathcal{L}_{B_2} = \varnothing$, and let φ be a sentence of $\mathcal{L}$ such that $\mathcal{L}_\varphi \subseteq \mathcal{L}_{B_1}$. Since, by definition, a φ-kernel of B is a $\subseteq$-minimal subset of B that entails φ, it follows that, for each element $X \in B \perp\!\!\!\perp \varphi$, $\mathcal{L}_X \subseteq \mathcal{L}_{B_1}$. This again entails that $\mathcal{L}_{\sigma(B \perp\!\!\!\perp \varphi)} \subseteq \mathcal{L}_{B_1}$. Consequently, we derive from condition (K$\dot{-}$) that *no* sentence of $B \cap \overline{\mathcal{L}_{B_1}}$ is affected by the $\dot{-}_\sigma$-contraction of φ from B. That is, $(B \dot{-}_\sigma \varphi) \cap \overline{\mathcal{L}_{B_1}} = B \cap \overline{\mathcal{L}_{B_1}}$, as desired. □

An important implication of Theorem 7, in view of Proposition 1 of Section 4, is the following corollary, which states that smooth kernel contraction and partial-meet contraction are relevance-sensitive as well.

Corollary 2 *Smooth kernel contraction functions and partial-meet contraction functions satisfy postulate (PB).*

Example 5 presents a kernel-contraction operation (on a belief base) that respects postulate (PB).

[15]Notice that the belief set K coincides with the logical closure $Cn(B)$ of the belief base B of Example 2 (Section 6).

Example 5 (Kernel Contraction & Relevance) Let $\mathcal{P} = \{a, b\}$ and let $B = \{a, b\}$ be a belief base. Clearly, B is the union of the belief bases $B_1 = \{a\}$ and $B_2 = \{b\}$. Moreover, let $\varphi = a$ be a sentence of $\mathcal{L}$, for which it is true that $\mathcal{L}_\varphi \subseteq \mathcal{L}_{B_1} = \mathcal{L}^{\{a\}}$. Then, we have that $B \perp\!\!\!\perp \varphi = \big\{\{a\}\big\}$. Notice that the minimal language of the single element of $B \perp\!\!\!\perp \varphi$ is the language $\mathcal{L}_\varphi = \mathcal{L}^{\{a\}}$. Thus, for any incision function σ, it holds that $\mathcal{L}_{\sigma(K \perp\!\!\!\perp \varphi)} \subseteq \mathcal{L}_{B_1}$. Therefore, the kernel contraction of φ from B respects postulate (PB), as expected due to Theorem 7.

We conclude this subsection noting that, in view of condition (LVB), a result analogous to Theorem 7 (and thus to Corollary 2) can straightforwardly be established for kernel revision as well.

8.5 Deductive belief change and base-generated change-operations

In the remainder of this section, we highlight some interesting connections between DBC and *base-generated* change-operations. As mentioned in Section 5, change-operations on a belief base B induce change-operations on its logical closure $K = Cn(B)$ [18]. Hence, if $\dot{-}$ is a contraction function for B, a base-generated contraction function $\div$ can be defined, such that, for any sentence $\varphi \in \mathcal{L}$,

$$K \div \varphi = Cn\big(B \dot{-} \varphi\big).$$

Then, it follows directly from conditions (DC1) & (DC2) that the contraction function $\div$ respects the following two properties, which encode the intuition of deductive belief contraction.

- If $\varphi \trianglelefteq_B \psi$, then $\psi \in K \div (\varphi \wedge \psi)$.
- If $\varphi \simeq_B \psi$, then $\varphi \notin K \div (\varphi \wedge \psi)$.

Notice that, in the above properties, the degree of support of sentences relative to the belief base B specifies the vulnerability of the sentences relative to the logical closure K of B. Analogous connections can be established between deductive belief revision and base-generated revision functions. Further relations between DBC and base-generated change-operations are left for future investigation.

9 Conclusion

Prompted by Sven Ove Hansson's observation that the relation between the justificatory structure of a belief corpus and the vulnerability of its beliefs remains an open issue [34], we introduced in the present article deductive belief change (DBC), a new well-behaved type of change-operation, identified by its simple and intuitive structure. DBC relates in a natural manner the deductive justification that the sentences of the language have, in the context of a belief base B, with their vulnerability relative to B. The core principle of DBC is that the more explicit B-beliefs imply a sentence φ, the more reluctant to change φ is, with respect to B.

The introduced type of belief revision was characterized both axiomatically (postulationally), as well as constructively, in terms of kernel belief change. Interestingly enough, we proved that DBC is implemented by kernel belief change, an important result pointing out that this fundamental kernel-based change-operation already encodes a strong coupling between justificatory structure and vulnerability. Immediate corollary of this outcome is that the well-established smooth kernel and partial-meet change-operations also implement DBC.

Beyond belief bases, DBC was studied in the realm of belief sets as well. As we proved, a proper adaptation of deductive belief revision for belief sets is indistinguishable from Parikh's relevance-sensitive revision, a central type of revision which, due to its favourable properties, constitutes a promising candidate for a variety of practical applications. Therefore, as shown, a refined version of deductive belief revision for belief sets and Parikh's relevance-sensitive revision are, as a matter of fact, two different manifestations of the same operation. Last but not least, we examined the notion of syntax-relevance in the context of belief bases, and showed that kernel belief change (on belief bases) is in fact relevance-sensitive, as it respects the counterpart of Parikh's axiom for belief bases.

As stated, DBC associates the deductive justification of sentences in a belief base B with their vulnerability relative to B. Of course, one can think of other properties of a belief base that could determine the vulnerability of its beliefs. For example, the vulnerability of the beliefs of B could be determined by the degree of inconsistency within B. Consider, for instance, the inconsistent belief base $B = \{a, b, \neg b\}$, where a and b are atoms of $\mathcal{P}$. One may argue that the belief a should be more resistant to withdrawal than b, since, contrary to a, the removal of b from B makes the logical closure of B consistent. By extending the intuition of the previous example, a type of belief-change could be defined in the context of which a sentence φ is more vulnerable than a sentence ψ, if the removal of φ from B makes B "more consistent" than the removal of ψ from B — measures of inconsistency of belief corpora have been widely discussed in the literature [37, 38, 41, 53]. Future investigation shall be devoted to this promising line of research.

Future work shall also be devoted to the exploration of potential interrelations of DBC with other notable types of change-operations, or frameworks such as the Truth Maintenance Systems [12] and the Argumentation Systems [52], in the context of which the justifications of beliefs play a cornerstone role (for example, in the explanations of the agents' actions).

Acknowledgements The author expresses his gratitude to Pavlos Peppas for his insightful remarks on this article, as well as to the anonymous referees for their invaluable input.

Data Availability Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Declarations

Conflict of Interests The author has no relevant financial or non-financial interests to disclose.

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. J. Symb. Log. **50**(2), 510–530 (1985)
2. Alchourrón, C., Makinson, D.: On the logic of theory change: Safe contractions. Stud. Logica. **44**, 405–422 (1985)
3. Aravanis, T.: On uniform belief revision. J. Log. Comput. **30**, 1357–1376 (2020)
4. Aravanis, T.: Generalizing Parikh's criterion for relevance-sensitive belief revision. ACM Trans. Comput. Logic **24**(2), 1–29 (2022). Article 18
5. Aravanis, T., Peppas, P.: Theory-relational belief revision. Ann. Math. Artif. Intell. **90**, 573–594 (2022)
6. Theofanis, I.: Aravanis. Relevance in belief update. J. Artif. Intell. Res. **72**, 251–283 (2021)
7. Aravanis, T.I., Peppas, P., Williams, M.-A.: Full characterization of Parikh's relevance-sensitive axiom for belief revision. J. Artif. Intell. Res. **66**, 765–792 (2019)

8. Areces, C., Becher, V.: Iterable AGM functions. In: Mary-Anne Williams and Hans Rott, editors, Frontiers in Belief Revision, volume 22 of Applied Logic Series, pp. 165–196. Springer (2001)
9. Dalal, M.: Investigations into theory of knowledge base revision: Preliminary report. In: Proceedings of the 7th National Conference of the American Association for Artificial Intelligence (AAAI 1988), pp. 475–479. The AAAI Press, Menlo Park (1988)
10. de Kleer, J.: An assumption-based TMS. Artificial Intelligence, **28**(127-162) (1986)
11. Dixon, S., Foo, N.: Connections between the ATMS and AGM belief revision. In: Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI 1993), pp. 534–539 (1993)
12. Doyle, J.: A truth maintenance system. Artif. Intell. **12**, 231–272 (1979)
13. Falappa, M.A., Fermé, E., Kern-Isberner, G.: On the logic of theory change: Relations between incision and selection functions. In: Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006), pp. 402–406 (2006)
14. Falappa, M.A., Kern-Isberner, G., Reis, M.D.L., Simari, G.R.: Prioritized and non-prioritized multiple change on belief bases. J. Philos. Log. **41**, 77–113 (2012)
15. Fermé, E., Hansson, S.O.: Selective revision. Stud. Logica. **63**, 331–342 (1999)
16. Fermé, E., Hansson, S.O.: Belief Change: Introduction and Overview. Springer, Berlin (2018)
17. Fermé, E., Mikalef, J., Taboada, J.: Credibility-limited functions for belief bases. J. Log. Comput. **13**, 99–110 (2003)
18. Fuhrmann, A.: Theory contraction through base contraction. J. Philos. Log. **20**(2), 175–203 (1991)
19. Fuhrmann, A.: An essay on contraction CSLI publications (1997)
20. Garapa, M.: Selective base revisions. J. Philos. Log. **51**, 1–26 (2022)
21. Garapa, M., Fermé, E., Reis, M.D.L.: Studies in credibility-limited base revision. In: Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018), pp. 240–247 (2018)
22. Garapa, M., Fermé, E., Reis, M.D.L.: Credibility-limited base revision: New classes and their characterizations. J. Artif. Intell. Res. **69**, 1023–1075 (2020)
23. Gärdenfors, P.: Knowledge in Flux – Modeling the Dynamics of Epistemic States. MIT Press, Cambridge (1988)
24. Gärdenfors, P.: Belief revision and relevance. In: PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, pp. 349–365 (1990)
25. Gärdenfors, P.: The dynamics of belief systems: Foundations versus coherence theories. Rev. Int. Philos. **44**, 24–46 (1990)
26. Di Giusto, P., Governatori, G.: A new approach to base revision. In: Barahona, P., Alferes, J. (eds.) Progress in Artificial Intelligence, volume 1695 of Lecture Notes in Computer Science, pp. 327–341. Springer (1999)
27. Greiner, R., Pearl, J., Subramanian, D.: (Eds.). Special issue on relevance. Artificial intelligence 97(1-2) (1997)
28. Grove, A.: Two modellings for theory change. J. Philos. Log. **17**(2), 157–170 (1988)
29. Hansson, S.O.: Belief Base Dynamics. PhD thesis, Uppsala University (1991)
30. Hansson, S.O.: Theory contraction and base contraction unified. The Journal of Symbolic Logic **58**(2), 602–625 (1993)
31. Hansson, S.O.: Kernel contraction. J. Symb. Log. **59**, 845–859 (1994)
32. Hansson, S.O.: Taking belief bases seriously. In: Prawitz, D., Westerståhl, D. (eds.) Logic and Philosophy of Science in Uppsala, volume 236 of Applied Logic Series, pp. 13–28. Springer, Berlin (1994)
33. Hansson, S.O.: A Textbook of Belief Dynamics: Theory change and database updating. Springer, Berlin (1999)
34. Hansson, S.O.: Ten philosophical problems in belief revision. J. Log. Comput. **13**, 37–49 (2003)
35. Hansson, S.O., Fermé, E., Cantwell, J., Falappa, M.A.: Credibility limited revision. J. Symb. Log. **66**(4), 1581–1596 (2001)
36. Hansson, S.O., Wassermann, R.: Local change. Studia Logica: An International Journal for Symbolic Logic **70**(1), 49–76 (2002)
37. Hunter, A., Konieczny, S.: Approaches to measuring inconsistent information. In: Bertossi, L., Hunter, A., Schaub, T. (eds.) Inconsistency Tolerance, volume 3300 of Lecture Notes in Computer Science, pp. 191–236. Springer (2005)
38. Hunter, A., Konieczny, S.: On the measure of conflicts: Shapley Inconsistency Values. Artif. Intell. **174**, 1007–1026 (2010)
39. Katsuno, H., Mendelzon, A.: Propositional knowledge base revision and minimal change. Artif. Intell. **52**(3), 263–294 (1991)
40. Kern-Isberner, G., Brewka, G.: Strong syntax splitting for iterated belief revision. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017), pp. 1131–1137 (2017)

41. Knight, K.: Measuring inconsistency. J. Philos. Log. **31**, 77–98 (2002)
42. Kourousias, G., Makinson, D.: Parallel interpolation, splitting, and relevance in belief change. J. Symb. Log. **72**(3), 994–1002 (2007)
43. Levi, I.: Subjunctives, dispositions and chances. Synthese **34**(4), 423–455 (1977)
44. Makinson, D.: Propositional relevance through letter-sharing. J. Appl. Log. **7**, 377–387 (2009)
45. Nebel, B.: Syntax-based approaches to belief revision. In: Gärdenfors, P. (ed.) Belief Revision, Theoretical Computer Science, pp. 52–88. Cambridge University Press (1992)
46. Nebel, B.: How hard is it to revise a belief base? In: Dubois, D., Prade, H. (eds.) Handbook of Defeasible Reasoning and Uncertainty Management Systems, volume 3 of Belief Change, pp. 77–145. Springer (1998)
47. Parikh, R.: Beliefs, belief revision, and splitting languages. In: Moss, L.S., Ginzburg, J., de Rijke, M. (eds.) Logic, Language and Computation, vol. 2, pp. 266–278. CSLI Publications (1999)
48. Parikh, R.: Beth definability, interpolation and language splitting. Synthese **179**, 211–221 (2011)
49. Peppas, P.: Belief revision. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) Handbook of Knowledge Representation, pp. 317–359. Elsevier Science (2008)
50. Peppas, P., Williams, M.-A.: Parametrised difference revision. In: Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018), pp. 277–286. The AAAI Press, Palo Alto (2018)
51. Peppas, P., Williams, M.-A., Chopra, S., Foo, N.: Relevance in belief revision. Artif. Intell. **229**, 126–138 (2015)
52. Simari, GR., Loui, R.P.: A mathematical treatment of defeasible reasoning and its implementation. Artificial Intelligence, **53**(125-157) (1992)
53. Thimm, M.: Inconsistency measurement. In: Amor, N.B., Quost, B., Theobald, M. (eds.) Scalable Uncertainty Management, Lecture Notes in Computer Science, pp. 9–23. Springer (2019)
54. Wassermann, R.: Resource bounded belief revision. Phd thesis University of Amsterdam (2000)
55. Wassermann, R.: Local diagnosis. J. Appl. Non-Classical Logics **11**(1–2), 107–129 (2001)
56. Wassermann, R.: On structured belief bases. In: Williams, M.-A., Rott, H. (eds.) Frontiers in Belief Revision, volume 22 of Applied Logic Series, pp. 349–367. Springer (2001)
57. Weydert, E.: Relevance and revision: About generalizing syntax-based belief revision. In: Pearce, D., Wagner, G. (eds.) Logics in AI, European Workshop, JELIA '92, volume 633 of Lecture Notes in Computer Science, pp. 126–138. Springer (1992)