Journal of Artificial Intelligence Research 72 (2021) 251-283 Submitted 02/2021; published 10/2021

Relevance in Belief Update

Theofanis I. Aravanis TARAVANIS@UPATRAS.GR
Department of Business Administration
School of Economics & Business
University of Patras
Patras 265 00, Greece

Abstract

It has been pointed out by Katsuno and Mendelzon that the so-called AGM revision operators, defined by Alchourrón, Gärdenfors and Makinson, do not behave well in dynamically-changing applications. On that premise, Katsuno and Mendelzon formally characterized a different type of belief-change operators, typically referred to as KM update operators, which, to this date, constitute a benchmark in belief update. In this article, we show that there exist KM update operators that yield the same counter-intuitive results as any AGM revision operator. Against this non-satisfactory background, we prove that a translation of Parikh's relevance-sensitive axiom (P), in the realm of belief update, suffices to block this liberal behaviour of KM update operators. It is shown, both axiomatically and semantically, that axiom (P) for belief update, essentially, encodes a type of relevance that acts at the possible-worlds level, in the context of which each possible world is locally modified, in the light of new information. Interestingly, relevance at the possible-worlds level is shown to be equivalent to a form of relevance that acts at the sentential level, by considering the building blocks of relevance to be the sentences of the language. Furthermore, we concretely demonstrate that Parikh's notion of relevance in belief update can be regarded as (at least a partial) solution to the frame, ramification and qualification problems, encountered in dynamically-changing worlds. Last but not least, a whole new class of well-behaved, relevance-sensitive KM update operators is introduced, which generalize Forbus' update operator and are perfectly-suited for real-world implementations.

1. Introduction

Belief revision (or, simply, revision) is the process by which a rational agent modifies her beliefs about a *static/unchanging* world, in the light of new information (Gärdenfors, 1988). A prominent approach that formalizes belief revision is that proposed by Alchourrón, Gärdenfors, and Makinson (1985), now known as the *AGM paradigm*, after the initials of its three founders. Within the AGM paradigm, the agent's belief corpus is represented by a logical theory of an underlying logic language, and a new piece of information is expressed by a logical sentence. A collection of postulates, called the *AGM postulates*, characterizes any rational revision operator, named *AGM revision operator*, which maps theories and sentences to theories.

In a cornerstone article, Katsuno and Mendelzon (1992) pinpointed an insufficiency of AGM revision operators in producing reasonable results in a real-world scenario, the so-called *book/magazine example*. They, in turn, argued that AGM revision operators were never meant to deal with situations analogous to the book/magazine example in the first place, and proposed as a solution to this problem a new type of belief change that is appropriate for *dynamically-changing* applications, known as *belief update*.[1]

Belief update (or, simply, update) refers to the type of belief change by which, in view of new information, a change *in the state of the world* —rather than in the *description/perception* of the state of the world— takes place. Following the AGM tradition, Katsuno and Mendelzon characterized belief update in terms of a set of rationality postulates, typically referred to as the *KM postulates*. The update operators that satisfy the KM postulates are called *KM update operators*, and, like the AGM revision operators, map theories and sentences to theories. Contrary to AGM revision operators, however, which treat theories as unified entities, KM update operators operate in a *pointwise* manner, by reducing the update of an arbitrary theory to the modification of *each* possible world that satisfies the theory.

Although KM update operators have been proven to be benchmark formal tools for belief update, we show, in this article, that they are liberal in their treatment of *relevance*. In response to this weakness, we propose a relevance-sensitive framework for belief update, which *strictly strengthens* Katsuno and Mendelzon's approach, based on Parikh's model for belief revision (1999). Specifically, the following contributions are made:

- We show that there are in fact KM update operators that yield the same counter-intuitive results as *any* AGM revision operator.

- Against this non-satisfactory background, we show that a translation of Parikh's relevance-sensitive axiom (P) in the realm of belief update suffices to block these undesirable behaviours of KM update operators. Parikh's axiom, originally formulated for capturing *local* revision, encodes the following intuitive principle: If a theory K is *splittable*, namely, it can be expressed in two *syntax-disjoint* compartments (representing independent subject matters), then the revision of K by a new piece of information φ affects only the part of K that is *syntactically relevant* to the *minimal language* of φ. Evidently, Parikh's axiom perceives relevance *at the theory level*, in the sense that it considers the building blocks of relevance to be the disjoint compartments of a splittable theory. Splittability is, in turn, a function of the *contingent* beliefs of an agent. Parikh's approach is often referred to as the *language-splitting model*, and, since its proposal, has been extensively studied; see, indicatively, the works by Chopra and Parikh (2000), Makinson and Kourousias (2006), Kourousias and Makinson (2007), Makinson (2009), Peppas, Williams, Chopra, and Foo (2015), and Aravanis, Peppas, and Williams (2019).

- Peppas et al. (2015) identified two different interpretations of Parikh's axiom (P), called the *weak* and the *strong* version of (P), both of which are plausible depending on the application. We prove that the weak version of axiom (P), when translated in the realm of belief update, *reduces* to a simple postulate, named (R1), which ensures that the modification of any possible world (representing a world state) should *not* affect the part of the world that is *irrelevant* to the new information φ. On the other hand, the strong version of (P) for belief update *strictly implies* (but it is not

1. Although the process of belief update was first noticed by Keller and Winslett (1985), it was elaborated and formalized by Katsuno and Mendelzon (1992).

equivalent to) a postulate, named (R2), which ensures that the modification of the φ-relevant part of possible worlds is *not* affected by their φ-irrelevant part; thus, (R2) makes the update-process *context-independent*. Both postulates (R1) and (R2) do *not* demand the splittability of theories, as they act directly *at the possible-worlds level* — this behaviour is aligned with the pointwise operation of KM postulates.

- We show that postulates (R1)–(R2) are, in turn, *equivalent* to postulates (S1)–(S2), respectively. Both (S1) and (S2) act *at the sentential level*, in the sense that they consider the building blocks of relevance to be the sentences (beliefs) of a theory.

- *Possible-worlds semantics* for all the relevance-sensitive postulates for belief update presented herein are provided, by formulating appropriate constraints on preorders over possible worlds which precisely characterize the postulates. We, also, discuss the resemblance between the semantic characterization of the weak version of axiom (P) for belief update and Winslett's *Possible Models Approach* (PMA) (1988), which constitutes a popular method suitable for *reasoning about action*.

- The ability to ignore unrelated information is a prerequisite for the implementation of real-world rational agents, equipped with high-level cognitive capabilities. Although this is, in general, a rather simple task for humans in their daily-life action, it turned out to be a challenging problem in formalizing dynamically-changing worlds; an appropriate formalization requires "too much" explicit information about environment dynamics. This challenge is formally described, to a large extent, in the three classical problems of (logic-based) Artificial Intelligence, namely, the *frame* (McCarthy & Hayes, 1969), *ramification* (Finger, 1987), and *qualification* (McCarthy, 1977) problems. As the main desideratum concerning these important problems is the distinction between the relevant and irrelevant fluents with respect to an action,[2] they are evidently in a close connection with relevance-sensitive belief update. Accordingly, we concretely show how a slight adjustment of the weak version of axiom (P) for belief update, which incorporates the notion of *causality*, can be regarded as (at least a partial) solution to these problems.

- A whole class of concrete KM update operators that respect *all* the presented relevance-sensitive postulates is introduced, both *axiomatically* and *semantically*; we call these new operators *parametrized-difference* (PD) *update operators*, or *PD update operators*, for short. PD update operators constitute a natural *generalization* of Forbus' update operator (1989), and are, essentially, a subsequent reformulation of parametrized-difference revision operators of Peppas and Williams (2016, 2018). We demonstrate that PD update operators are perfectly-suited for real-world implementations, as they can be compactly specified.

Although relevance in belief revision has been extensively studied —beyond the works mentioned earlier, see also (Gärdenfors, 1990; Hansson & Wassermann, 2002; Wassermann, 2001b, 2001a; Kern-Isberner & Brewka, 2017; Delgrande & Peppas, 2018)— the role that this important notion plays in the realm of belief update has not gained the analogous

2. A fluent is a condition of a world state that can change over time.

attention. Indeed, to the best of our knowledge, the only work that directly bears on the approach adopted herein is that by Perrussel, Marchi, Thévenin, and Zhang (2012). Nevertheless, that work departs from the classical framework of Katsuno and Mendelzon, as it considers belief update by handling *prime implicants*; thus, a significant deviation from the KM postulates, as well as their semantic characterization, was required. The present work, on the other hand, is *fully aligned* with the standard Katsuno and Mendelzon's approach, and essentially provides a *strict strengthening* of the KM postulates, by characterizing a *proper sub-class* of KM update operators —namely, the sub-class satisfying the proposed postulates— that respect relevant change.

The remainder of this article is structured as follows. The next section sets the formal background for our discussion. Thereafter, Sections 3 and 4 discuss belief revision and belief update, respectively, followed by Section 5, which introduces Parikh's notion of relevance. Sections 6 and 7 investigate (both axiomatically and semantically) the weak and the strong version of Parikh's axiom (P), respectively, in the realm of belief update. In Section 8, PD update operators are introduced. A brief concluding section closes the paper.

2. Formal Background

This section sets the formal background material that shall be used throughout this article.

Language. 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 $\models$.

Sentences. A sentence φ of $\mathcal{L}$ is *contingent* iff it is neither a tautology, nor a contradiction. For a set of sentences Γ of $\mathcal{L}$, $Cn(\Gamma)$ denotes the set of all logical consequences of Γ; i.e., $Cn(\Gamma) = \{\varphi \in \mathcal{L} : \Gamma \models \varphi\}$. For a sequence of sentences $x_1, \ldots, x_n$ of $\mathcal{L}$, we shall write $Cn(x_1, \ldots, x_n)$, as an abbreviation of $Cn(\{x_1, \ldots, x_n\})$. For a set of sentences Γ of $\mathcal{L}$, $\bigwedge \Gamma$ denotes the single sentence of $\mathcal{L}$ resulting from the conjunction of all sentences in Γ.

Theories. A theory (also referred to as *belief set*) K is any set of sentences of $\mathcal{L}$ closed under logical consequence; i.e., $K = Cn(K)$. A theory K is *complete* iff, for all sentences $\varphi \in \mathcal{L}$, either $\varphi \in K$ or $\neg\varphi \in K$. For a theory K and a sentence φ of $\mathcal{L}$, the *expansion* of K by φ, denoted by $K + \varphi$, is defined to be the deductive closure of the set $K \cup \{\varphi\}$; i.e.,

$$K + \varphi = Cn(K \cup \{\varphi\}).$$

Literals. A *literal* is an atom $p \in \mathcal{P}$ or its negation. For a set of literals Q, $|Q|$ denotes the cardinality of Q, whereas, $\overline{Q}$ denotes the set of all the negated literals in Q; i.e., $\overline{Q} = \{\neg q : q \in Q\}$.

Possible Worlds. A *possible world* (or, simply, *world*) r is defined to be a consistent set of literals, such that, for any atom $p \in \mathcal{P}$, either $p \in r$ or $\neg p \in r$. The set of all possible worlds is denoted by $\mathbb{M}$. By the definition of complete theories, it follows that there is a *one-to-one correspondence* between consistent complete theories and possible worlds, such that, for any consistent complete theory K, there exists a world $w \in \mathbb{M}$ such that $[K] = \{w\}$. For a sentence (set of sentences) φ of $\mathcal{L}$, $[\varphi]$ is the set of all worlds satisfying φ.

Sublanguages. Let Q be a subset of $\mathcal{P}$. We denote by $\mathcal{L}^Q$ the *sublanguage* of $\mathcal{L}$ defined over Q, using the standard Boolean connectives. For a sentence φ of $\mathcal{L}$, we denote by $\mathcal{P}_\varphi$ the (*unique*) *minimal* subset of $\mathcal{P}$, through which a sentence that is logically equivalent to φ can be formulated. If φ is not contingent, we take $\mathcal{P}_\varphi$ to be the empty set. We define the *minimal language* $\mathcal{L}_\varphi$ of φ to be the propositional (sub)language defined over $\mathcal{P}_\varphi$, using the standard Boolean connectives. We, also, define $\overline{\mathcal{L}_\varphi}$ to be the propositional (sub)language defined over $\mathcal{P}-\mathcal{P}_\varphi$, using the standard Boolean connectives. If $\mathcal{P}_\varphi$ (resp., $\mathcal{P}-\mathcal{P}_\varphi$) is empty, then $\mathcal{L}_\varphi$ (resp., $\overline{\mathcal{L}_\varphi}$) is empty. Lastly, for a possible world r of $\mathbb{M}$ and a contingent sentence φ of $\mathcal{L}$, r_φ denotes the *restriction* of r to the minimal language $\mathcal{L}_\varphi$; that is, $r_\varphi = r \cap \mathcal{L}_\varphi$.[3]

Preorders. A *partial preorder* over a set V is a reflexive, transitive binary relation in V. A partial preorder $\preceq$ is *total* iff, for all $r, r' \in V$, $r \preceq r'$ or $r' \preceq r$. The strict part of $\preceq$ is denoted by $\prec$; i.e., $r \prec r'$ iff $r \preceq r'$ and $r' \not\preceq r$. The symmetric part of $\preceq$ is denoted by $\approx$; i.e., $r \approx r'$ iff $r \preceq r'$ and $r' \preceq r$. Also, $min(V, \preceq)$ denotes the set of all minimal elements of V, with respect to $\preceq$; that is,

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

3. Belief Revision

In this section, we present the axiomatic and semantic characterization of the process of belief revision, along with a benchmark scenario of the belief-update literature that served as a motivation to this work.

3.1 Axiomatic Characterization

A *revision operator* is a function $*$ that maps a theory K and a sentence φ to a new theory $K * \varphi$, representing the result of revising K by φ. We shall say that a revision operator $*$ is an *AGM revision operator* iff it satisfies the following postulates, known as *AGM postulates*:

($\mathbf{K*1}$) $K * \varphi$ is a theory of $\mathcal{L}$.

($\mathbf{K*2}$) $\varphi \in K * \varphi$.

($\mathbf{K*3}$) $K * \varphi \subseteq K + \varphi$.

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

($\mathbf{K*5}$) $K * \varphi = \mathcal{L}$ iff φ is inconsistent.

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

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

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

The guiding principle behind the AGM postulates —whose rationale is discussed by Gärdenfors (1988, Section 3.3) and Peppas (2008, Section 8.3.1)— is that of *economy of*

3. Notation and terminology on sublanguages is borrowed from Parikh (1999), Peppas et al. (2015), and Makinson (2009).

information; namely, the new information φ must be consistently incorporated into the initial belief set K, changing it as little as possible. Note that, in the special case where φ is *consistent* with K, it follows from postulates $(K*3)$–$(K*4)$ that revision reduces to expansion; that is, $K * \varphi = K + \varphi$.

3.2 Semantic Characterization

Let us, now, proceed to the semantic characterization of belief revision. It turns out that the revision operators that satisfy postulates $(K*1)$–$(K*8)$ are precisely those that are induced by means of *total* preorders over all possible worlds.

Definition 1 (Faithful Preorder Associated with Theories, Katsuno & Mendelzon, 1991)**.** *A total preorder $\preceq_K$ over $\mathbb{M}$ is faithful to a theory K iff, for any $r, r' \in \mathbb{M}$, the following two conditions hold:*

- **(i)** *If $r \in [K]$, then $r \preceq_K r'$.*
- **(ii)** *If $r \in [K]$ and $r' \notin [K]$, then $r \prec_K r'$.*

Intuitively, $r \preceq_K r'$ precisely when the world r is at least as plausible (relative to K) as the world r'.

Definition 2 (Faithful Assignment, Katsuno & Mendelzon, 1991)**.** *A faithful assignment is a function that maps each theory K of $\mathcal{L}$ to a total preorder $\preceq_K$ over $\mathbb{M}$, that is faithful to K.*

The following theorem characterizes the class of revision operators induced by faithful assignments based on *total* preorders.

Theorem 1 (Katsuno & Mendelzon, 1991)**.** *A revision operator $*$ satisfies postulates $(K*1)$–$(K*8)$ iff there exists a faithful assignment that maps each theory K of $\mathcal{L}$ to a total preorder $\preceq_K$ over $\mathbb{M}$, such that, for any $\varphi \in \mathcal{L}$:*

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

3.3 The Book/Magazine Example

As pointed out by Katsuno and Mendelzon (1992), the type of belief change encoded in the AGM revision operators is inadequate to produce reasonable outcomes in dynamically-changing applications. The scenario they use to make their case is known as the *book/magazine example*, presented below.

Example 1 (The Book/Magazine Example, Katsuno & Mendelzon, 1992)**.** *Consider a room with a table, a magazine and a book. We look through an open door at the room, and see that either the magazine or the book is on the table, but not both. However, because of poor lighting, we cannot distinguish which is which. Supposing that atom b represents the fact that "the book is on the table", and that atom m represents the fact that "the magazine is on the table", the initial state of the world can be described by a theory K, such that:*

$$K = Cn\Big((b \wedge \neg m) \vee (\neg b \wedge m)\Big).$$

Now, we instruct a robot to place the book on the floor; that is, if the book is on the table, the robot will place it on the floor, otherwise, it will do nothing. After the (successful) action of the robot, the new state of the world will, presumably, be the initial theory K modified by the new information $\varphi = \neg b$.

Assuming that $\mathcal{L}$ contains only the atoms b and m, i.e., $\mathcal{P} = \{b, m\}$, the set $[K]$ contains exactly the two worlds $w_1 = \{b, \neg m\}$ and $w_2 = \{\neg b, m\}$, whereas, the set $[\varphi]$ contains exactly the two worlds w_2 and $r = \{\neg b, \neg m\}$. If we use an *arbitrary* AGM revision operator $*$ to modify theory K, we obtain that

$$K * \varphi = K + \varphi = Cn(\neg b \wedge m),$$

as the new information φ is *consistent* with K; that is to say, K is expanded by φ. This is clearly a counter-intuitive result, since, if the magazine was initially on the floor ($\neg m$), putting the book on the floor ($\neg b$), somehow the magazine jumps onto the table (m).

4. Belief Update

In this section, an overview of the process of belief update is presented, as defined by Katsuno and Mendelzon (1992), along with a concrete KM update operator, introduced by Winslett (1988). The liberal behaviour of KM postulates in the context of the book/magazine example is formally pointed out as well.

4.1 Axiomatic Characterization

An *update operator* is a function $\diamond$ that maps a theory K and a sentence φ to a new theory $K \diamond \varphi$, representing the result of updating K by φ. We shall say that an update operator $\diamond$ is a *KM update operator* iff it satisfies the following postulates, known as *KM postulates*:

(**K$\diamond$1**) $K \diamond \varphi$ is a theory of $\mathcal{L}$.

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

(**K$\diamond$3**) If $\varphi \in K$, then $K \diamond \varphi = K$.

(**K$\diamond$4**) $K \diamond \varphi = \mathcal{L}$ iff K or φ is inconsistent.

(**K$\diamond$5**) If $Cn(\varphi) = Cn(\psi)$, then $K \diamond \varphi = K \diamond \psi$.

(**K$\diamond$6**) $K \diamond (\varphi \wedge \psi) \subseteq (K \diamond \varphi) + \psi$.

(**K$\diamond$7**) If $\psi \in K \diamond \varphi$ and $\varphi \in K \diamond \psi$, then $K \diamond \varphi = K \diamond \psi$.

(**K$\diamond$8**) If K is complete, then $K \diamond (\varphi \vee \psi) \subseteq Cn\big((K \diamond \varphi) \cup (K \diamond \psi)\big)$.

(**K$\diamond$9**) If $[K] \neq \varnothing$, then $K \diamond \varphi = \bigcap_{w \in [K]} Cn(w) \diamond \varphi$.

For ease of comparison, the KM postulates have been rephrased in the AGM notation. That is to say, postulates $(K \diamond 1)$–$(K \diamond 9)$ are equivalent to postulates (U1)–(U8) introduced by Katsuno and Mendelzon (1992), when sentences are replaced by theories in the representation of states of belief — details on this transformation can be found in the works of Peppas, Nayak, Pagnucco, Foo, Kwok, and Prokopenko (1996), and Peppas (1993, Section 5.5), whereas, a discussion on $(K \diamond 1)$–$(K \diamond 9)$ has been conducted by Peppas (2008, Section 8.8). As in the case of the AGM postulates $(K*1)$–$(K*8)$, postulates $(K\diamond 1)$–$(K\diamond 9)$ have been formulated according to the dictates of the principle of *minimal change*; hence, the initial theory K is updated so that the modification be as little as possible.

4.2 Semantic Characterization

Let us, now, proceed to the semantic characterization of belief update. It turns out that the update operators that satisfy postulates $(K \diamond 1)$–$(K \diamond 9)$ are precisely those that are induced by means of *partial* preorders over all possible worlds.

Definition 3 (Faithful Preorder Associated with Worlds, Katsuno & Mendelzon, 1992)**.** *A partial preorder $\preceq_w$ over $\mathbb{M}$ is faithful to a world w iff, for any $r \in \mathbb{M}$, $w \neq r$ implies $w \prec_w r$.*

As in the case of faithful preorders associated with theories, $r \preceq_w r'$ states that the world r is at least as plausible (relative to w) as the world r'.

Definition 4 (Faithful Pointwise Assignment, Katsuno & Mendelzon, 1992)**.** *A faithful pointwise assignment is a function that maps each world w of $\mathbb{M}$ to a partial preorder $\preceq_w$ over $\mathbb{M}$, that is faithful to w.*

The following theorem characterizes the class of update operators induced by faithful pointwise assignments based on *partial* preorders.

Theorem 2 (Katsuno & Mendelzon, 1992)**.** *An update operator $\diamond$ satisfies postulates $(K \diamond 1)$–$(K \diamond 9)$ iff there exists a faithful pointwise assignment that maps each world w of $\mathbb{M}$ to a partial preorder $\preceq_w$ over $\mathbb{M}$, such that, for any theory K and any $\varphi \in \mathcal{L}$:*

$$\textbf{(F}\diamond\textbf{)} \quad \big[K \diamond \varphi\big] = \bigcup_{w \in [K]} min([\varphi], \preceq_w).$$

The semantic characterizations of revision and update, although similar, point out two major differences between these important processes of belief change. Firstly, and perhaps more importantly, to a fixed theory K, a whole *family* of preorders over worlds is assigned in belief update (one for each K-world), contrary to a *single* preorder assigned in belief revision. This pointwise behaviour of update is due to postulate $(K \diamond 9)$, which ensures that the deductive closure of every K-world —which, essentially, constitutes a consistent complete theory— is modified *separately.* Secondly, the preorders corresponding to the AGM postulates $(K * 1)$–$(K * 8)$ are *total*, whereas, the preorders corresponding to the KM postulates $(K \diamond 1)$–$(K \diamond 9)$ are *partial* (and, thus, not necessarily total). For a more detailed discussion on the distinction between revision and update, the interested reader is referred to the work of Peppas et al. (1996).

It proves to be the case that, replacing postulates $(K \diamond 7)$ and $(K \diamond 8)$ by the following postulate $(K \diamond 10)$, we restrict the class of KM update operators only to those that are induced by *total* preorders over possible worlds.[4] For the sake of comparison with belief revision, we shall confine ourselves to this type of KM update operators, unless explicitly stated otherwise. In order to avoid a cumbersome presentation, we shall, also, focus on the *principal* case of consistent theories (belief sets) and contingent new information.

(K ◇ 10) If K is complete and $\neg\psi \notin K \diamond \varphi$, then $(K \diamond \varphi) + \psi \subseteq K \diamond (\varphi \wedge \psi)$.

We conclude this subsection with the next remark, which follows directly from the semantic conditions (F∗) and (F◇).

Remark 1. *Let $*$ be an AGM revision operator, $\diamond$ be a KM update operator, and K be a consistent complete theory of $\mathcal{L}$, such that, for a world $w \in \mathbb{M}$, $[K] = \{w\}$. Moreover, let $\preceq_K$ be a total preorder over $\mathbb{M}$ that $*$ assigns at K, via (F∗), and let $\preceq_w$ be a total preorder over $\mathbb{M}$ that $\diamond$ assigns at w, via (F◇). Then, for any sentence φ of $\mathcal{L}$, it is true that:*

$$K * \varphi = K \diamond \varphi \quad \textit{iff} \quad \preceq_K \, = \, \preceq_w .$$

4.3 Possible Models Approach

Possible Models Approach (PMA) is a popular method, proposed by Winslett (1988), that is suitable for certain applications of reasoning about action. Restricting PMA to a propositional framework, we end up with a (unique) update operator, herein denoted by $\diamond_W$.[5] In order to present Winslett's operator, we need the following definition concerning a notion of *difference* between possible worlds.

Definition 5 (Difference between Worlds)**.** *The difference between two worlds w, r of $\mathbb{M}$, denoted by Diff (w, r), is the set of atoms over which the two worlds disagree. That is:*

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

Then, the operator $\diamond_W$ is induced, via condition (F◇), by means of the family of *partial* preorders $\{\sqsubseteq_w\}_{w \in \mathbb{M}}$, defined, for any $r, r' \in \mathbb{M}$, as follows:

$$\textbf{(W)} \quad r \sqsubseteq_w r' \quad \text{iff} \quad \textit{Diff}\,(w, r) \subseteq \textit{Diff}\,(w, r').$$

It is not hard to verify that the preorder $\sqsubseteq_w$ is *not total*, and faithful to w. Hence, $\diamond_W$ is a KM update operator that satisfies the KM postulates $(K \diamond 1)$–$(K \diamond 9)$, but does not satisfy postulate $(K \diamond 10)$.

It turns out that the use of the KM update operator $\diamond_W$ in Example 1 to modify theory $K = Cn\Big((b \wedge \neg m) \vee (\neg b \wedge m)\Big)$, in the light of the information $\varphi = \neg b$, leads to the desired outcome. In particular, from the fact that $\textit{Diff}\,(w_1, w_2) = \{b, m\}$ and $\textit{Diff}\,(w_1, r) = \{b\}$, we derive that $r \sqsubset_{w_1} w_2$.[6] Consequently, we obtain from condition (F◇) that $\big[K \diamond_W \varphi\big] = \{w_2, r\}$, that is, $K \diamond_W \varphi = Cn(\neg b)$, which is obviously a reasonable result, as all we know after the (successful) action of the robot is that the book is on the floor.

4. The counterpart of $(K \diamond 10)$ in (Katsuno & Mendelzon, 1992) is postulate (U9) — for details on the relation between postulates $(K \diamond 1)$–$(K \diamond 10)$ and (U1)–(U9), the reader is referred to the works of Peppas et al. (1996) and Peppas (1993, Section 5.5).
5. The original definition of PMA has been provided by Winslett for the case of first-order calculus.
6. $\sqsubset_{w_1}$ denotes the strict part of $\sqsubseteq_{w_1}$.

4.4 Counter-Intuitive KM Update Operators

Unfortunately, as it is shown in the next remark, not every KM update operator leads to intuitive results as far as the well-known book/magazine example is concerned.

Remark 2. *Let $\mathcal{P} = \{b, m\}$. Moreover, let K and φ be a theory and a sentence of $\mathcal{L}$, respectively, such that $K = Cn\Big((b \wedge \neg m) \vee (\neg b \wedge m)\Big)$ and $\varphi = \neg b$ (cf. Example 1). There exists a KM update operator $\diamond$ such that, for any AGM revision operator $*$, $K \diamond \varphi = K * \varphi$.*

Proof. The set $[K]$ contains exactly the two worlds $w_1 = \{b, \neg m\}$ and $w_2 = \{\neg b, m\}$, whereas, the set $[\varphi]$ contains exactly the two worlds w_2 and $r = \{\neg b, \neg m\}$.[7] Let $\diamond$ be a KM update operator, induced (via condition (F$\diamond$)) by a family of total preorders over worlds, such that $w_2 \prec_{w_1} r$. Then, from (F$\diamond$), we have that $K \diamond \varphi = Cn(\neg b \wedge m) = K * \varphi$. ■

Evidently, the KM update operator $\diamond$ of Remark 2 has the same counter-intuitive behaviour, in the context of the book/magazine example, as *any* AGM revision operator — since revision reduces, in this case, to *expansion*. The reason why $\diamond$ fails to produce the desired outcome is that it does not take into account some form of *relevance*; since the robot is instructed to change *only* book's place, there is no reason that the magazine is affected by this modification. It seems that the KM postulates interpret the notion of minimal change as follows: As few formulas of the initial theory K as possible are given up so that the new information φ is (consistently) incorporated into K, yet, it is not ensured that those formulas are, actually, *related* to φ.[8]

One can devise a plethora of update-scenarios, analogous to the book/magazine example, in which KM update operators behave in a counter-intuitive manner. In response to this weakness, we shall present a *relevance-sensitive* framework for belief update, which essentially *strictly strengthens* Katsuno and Mendelzon's framework, and remedies the liberal behaviour of KM update operators. Our proposal is based on Parikh's approach for relevance-sensitive belief revision (1999), introduced in the subsequent section.

5. Parikh's Notion of Relevance

This section is devoted to the presentation of Parikh's notion of relevance, both *axiomatically* and *semantically*.

5.1 Axiomatic Side

After the observation that the severe *full-meet revision* (Alchourrón et al., 1985) —a type of revision that discards *all* prior beliefs of a theory K and retains only the (deductive closure of the) new information φ, in the principal case where φ contradicts K— satisfies the AGM postulates, Parikh (1999) proposed an additional axiom, named (P) and presented below, that addresses the problem of relevance-sensitive belief revision.

7. For convenience, the notation of Example 1 is maintained.
8. In the next section (Section 5), it will be shown that this type of "misbehaviour", as far as the AGM postulates are concerned, has been noted by Parikh (1999). It should be mentioned that an interesting discussion on the interpretation of the notion of minimal change in the AGM setting has been conducted by Rott (2000).

(P) If $K = Cn(x, y)$, where x, y are sentences of disjoint sublanguages $\mathcal{L}_x, \mathcal{L}_y$, respectively, and $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, then $K * \varphi = \big(Cn_{\mathcal{L}_x}(x) \circ \varphi\big) + y$, where $\circ$ is a local revision operator defined over the sublanguage $\mathcal{L}_x$.

In the above condition, $Cn_{\mathcal{L}_x}(x)$ denotes the deductive closure of x *in* the sublanguage $\mathcal{L}_x$; i.e., $Cn_{\mathcal{L}_x}(x) = Cn(x) \cap \mathcal{L}_x$.

Axiom (P) asserts that, *if it happens* that the initial belief set K is *splittable* (i.e., it can be expressed in two *syntax-disjoint* compartments $Cn(x)$ and $Cn(y)$), then this can be exploited in the revision-process, during which the portion of K that is *syntactically irrelevant* to the new information remains intact. Hence, Parikh's axiom perceives relevance *at the theory level*, in the sense that it considers the building blocks of relevance to be the disjoint compartments of a splittable theory; splittability is a property of theories which, in turn, depends on the *contingent* beliefs of an agent.

In subsequent works by Peppas, Chopra, and Foo (2004) and Peppas et al. (2015), two distinct interpretations of axiom (P) were identified, namely, its *weak* and *strong* version. For presenting these two versions of (P), consider the following two conditions that make no reference to a local revision operator.

(P1) If $K = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, and $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, then $(K * \varphi) \cap \overline{\mathcal{L}_x} = K \cap \overline{\mathcal{L}_x}$.

(P2) If $K = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, and $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, then $(K * \varphi) \cap \mathcal{L}_x = \big(Cn(x) * \varphi\big) \cap \mathcal{L}_x$.

Condition (P1) corresponds to the weak version of axiom (P), and says that, if a theory K can be expressed in two *syntax-disjoint* compartments $Cn(x)$ and $Cn(y)$, its revision by a sentence that can be formulated within $\mathcal{L}_x$ should *not* affect anything outside $\mathcal{L}_x$. Appending (P1) with condition (P2), which demands that the $\mathcal{L}_x$-part of a theory is modified *independently* of its $\overline{\mathcal{L}_x}$-part, we get the strong version of axiom (P).

5.2 Semantic Characterization

Peppas et al. (2015) characterized both postulates (P1) and (P2) in terms of total preorders over possible worlds; we present this semantic characterization, since we shall rely on it later in this article. First, however, let us fix the appropriate definitions.[9]

Definition 6 (Theory-Splitting, Parikh, 1999)**.** *Let K be a theory of $\mathcal{L}$, 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 $x_1 \in \mathcal{L}^{Q_1}, \ldots, x_n \in \mathcal{L}^{Q_n}$, such that $K = Cn(x_1, \ldots, x_n)$.*

Parikh (1999) showed that, for every theory K of $\mathcal{L}$, there is a *unique finest* K-splitting, denoted by $\mathcal{F}_K$, which refines every (other) K-splitting.[10] That is to say, there is a *unique* way to consider theory K as being composed of disjoint compartments.

9. For a characterization of postulates (P1)–(P2) in terms of all the well-known constructive models for belief revision, the interested reader is referred to the work of Aravanis et al. (2019).

10. A partition Q' refines another partition Q iff, for every $Q'_i \in Q'$, there is a $Q_j \in Q$, such that $Q'_i \subseteq Q_j$.

Definition 7 (Difference between Theories and Possible Worlds, Peppas et al., 2015)**.** *Let K be a consistent theory of $\mathcal{L}$, and let $\mathcal{F}_K = \{F_1, \ldots, F_n\}$ be the finest K-splitting. Moreover, 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_i of $\mathcal{F}_K$, for which there exists a sentence φ that can be expressed in the sublanguage $\mathcal{L}^{F_i}$, on that K and r disagree. In symbols:*

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

Notice that, in the special case of a consistent *complete* theory K, i.e., whenever $[K] = \{w\}$, for some world $w \in \mathbb{M}$, it is true that, for any $r \in \mathbb{M}$, $\mathit{Diff}\,(K, r) = \mathit{Diff}\,(w, r)$.

Definition 8 (Total-Preorder Filtering, Peppas et al., 2015)**.** *Let $\preceq$ be a total preorder over the possible worlds of $\mathbb{M}$, and let x be a contingent sentence of $\mathcal{L}$. The x-filtering of $\preceq$, denoted by $\preceq^x$, is the (unique) total preorder over $\mathbb{M}$, such that, for any $r, r' \in \mathbb{M}$:*

$$r \preceq^x r' \quad \textit{iff} \quad \textit{there is a world } z \in [r_x], \textit{ such that } z \preceq z', \textit{ for all worlds } z' \in [r'_x].$$

Intuitively, the total preorder $\preceq^x$ can be regarded as a "projection" of the total preorder $\preceq$ to the minimal language $\mathcal{L}_x$ of x, treating the atoms outside $\mathcal{L}_x$ as *invisible*. Notice that, if $\mathcal{L}_x = \mathcal{L}$, then $\preceq^x = \preceq$.

The following example will help us clarify the above definitions.

Example 2. *Let $\mathcal{P} = \{a, b, c\}$ and $K = Cn\big(a \leftrightarrow b, c\big)$. Clearly, theory K is splittable, and the finest K-splitting is $\mathcal{F}_K = \big\{\{a, b\}, \{c\}\big\}$. Consider the world $r = \{\neg a, b, c\}$. Then, Diff $(K, r) = \{a, b\}$, since K and r disagree on the sentence $\varphi = a \vee \neg b \in \mathcal{L}^{\{a,b\}}$; i.e., $K \models \varphi$ and $r \models \neg\varphi$. Thereafter, let $\preceq_K$ be the total preorder over the possible worlds of $\mathbb{M}$, which is faithful to K, shown below:*[11]

$$\begin{array}{ccccccccc} abc & & a\overline{b}c & & & & \overline{a}bc & & \\ & \prec_K & & \prec_K & \overline{a}b\overline{c} & \prec_K & & \prec_K & ab\overline{c} \\ \overline{a}\overline{b}c & & a\overline{b}\overline{c} & & & & \overline{a}\overline{b}\overline{c} & & \end{array}$$

Then, for a sentence $x = a \vee b$ (for which $\mathcal{L}_x = \mathcal{L}^{\{a,b\}}$), the x-filtering $\preceq^x_K$ of $\preceq_K$ is the following total preorder over $\mathbb{M}$:

$$\begin{array}{ccccc} abc & & & & \\ ab\overline{c} & \prec^x_K & a\overline{b}c & \prec^x_K & \overline{a}bc \\ \overline{a}\overline{b}c & & a\overline{b}\overline{c} & & \overline{a}b\overline{c} \\ \overline{a}\overline{b}\overline{c} & & & & \end{array}$$

Against this background, consider the following conditions (Q1)–(Q3), which constrain total preorders over possible worlds.

- **(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 $r \approx_K r'$.
- **(Q3)** If $K = Cn(x, y)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, then $\preceq^x_K = \preceq^x_{Cn(x)}$.

11. For the sake of readability, possible worlds are represented as sequences (rather than sets) of literals, and the negation of an atom p is represented as $\overline{p}$ (instead of $\neg p$).

Theorem 3 states that the semantic characterization of postulate (P1) is the conjunction of conditions (Q1)–(Q2), whereas, the semantic characterization of postulate (P2) corresponds to condition (Q3).

Theorem 3 (Peppas et al., 2015)**.** *Let $*$ be an AGM revision operator, K be a theory of $\mathcal{L}$, and $\preceq_K$ be a total preorder over $\mathbb{M}$ that is faithful to K, corresponding to $*$ by means of (F$*$). Then, $*$ satisfies postulate (P1) iff $\preceq_K$ satisfies conditions (Q1)–(Q2), and, moreover, $*$ satisfies postulate (P2) iff $\preceq_K$ satisfies condition (Q3).*

5.3 A Naïve Translation in Belief Update

For applying Parikh's notion of relevance in the context of belief update, we need to recast postulates (P1)–(P2) so that they constrain the behaviour of update operators. Accordingly, by the replacement of the revision operator $*$ of (P1)–(P2) with an update operator $\diamond$, we end up with the following conditions (P1$\diamond$)–(P2$\diamond$), which correspond to a naïve translation of (P1)–(P2), respectively, in the realm of belief update.

(P1$\diamond$) If $K = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, and $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, then $(K \diamond \varphi) \cap \overline{\mathcal{L}_x} = K \cap \overline{\mathcal{L}_x}$.

(P2$\diamond$) If $K = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, and $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, then $(K \diamond \varphi) \cap \mathcal{L}_x = \big(Cn(x) \diamond \varphi\big) \cap \mathcal{L}_x$.

At first glance, it seems that postulates (P1$\diamond$)–(P2$\diamond$), like (P1)–(P2), act at the theory level, since they refer to splittable theories. A closer look at (P1$\diamond$)–(P2$\diamond$), however, against the background of Katsuno and Mendelzon's framework for belief update, reveals a somewhat different reality, discussed in the subsequent two sections.

6. The Weak Version of Axiom (P) in Belief Update

We start our analysis with the *weak* version of axiom (P) for belief update, namely, postulate (P1$\diamond$), from both an *axiomatic* and a *semantic* perspective.

6.1 Axiomatic Side

It turns out that, in the presence of postulates $(K \diamond 1)$–$(K \diamond 10)$, postulate (P1$\diamond$) becomes *equivalent* to postulate (S1), which is, in turn, *equivalent* to postulate (R1); both (S1) and (R1) are shown below.

(S1) $(K \diamond \varphi) \cap \overline{\mathcal{L}_\varphi} = K \cap \overline{\mathcal{L}_\varphi}$.

(R1) If K is complete, then $(K \diamond \varphi) \cap \overline{\mathcal{L}_\varphi} = K \cap \overline{\mathcal{L}_\varphi}$.

Before presenting the formal result that establishes the equivalence between (P1$\diamond$), (S1) and (R1), let us discuss the intuition behind the latter two conditions. Postulate (S1) states that the new information φ should *not* affect any sentence (belief) of a theory K that is syntactically irrelevant to φ. Contrary to (P1$\diamond$), which encodes relevance at the theory level, postulate (S1) encodes relevance *at the sentential level*, in the sense that it considers the building blocks of relevance to be the sentences of an arbitrary theory, which is *not necessarily splittable*.

Postulate (R1), on the other hand, is identical to (S1) when restricted to consistent *complete* theories, encoding the following simple rule: The update of any consistent complete theory K by a new piece of information φ should *not* affect the φ-irrelevant part of K. Since there is a *one-to-one correspondence* between consistent complete theories and possible worlds,[12] postulate (R1), essentially, acts *at the possible-worlds level*; thus, it is aligned with the *pointwise* operation of the update-process.

The alluded connection between (P1$\diamond$), (S1) and (R1) is established in Theorem 4.

Theorem 4. *Postulate (P1$\diamond$) is equivalent to postulate (S1), and postulate (S1) is equivalent to postulate (R1).*

Proof. Let $\diamond$ be a KM update operator. We need to prove that $\diamond$ satisfies (P1$\diamond$) iff $\diamond$ satisfies (S1) iff $\diamond$ satisfies (R1).

We, first, show that $\diamond$ satisfies (S1) iff $\diamond$ satisfies (R1). The left-to-right implication follows directly, by restricting the application of (S1) to consistent complete theories. As far as the right-to-left implication is concerned, suppose that $\diamond$ satisfies (R1). To show that $\diamond$ satisfies (S1), let K be a theory of $\mathcal{L}$, and let φ be a sentence of $\mathcal{L}$. From (R1), it follows that the $\overline{\mathcal{L}_\varphi}$-part of *all* K-worlds remains unaffected by the $\diamond$-update of K by φ. This again entails that $(K \diamond \varphi) \cap \overline{\mathcal{L}_\varphi} = K \cap \overline{\mathcal{L}_\varphi}$. Hence, we have shown that $\diamond$ satisfies (S1) iff $\diamond$ satisfies (R1), as desired.

Next, we show that (S1) entails (P1$\diamond$). Assume, therefore, that $\diamond$ satisfies (S1). We prove that $\diamond$ satisfies (P1$\diamond$). Let K be a theory of $\mathcal{L}$, such that, for some sentences $x, y \in \mathcal{L}$, $K = Cn(x, y)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$. Moreover, let φ be a sentence of $\mathcal{L}$, such that $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$. Consider any consistent sentence $\psi \in \overline{\mathcal{L}_x}$, such that $\psi \in K$ (i.e., $\psi \in K \cap \overline{\mathcal{L}_x}$). From $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$, we derive that $\psi \in K \cap \overline{\mathcal{L}_\varphi}$. Hence, from (S1), it follows that $\psi \in (K \diamond \varphi) \cap \overline{\mathcal{L}_\varphi}$. Given that $\psi \in \overline{\mathcal{L}_x}$, this again entails that $\psi \in (K \diamond \varphi) \cap \overline{\mathcal{L}_x}$. Therefore, $K \cap \overline{\mathcal{L}_x} \subseteq (K \diamond \varphi) \cap \overline{\mathcal{L}_x}$. With a totally symmetric line of reasoning, we can show that $(K \diamond \varphi) \cap \overline{\mathcal{L}_x} \subseteq K \cap \overline{\mathcal{L}_x}$. Consequently, we have that $(K \diamond \varphi) \cap \overline{\mathcal{L}_x} = K \cap \overline{\mathcal{L}_x}$; that is, (P1$\diamond$).

To conclude the proof, it suffices to show that (P1$\diamond$) entails (R1). Assume, therefore, that $\diamond$ satisfies (P1$\diamond$). We prove that $\diamond$ satisfies (R1). Let K be a consistent complete theory, and let φ be any sentence of $\mathcal{L}$. Then, there exists a world $w \in \mathbb{M}$, such that $K = Cn(w)$. Therefore, there exist two sentences $x, y \in \mathcal{L}$, such that $K = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, and $\mathcal{L}_\varphi = \mathcal{L}_x$. Hence, it follows from (P1$\diamond$) that $(K \diamond \varphi) \cap \overline{\mathcal{L}_\varphi} = K \cap \overline{\mathcal{L}_\varphi}$, as desired. ■

The results of the present subsection are summarized in the following relationship (where $\Longleftrightarrow$ denotes equivalence):

$$(\text{P1}\diamond) \quad \Longleftrightarrow \quad (\text{S1}) \quad \Longleftrightarrow \quad (\text{R1})$$

6.2 Semantic Characterization

We, now, turn to the semantic characterization of postulate (R1), which, in the presence of Theorem 4, coincides with the semantic characterization of postulate (S1) and that of

12. Recall that a consistent theory K is complete iff there exists a world $w \in \mathbb{M}$ such that $[K] = \{w\}$.

postulate (P1$\diamond$). Accordingly, the semantic counterpart of (R1) is the subsequent constraint on a *total* preorder $\preceq_w$ that is faithful to a world w (for any $r, r' \in \mathbb{M}$).

(SR1) If $\mathit{Diff}\,(w,r) \subset \mathit{Diff}\,(w,r')$, then $r \prec_w r'$.

According to (SR1), the less a world differs in atoms from a world w, the more plausible it is, relative to w. In case, however, two worlds r, r' are *Diff*-incomparable (i.e., $\mathit{Diff}\,(w,r) \not\subset \mathit{Diff}\,(w,r')$ and $\mathit{Diff}\,(w,r') \not\subset \mathit{Diff}\,(w,r)$), (SR1) places no constraints on their relative order, with respect to $\preceq_w$. As a consequence, contrary to Winslett's condition (W), constraint (SR1) can be applied to *total* preorders over worlds. Note, lastly, that (SR1) constitutes a strict weakening of (W), and it has been studied by Peppas, Foo, and Nayak (2000), as a measure of similarity between worlds, in the context of belief revision.

The next theorem establishes the correspondence between conditions (R1) and (SR1).

Theorem 5. *Let $\diamond$ be a KM update operator, and let $\{\preceq_w\}_{w\in\mathbb{M}}$ be the family of total preorders over $\mathbb{M}$, corresponding to $\diamond$ by means of (F$\diamond$). Then, $\diamond$ satisfies (R1) iff $\{\preceq_w\}_{w\in\mathbb{M}}$ satisfies (SR1).*

Proof. The proof is analogous to the proof of Theorem 2 by Peppas et al. (2015); we present it herein as well, for completeness.

($\Rightarrow$)

Let K be a consistent complete theory; clearly then, for some world w, $[K] = \{w\}$. Assume that $\diamond$ satisfies postulate (R1) at K. Moreover, assume that, contrary to the theorem, $\preceq_w$ violates (SR1). Then, for some worlds $r, r' \in \mathbb{M}$, $\mathit{Diff}\,(w,r) \subset \mathit{Diff}\,(w,r')$ and $r' \preceq_w r$. Since $\preceq_w$ is faithful to w, this entails that $r \neq w$ and $r' \neq w$.

Define φ to be the conjunction of all literals in r that are not in w; i.e., $\varphi = \bigwedge(r - w)$. Then, from $\mathit{Diff}\,(w,r) \subset \mathit{Diff}\,(w,r')$, we derive that $r' \in [\varphi]$. Given that $r' \preceq_w r$, it follows that either r is not $\preceq_w$-minimal in $[\varphi]$, or, if it is $\preceq_w$-minimal in $[\varphi]$, then so is r'. In either case, there is a world $r'' \in min([\varphi], \preceq_w)$, such that $r'' \neq r$; thus, $\mathit{Diff}\,(w,r) \subset \mathit{Diff}\,(w,r'')$. Consider, now, any atom $p \in \big(\mathit{Diff}\,(w,r'') - \mathit{Diff}\,(w,r)\big)$, and define $q = p$ if $p \in r''$, and $q = \neg p$ otherwise. Clearly, $q \in r''$ and $\neg q \in w$. Furthermore, by construction, $q \notin \mathcal{L}_\varphi$.[13] From $q \in r''$, and since r'' is a $\preceq_w$-minimal φ-world, it follows, from condition (F$\diamond$), that $\neg q \notin K \diamond \varphi$. On the other hand, from $\neg q \in K$ and $\neg q \notin \mathcal{L}_\varphi$, (R1) entails that $\neg q \in K \diamond \varphi$. Contradiction.

($\Leftarrow$)

Assume that, for any $w \in \mathbb{M}$, $\preceq_w$ satisfies condition (SR1), and let K be a consistent complete theory; clearly then, for some world w, $[K] = \{w\}$. Consider a sentence φ of $\mathcal{L}$, and let r be any world in $[K \diamond \varphi]$; thus, from condition (F$\diamond$), $r \in min([\varphi], \preceq_w)$. We will show that $\mathit{Diff}\,(w,r) \subseteq \mathcal{L}_\varphi$. Assume, on the contrary, that there is a literal $l \in w \cap \overline{\mathcal{L}_\varphi}$, such that $l \notin r$. Let r' be the world that agrees with r in all literals, except l. Clearly then, since $r \in [\varphi]$ and $l \notin \mathcal{L}_\varphi$, we derive that $r' \in [\varphi]$. Moreover, by the construction of r', $\mathit{Diff}\,(w,r') \subset \mathit{Diff}\,(w,r)$. Consequently, by condition (SR1), $r' \preceq_w r$. This clearly contradicts our assumption that r is a $\preceq_w$-minimal φ-world. Hence, we have shown that

13. To see that $q \notin \mathcal{L}_\varphi$ (or, equivalently, $p \notin \mathcal{L}_\varphi$), observe that $\mathcal{L}^{\mathit{Diff}\,(w,r)} = \mathcal{L}_\varphi$ and $\mathit{Diff}\,(w,r) \subset \mathit{Diff}\,(w,r'')$; thus, $\mathit{Diff}\,(w,r'')$ contains atoms outside $\mathcal{L}_\varphi$.

Diff $(w, r) \subseteq \mathcal{L}_\varphi$. This shows that all worlds r in $[K \diamond \varphi]$ agree with w on all atoms outside $\mathcal{L}_\varphi$. Therefore, $(K \diamond \varphi) \cap \overline{\mathcal{L}_\varphi} = K \cap \overline{\mathcal{L}_\varphi}$, as desired. ■

It is not hard to verify that the following important remark is true.

Remark 3. *Any KM update operator $\diamond$ satisfying postulate (R1) produces the appropriate outcome in the book/magazine example (Example 1); that is, it holds that*

$$K \diamond \varphi = Cn(\neg b).$$

Hence, the addition of (R1) —or, equivalently, (P1$\diamond$) or (S1)— to the KM postulates excludes the type of undesirable KM update operators identified in Remark 2 of Subsection 4.4 — notice that the KM update operator $\diamond$ in the proof of Remark 2 does not satisfy postulate (R1), as the preorder $\preceq_{w_1}$ violates constraint (SR1).

6.3 Relevance against the Frame, Ramification and Qualification Problem

A classical and significant problem in dynamically-changing worlds, namely, the *frame problem* (McCarthy & Hayes, 1969), is closely connected to relevance-sensitive belief update. The frame problem is the challenge of *succinctly* representing the effects of actions, without having to represent explicitly a large number of intuitively obvious non-effects. Evidently then, the main desideratum of the frame problem is the distinction between the relevant and irrelevant fluents, with respect to an action. Two other widely-studied problems are associated with the frame problem; namely, the *ramification problem* (Finger, 1987), which denotes the problem of handling the indirect effects of actions, and the *qualification problem* (McCarthy, 1977), which is concerned with the impossibility of listing all the preconditions required for a real-world action to have its intended effect.

In notable works on reasoning about action (McCain & Turner, 1995; Thielscher, 1996, 2001), *causality* —which indicates how the world progresses— has been identified as a crucial concept in avoiding unintended ramifications and inferring derived qualifications. Accordingly, Thielscher (1996, 2001) proposes that the incorporation of causal information, concerning a domain to be encoded, can be accomplished by statements, known as *causal relationships*, of the following form (where a, b, c are atoms of the language):

"A change of $\neg a$ to a causes a change of $\neg b$ to b, provided that c is true."

It is important to mention that causal relationships should *not* be considered as pure logical formulas, identical to those contained in belief sets; logical formulas do not include causal information, in contrast to causal relationships which encode the circumstances under which the occurrence of an effect *causes* another effect.

As noted by Thielscher, however, mere causality is insufficient to properly handle ramifications and qualifications. He argues that what is additionally needed is some kind of *persistence law*, which specifies those aspects of the world that are *unaffected* by an action's direct and indirect effects, and, moreover, distinguishes the *normal* from *abnormal* qualifications. As we sketch in the remainder of this subsection, the role of an *axiom of persistence* can be performed by a slight adjustment of the *weak* version of Parikh's axiom for belief update (namely, postulate (R1) or (S1) or (P1$\diamond$)), which incorporates causality.

Let us, first, fix the appropriate notation and terminology. For a causal relationship ρ, $\mathcal{P}_\rho$ denotes the set of atoms occurring in ρ. Let φ be a sentence of $\mathcal{L}$, which represents the occurrence of a new action. We shall say that the causal relationships $\rho_1, \rho_2, \ldots, \rho_{n-1}, \rho_n$ are *φ-dependent* iff $\mathcal{P}_\varphi \cap \mathcal{P}_{\rho_1} \neq \varnothing$ and $\mathcal{P}_{\rho_1} \cap \mathcal{P}_{\rho_2} \neq \varnothing$ and $\ldots$ and $\mathcal{P}_{\rho_{n-1}} \cap \mathcal{P}_{\rho_n} \neq \varnothing$. Intuitively, the fluents represented by the atoms contained in the φ-dependent causal relationships have a *potential causal influence* by the action represented by φ. Notice that this type of causal dependency is syntax-based. As an example, consider the following two causal relationships: "A change of $\neg a$ to a causes either a change of $\neg b$ to b or a change of $\neg c$ to c" and "A change of $\neg c$ to c causes a change of $\neg d$ to d", where a, b, c, d are atoms of the language. Then, assuming that $\varphi = a$, both the aforementioned causal relationships are φ-dependent (although a change of $\neg a$ to a does not, necessarily, cause a change of $\neg d$ to d). Thereafter, let us denote by $\mathcal{P}^{\mathcal{C}}_\varphi$ the set containing the atoms occurring in *all* the φ-dependent causal relationships. Then, we define $\mathcal{L}^{\mathcal{C}}_\varphi$ to be the language generated from the set of atoms $\mathcal{P}_\varphi \cup \mathcal{P}^{\mathcal{C}}_\varphi$, using the standard Boolean connectives; i.e., $\mathcal{L}^{\mathcal{C}}_\varphi = \mathcal{L}^{(\mathcal{P}_\varphi \cup \mathcal{P}^{\mathcal{C}}_\varphi)}$. Obviously, $\mathcal{L}_\varphi \subseteq \mathcal{L}^{\mathcal{C}}_\varphi$.

Against this background, consider the following rule (RC), which, essentially, encodes a *weakening* of the modification-policy encoded in the weak version of Parikh's axiom for belief update (since $\mathcal{L}_\varphi \subseteq \mathcal{L}^{\mathcal{C}}_\varphi$).

(RC) "The modification of a world state, in response to a sentence φ that represents the occurrence of a new action, should *not* affect, or affected by, any fluent which is represented by atoms contained in $\overline{\mathcal{L}^{\mathcal{C}}_\varphi}$."

The above principle ensures, in a simple and compact manner, that only things *relevant to an action* change in a dynamic environment, whereas, the rest remain unaffected. Thus, for instance, rule (RC) would guarantee that moving does not affect color and painting does not affect location. It follows, then, that a modification-policy adhering to the relevance-sensitive principle of (RC) ensures that:

- A new action, represented by a sentence φ, *cannot* initiate any change of fluent represented by atoms of the language $\overline{\mathcal{L}^{\mathcal{C}}_\varphi}$. That is to say, any such fluent is considered *irrelevant*, and it is not affected during the modification-process. This, in turn, entails that only the relevant ramifications (with respect to a new action) will be accommodated in the new state of the world.

- During the modification of a world state, in response to a new action represented by a sentence φ, any qualifications (preconditions) *preventing* the successful occurrences of effects, which are represented by atoms of the language $\overline{\mathcal{L}^{\mathcal{C}}_\varphi}$, are assumed *abnormal* (or *irrelevant*).

The formal scheme, described above, that handles (at least to a certain degree) the frame, ramification and qualification problems, by means of relevance-sensitive modification-policies, is depicted in Figure 1. The following concrete example illustrates its application to a representative real-world dynamic scenario.

Example 3. *Consider the simple electric circuit of Figure 2, which consists of a battery, two switches ($switch_1$, $switch_2$) and two light bulbs ($light\_bulb_1$, $light\_bulb_2$). We shall refer*

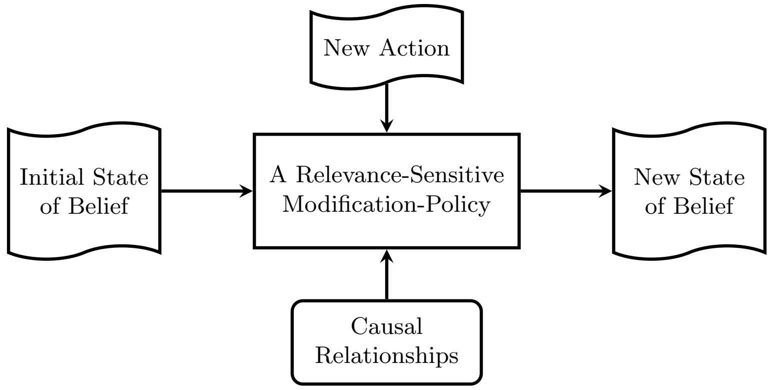

Figure 1: A formal scheme that handles (at least to a certain degree) the frame, ramification and qualification problems, by means of relevance-sensitive modification-policies.

to the filaments of light_bulb$_1$ *and light_bulb*$_2$ *as* $filament_1$ *and* $filament_2$*, respectively. Suppose that the language used to represent the state of the circuit is built over the set of atoms* $\mathcal{P} = \{s_1, l_1, f_1, s_2, l_2, f_2\}$*, where* s_i*,* l_i*,* f_i *denote the propositions "switch*$_i$ *is on", "light_bulb*$_i$ *is on", "the filament of light_bulb*$_i$ *is damaged", respectively, and* $i \in \{1, 2\}$*. The topology of the circuit entails that the left side of the circuit (i.e., the condition of components switch*$_1$*, light_bulb*$_1$ *and* $filament_1$*) is causally independent from the right side of the circuit (i.e., the condition of components switch*$_2$*, light_bulb*$_2$ *and* $filament_2$*); therefore, each light bulb is solely controlled by its respective switch, and, moreover, for the sentence* $\varphi = s_1$*, it is true that* $\mathcal{L}^{\mathcal{C}}_{\varphi} = \mathcal{L}^{\{s_1, l_1, f_1\}}$*.*

In the initial state of the world, we know that both switches are turned off, thus, both light bulbs are off. We have, also, no information of a damaged $filament_1$*. Furthermore, assume that the causal properties of the domain are encoded in the following two causal relationships.*

"A change of $\neg s_1$ *to* s_1 *causes a change of* $\neg l_1$ *to* l_1*, provided that* f_1 *is not true."*

"A change of $\neg s_2$ *to* s_2 *causes a change of* $\neg l_2$ *to* l_2*, provided that* f_2 *is not true."*

Now, we instruct a robot to toggle switch$_1$ *on. In order to estimate the new state of the world, after the (successful) action of the robot, we need to modify the initial state of the world, taking into account the new information* $\varphi = s_1$*. Supposing that our modification-policy respects rule (RC), we expect, under the current assumptions, that* s_1 *will become true (direct effect), and that light_bulb*$_1$ *will turn on (indirect effect), in the new state of the world. On the other hand, rule (RC) dictates that nothing is changed in the right side of the circuit, during the operation of the robot; hence, for example, we have no reason to believe that, in the new state of the world, light_bulb*$_2$ *turns on. Note, also, that we have no reason to believe that light_bulb*$_1$ *is off, since we have no information of a damaged* $filament_1$ *in the initial state of the world (thus,* f_1 *is not true), and, moreover, any qualifications (preconditions) preventing the successful occurrences of effects, which are represented by atoms of* $\overline{\mathcal{L}^{\mathcal{C}}_{\varphi}}$*, are assumed irrelevant.*

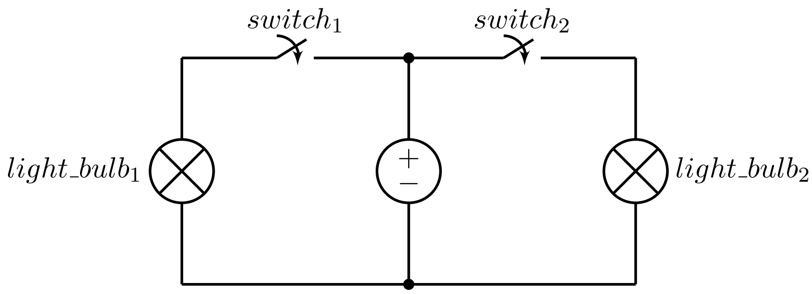

Figure 2: A simple electric circuit, consisting of a battery, two switches and two light bulbs.

Consequently, only the plausible (relevant) ramifications of the action of the robot are accommodated in the new state of the world, and, moreover, any implausible qualifications that prevent the successful occurrences of effects are not taken into account.

In this subsection, we essentially demonstrated how a slight modification of postulate (R1) —namely, rule (RC), which, as a matter of fact, does not refer to KM update operators— can be used as an axiom of persistence, in the context of dynamically-changing worlds. The reason for adopting a variance of (R1), and not (R1) itself, is that the syntax-relevance encoded in (R1) (contrary to that encoded in (RC)) does not take into account the causal properties of a domain. In the absence of causal information, the syntax-relevance of (R1) can be restrictive, leading to counter-intuitive results. To see this, consider Example 3, in the context of which a robot performs a new action, represented by a sentence φ, namely, toggles $switch_1$ on; i.e., $\varphi = s_1$. On that premise, postulate (R1) would lead to a new state of the world in which $light\_bulb_1$ would be off, since, in the initial state of the world, $\neg l_1$ holds, and, according to (R1), no fluent represented by atoms of the language $\overline{\mathcal{L}_\varphi}$ should be modified, during the operation of the robot. This is, clearly, a counter-intuitive outcome, as the fluents represented by s_1 and l_1, although syntax-irrelevant, are causally dependent.

Undoubtedly, a deeper investigation of how syntax-relevance in belief update can properly accommodate causal information constitutes an appealing avenue for future research.[14]

7. The Strong Version of Axiom (P) in Belief Update

In this section, we study the *strong* version of Parikh's axiom (P) for belief update, namely, the conjunction of postulates (P1$\diamond$) and (P2$\diamond$). Given that (P1$\diamond$) was analysed earlier, we study, in this section, postulate (P2$\diamond$), both *axiomatically* and *semantically.*

7.1 Axiomatic Side

It turns out that, in the presence of postulates $(K \diamond 1)$–$(K \diamond 10)$, (P2$\diamond$) *strictly implies* postulate (R2), which is, in turn, *equivalent* to postulate (S2); both (S2) and (R2) are shown below.

14. An interesting approach that studies the relation between causality and the notion of minimal change —which is directly associated with relevance— has been followed by Pagnucco and Peppas (2001).

(S2) If $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$, then $(K \diamond \varphi) \cap \mathcal{L}_\varphi = (H \diamond \varphi) \cap \mathcal{L}_\varphi$.

(R2) If K, H are complete and $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$, then $(K \diamond \varphi) \cap \mathcal{L}_\varphi = (H \diamond \varphi) \cap \mathcal{L}_\varphi$.

Before presenting the formal results that establish the connections between (P2$\diamond$), (S2) and (R2), let us discuss the intuition behind the latter two conditions. Postulate (S2) states that, if two theories K and H agree on their $\mathcal{L}_\varphi$-parts (namely, on all sentences of $\mathcal{L}_\varphi$), then the updated theories $K \diamond \varphi$ and $H \diamond \varphi$ should agree on their $\mathcal{L}_\varphi$-parts as well. Hence, according to (S2), the φ-irrelevant part of a theory does not affect the way that the φ-relevant part of it is modified; thus, the update-process becomes *context-independent*. As in the case of (S1), and contrary to (P2$\diamond$), condition (S2) encodes relevance *at the sentential level*, and does not rely on properties of theories (e.g., splittability), since it refers to *arbitrary* theories of the language.

Postulate (R2), on the other hand, encodes the exact principle that (S2) encodes, but restricted only to consistent *complete* theories. Since there is a *one-to-one correspondence* between consistent complete theories and possible worlds, postulate (R2) —like (R1)— acts *at the possible-worlds level*; thus, it is aligned with the *pointwise* operation of belief update.

The alluded connection between (P2$\diamond$), (S2) and (R2) is established in the subsequent three theorems.

Theorem 6. *Postulate (S2) is equivalent to postulate (R2).*

Proof. Let $\diamond$ be a KM update operator. We need to prove that $\diamond$ satisfies (S2) iff $\diamond$ satisfies (R2). The left-to-right implication follows directly, by restricting the application of (S2) to consistent complete theories. As far as the right-to-left implication is concerned, suppose that $\diamond$ satisfies (R2). To show that $\diamond$ satisfies (S2), let K, H be two theories of $\mathcal{L}$, and let φ be any sentence of $\mathcal{L}$, such that $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$. From $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$, we derive that $\{r_\varphi : r \in [K]\} = \{r_\varphi : r \in [H]\}$. This, together with (R2), entails that $\{r_\varphi : r \in [K \diamond \varphi]\} = \{r_\varphi : r \in [H \diamond \varphi]\}$. Therefore, $(K \diamond \varphi) \cap \mathcal{L}_\varphi = (H \diamond \varphi) \cap \mathcal{L}_\varphi$, as desired. ■

Theorem 7. *Postulate (P2$\diamond$) entails postulate (R2).*

Proof. Let $\diamond$ be a KM update operator, and assume that $\diamond$ satisfies (P2$\diamond$). We prove that $\diamond$ satisfies (R2). Let K, H be two consistent complete theories and φ be a sentence of $\mathcal{L}$, such that $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$. Since K, H are consistent complete theories, there exist two worlds $w, w' \in \mathbb{M}$, such that $K = Cn(w)$ and $H = Cn(w')$. This, together with the fact that $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$, entails that there exist sentences $x, y, z \in \mathcal{L}$, such that $K = Cn(x, y)$, $H = Cn(x, z)$, $\mathcal{L}_x \cap \mathcal{L}_y = \mathcal{L}_x \cap \mathcal{L}_z = \varnothing$, and $\mathcal{L}_\varphi = \mathcal{L}_x$. Therefore, it follows from (P2$\diamond$) that $(K \diamond \varphi) \cap \mathcal{L}_\varphi = (H \diamond \varphi) \cap \mathcal{L}_\varphi$, as desired. ■

The converse of Theorem 7 is not, in general, true. This is shown in Theorem 8, which, together with Theorem 7, proves that postulate (P2$\diamond$) is, in fact, a *strict strengthening* of postulate (R2).

Theorem 8. *There exists a KM update operator that satisfies postulate (R2), but does not satisfy postulate (P2$\diamond$).*

Proof. Assume that $\mathcal{P} = \{a, b, c\}$. Let x be a sentence of $\mathcal{L}$, such that $x = (a \wedge b) \vee (\neg a \wedge \neg b)$, and let $K = Cn(x, c)$ and $H = Cn(x, \neg c)$; clearly, both K and H split in two disjoint compartments. Moreover, the set $[K]$ contains exactly the two worlds $w_1 = \{a, b, c\}$ and $w_2 = \{\neg a, \neg b, c\}$, whereas, the set $[H]$ contains exactly the two worlds $u_1 = \{a, b, \neg c\}$ and $u_2 = \{\neg a, \neg b, \neg c\}$.

Let φ be a sentence of $\mathcal{L}$, such that $\varphi = \neg a$. Hence, $\mathcal{L}_\varphi \subset \mathcal{L}_x$ and, thus, $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$. Moreover, let $\diamond$ be a KM update operator that satisfies (R2), induced (via (F$\diamond$)) by a family $\{\preceq_w\}_{w \in \mathbb{M}}$ of total preorders, such that $min([\varphi], \preceq_{w_1}) = \big\{\{\neg a, \neg b, c\}\big\} = \big[Cn(w_1) \diamond \varphi\big]$, and $min([\varphi], \preceq_{u_1}) = \big\{\{\neg a, b, \neg c\}\big\} = \big[Cn(u_1) \diamond \varphi\big]$. Then, condition (F$\diamond$) entails that $[K \diamond \varphi] = \big\{\{\neg a, \neg b, c\}\big\}$ and $[H \diamond \varphi] = \big\{\{\neg a, b, \neg c\}, \{\neg a, \neg b, \neg c\}\big\}$; that is, $K \diamond \varphi = Cn(\neg a, \neg b, c)$ and $H \diamond \varphi = Cn(\neg a, \neg c)$.[15] Observe that $(K \diamond \varphi) \cap \mathcal{L}_x \neq (H \diamond \varphi) \cap \mathcal{L}_x$, even though $K \cap \mathcal{L}_x = H \cap \mathcal{L}_x$, hence, $\diamond$ violates postulate (P2$\diamond$). ■

A key element in the proof of the above theorem is the consideration of a new piece of information φ whose minimal language $\mathcal{L}_\varphi$ is a *proper* subset of the minimal language $\mathcal{L}_x$ of the sentence x, which identifies the first compartment of the splittable theories K and H. Under this situation, postulate (R2) —or, equivalently, (S2)— constrains (only) the $\mathcal{L}_\varphi$-part of theories K, H, $K \diamond \varphi$ and $H \diamond \varphi$, but not the (wider) $\mathcal{L}_x$-part of these theories, which postulate (P2$\diamond$) constrains. This is the root cause of the fact that the converse of Theorem 7 does not, in general, hold.

The results of the present subsection are summarized in the following relationship (where $\Longrightarrow$, $\Longleftrightarrow$ denote strict implication and equivalence, respectively):

$$\text{(P2}\diamond\text{)} \quad \Longrightarrow \quad \text{(S2)} \quad \Longleftrightarrow \quad \text{(R2)}$$

7.2 Semantic Characterization

In this subsection, we formulate the semantic characterizations of postulate (R2) —or, equivalently, (S2)— and postulate (P2$\diamond$). First, we introduce the semantic counterpart of (R2), which is the following constraint on *total* preorders over worlds.

(SR2) If $w_\varphi = w'_\varphi$, then $\preceq_w^\varphi = \preceq_{w'}^\varphi$.

According to (SR2), the total preorders that are faithful to any two worlds that agree on all atoms of $\mathcal{L}_\varphi$, ought to have identical φ-filterings.

The next theorem establishes the correspondence between conditions (R2) and (SR2).

Theorem 9. *Let $\diamond$ be a KM update operator, and let $\{\preceq_w\}_{w \in \mathbb{M}}$ be the family of total preorders over $\mathbb{M}$, corresponding to $\diamond$ by means of (F$\diamond$). Then, $\diamond$ satisfies (R2) iff $\{\preceq_w\}_{w \in \mathbb{M}}$ satisfies (SR2).*

Proof. The proof is analogous to the proof of Theorem 5 by Peppas et al. (2015); we present it herein as well, for completeness.
($\Rightarrow$)
Assume that $\diamond$ satisfies postulate (R2). Let w, w' be two worlds of $\mathbb{M}$ and φ be a sentence of $\mathcal{L}$, such that $w_\varphi = w'_\varphi$. First, we show that $\preceq_w^\varphi \subseteq \preceq_{w'}^\varphi$.

15. Notice that the worlds w_2 and u_2 satisfy φ.

Let r, r' be any two worlds such that $r \preceq_w^\varphi r'$. Define ψ to be the sentence $\psi = \big(\bigwedge r_\varphi\big) \vee \big(\bigwedge r'_\varphi\big)$. Clearly, $\psi \in \mathcal{L}_\varphi$. From the definition of $\preceq_w^\varphi$, we derive that there is a world $u \in [r_\varphi]$ which is $\preceq_w$-minimal in $[r_\varphi] \cup [r'_\varphi]$. This again entails, from condition (F$\diamond$), that $\neg\big(\bigwedge r_\varphi\big) \notin Cn(w) \diamond \psi$. Then, given that $Cn(w) \cap \mathcal{L}_\varphi = Cn(w') \cap \mathcal{L}_\varphi$ (as $w_\varphi = w'_\varphi$), it follows from postulate (R2) that $\neg\big(\bigwedge r_\varphi\big) \notin Cn(w') \diamond \psi$. Hence, there is a world $z \in [r_\varphi]$ which is $\preceq_{w'}$-minimal in $[r_\varphi] \cup [r'_\varphi]$. This again entails that $r \preceq_{w'}^\varphi r'$. Therefore, $\preceq_w^\varphi \subseteq \preceq_{w'}^\varphi$.

The proof of the converse, i.e., $\preceq_{w'}^\varphi \subseteq \preceq_w^\varphi$, is totally symmetric.

($\Leftarrow$)

Assume that condition (SR2) is satisfied. Let K, H be two consistent complete theories and let φ be a sentence of $\mathcal{L}$, such that $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$. Observe that, since K and H are consistent complete theories, there are worlds $w, w' \in \mathbb{M}$, such that $[K] = \{w\}$ and $[H] = \{w'\}$. First, we show that $(H \diamond \varphi) \cap \mathcal{L}_\varphi \subseteq (K \diamond \varphi) \cap \mathcal{L}_\varphi$.

Consider any consistent sentence ψ such that $\psi \notin (K \diamond \varphi) \cap \mathcal{L}_\varphi$. Then, from condition (F$\diamond$), there is a $\preceq_w$-minimal φ-world r, such that $r \models \neg\psi$. Moreover, it is true that $\bigcup\big\{[r'_\varphi] : r' \in [\varphi]\big\} = [\varphi]$. Hence, since r is $\preceq_w$-minimal in $[\varphi]$, we derive that r is, also, $\preceq_w$-minimal in $\bigcup\big\{[r'_\varphi] : r' \in [\varphi]\big\}$. This again entails, from the definition of $\preceq_w^\varphi$, that r is $\preceq_w^\varphi$-minimal in $\bigcup\big\{[r'_\varphi] : r' \in [\varphi]\big\}$. Then, from condition (SR2), we derive that r is $\preceq_{w'}^\varphi$-minimal in $\bigcup\big\{[r'_\varphi] : r' \in [\varphi]\big\}$. Hence, r is $\preceq_{w'}$-minimal in $[\varphi]$, and, therefore, from condition (F$\diamond$), $r \in [H \diamond \varphi]$. Given that $r \models \neg\psi$, it follows that $\psi \notin H \diamond \varphi$. Consequently, we have shown that $(H \diamond \varphi) \cap \mathcal{L}_\varphi \subseteq (K \diamond \varphi) \cap \mathcal{L}_\varphi$.

The proof of the converse, i.e., $(K \diamond \varphi) \cap \mathcal{L}_\varphi \subseteq (H \diamond \varphi) \cap \mathcal{L}_\varphi$, is totally symmetric. ■

As far as postulate (P2$\diamond$) is concerned, its semantic counterpart is the following condition (SP2$\diamond$), which constrains *total* preorders over worlds. In the special case where theories K and H in condition (SP2$\diamond$) are consistent complete theories, (SP2$\diamond$) reduces to condition (SR2); this is to be expected, since, in that case, postulate (P2$\diamond$) reduces to postulate (R2) — for the rationale behind this latter statement, the reader is referred to the proof of Theorem 7.

(SP2$\diamond$) If $K = Cn(x, y)$, $H = Cn(x, z)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \mathcal{L}_x \cap \mathcal{L}_z = \varnothing$, then, for every $w \in [K]$, there is a $w' \in [H]$, such that $\preceq_w^x \subseteq \preceq_{w'}^x$.

Theorem 10 establishes the correspondence between conditions (P2$\diamond$) and (SP2$\diamond$).

Theorem 10. *Let $\diamond$ be a KM update operator, and let $\{\preceq_w\}_{w \in \mathbb{M}}$ be the family of total preorders over $\mathbb{M}$, corresponding to $\diamond$ by means of (F$\diamond$). Then, $\diamond$ satisfies (P2$\diamond$) iff $\{\preceq_w\}_{w \in \mathbb{M}}$ satisfies (SP2$\diamond$).*

Proof.

($\Rightarrow$)

Assume that $\diamond$ satisfies postulate (P2$\diamond$). Let K, H be two theories of $\mathcal{L}$, such that, for some sentences $x, y, z \in \mathcal{L}$, $K = Cn(x, y)$, $H = Cn(x, z)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \mathcal{L}_x \cap \mathcal{L}_z = \varnothing$. We show that, for every $w \in [K]$, there is a $w' \in [H]$, such that $\preceq_w^x \subseteq \preceq_{w'}^x$.

Let w be any world in $[K]$, and let r, r' be any two worlds of $\mathbb{M}$, such that $r \preceq_w^x r'$. Define φ to be the sentence $\varphi = \big(\bigwedge r_x\big) \vee \big(\bigwedge r'_x\big)$. Clearly, $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$. From the definition of $\preceq_w^x$, we derive that there is a world $z \in [r_x]$ which is $\preceq_w$-minimal in $[r_x] \cup [r'_x]$. This again entails, from (F$\diamond$), that $\neg\big(\bigwedge r_x\big) \notin K \diamond \varphi$. Then, from postulate (P2$\diamond$), $\neg\big(\bigwedge r_x\big) \notin H \diamond \varphi$.[16] Hence, (F$\diamond$) entails that, for some world $w' \in [H]$, there is a world $u \in [r_x]$ which is $\preceq_{w'}$-minimal in $[r_x] \cup [r'_x]$. Therefore, from the definition of $\preceq^x_{w'}$, we derive that $r \preceq^x_{w'} r'$, as desired.

$(\Leftarrow)$

Assume that condition (SP2$\diamond$) is satisfied. Let K be a theory of $\mathcal{L}$, such that, for some sentences $x, y \in \mathcal{L}$, $K = Cn(x, y)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$. Moreover, let φ be a sentence, such that $\mathcal{L}_\varphi \subseteq \mathcal{L}_x$; thus, it is true that $\bigcup\big\{[r'_x] : r' \in [\varphi]\big\} = [\varphi]$. First, we show that $\big(Cn(x)\diamond\varphi\big)\cap\mathcal{L}_x \subseteq (K\diamond\varphi)\cap\mathcal{L}_x$. Consider any consistent sentence ψ such that $\psi \notin (K\diamond\varphi)\cap\mathcal{L}_x$. We will prove that $\psi \notin (Cn(x)\diamond\varphi)\cap\mathcal{L}_x$. If $\psi \notin \mathcal{L}_x$, this is trivially true. Assume, therefore, that $\psi \in \mathcal{L}_x$ and $\psi \notin (K \diamond \varphi)$. Then, from condition (F$\diamond$), for some world $w \in [K]$, there is a $\preceq_w$-minimal φ-world r, such that $r \models \neg\psi$.

From $\bigcup\big\{[r'_x] : r' \in [\varphi]\big\} = [\varphi]$ and since r is $\preceq_w$-minimal in $[\varphi]$, we derive that r is, also, $\preceq_w$-minimal in $\bigcup\big\{[r'_x] : r' \in [\varphi]\big\}$. From the definition of $\preceq^x_w$, this entails that r is $\preceq^x_w$-minimal in $\bigcup\big\{[r'_x] : r' \in [\varphi]\big\}$. Then, from condition (SP2$\diamond$), we derive that there is a world $w' \in [Cn(x)]$, such that r is $\preceq^x_{w'}$-minimal in $\bigcup\big\{[r'_x] : r' \in [\varphi]\big\}$. From the definition of $\preceq^x_{w'}$ and since $\bigcup\big\{[r'_x] : r' \in [\varphi]\big\} = [\varphi]$, we have that r is $\preceq_{w'}$-minimal in $[\varphi]$, and, therefore, from condition (F$\diamond$), $r \in [Cn(x)\diamond\varphi]$. Given that $r \models \neg\psi$, it follows that $\psi \notin Cn(x)\diamond\varphi$, thus, $\psi \notin (Cn(x) \diamond \varphi) \cap \mathcal{L}_x$. Consequently, we have shown that $\big(Cn(x) \diamond \varphi\big) \cap \mathcal{L}_x \subseteq (K \diamond \varphi) \cap \mathcal{L}_x$, as desired.

The proof of the converse, i.e., $(K \diamond \varphi) \cap \mathcal{L}_x \subseteq \big(Cn(x) \diamond \varphi\big) \cap \mathcal{L}_x$, is totally symmetric, since condition (SP2$\diamond$) entails that, for every $w \in [Cn(x)]$, there is a $w' \in [K]$, such that $\preceq^x_w \subseteq \preceq^x_{w'}$. ■

It should be evident that postulates (R1) (equivalently, (S1) or (P1$\diamond$)), (R2) (equivalently, (S2)) and (P2$\diamond$) constitute reasonable constraints of the behaviour of arbitrary KM update operators; thus, each one of these postulates circumscribes a particular *proper subclass* of the whole class of KM update operators. Furthermore, although (R2) and (P2$\diamond$) are related, they are both quite independent from (R1). In any case, depending on the underlying application, one may argue in favour of many types of constraints; postulates (R1), (R2) and (P2$\diamond$), however, may be seen as core *domain-independent* rules for relevance-sensitive belief update.

We close this section noting that the characterization results of Aravanis (2019) and Aravanis et al. (2019) can be utilized to, straightforwardly, translate conditions (SR1) and (SR2) in the realm of the other two well-known semantic models for belief change — namely, the *epistemic-entrenchment* (Gärdenfors & Makinson, 1988) and *partial-meet* (Alchourrón et al., 1985) models.

8. Parametrized-Difference Update Operators

Having introduced the notion of relevance in the realm of belief update, we turn to the (*semantic* and *axiomatic*) introduction of a well-behaved family of concrete KM update operators, based on the approach of *parametrized-difference* belief revision by Peppas and

16. Notice that postulate (P2$\diamond$) entails that $(K \diamond \varphi) \cap \mathcal{L}_x = \big(Cn(x) \diamond \varphi\big) \cap \mathcal{L}_x = (H \diamond \varphi) \cap \mathcal{L}_x$.

Williams (2016, 2018).[17] We call these new operators *parametrized-difference* (PD) *update operators*; for short, *PD update operators*. As it will be shown, the alluded operators are quite natural and intuitive, compactly-specified, as they are induced from a *fixed* ordering over atoms, and relevance-sensitive, since they satisfy *all* the relevance-sensitive postulates presented herein; this latter feature of PD update operators proves that the conjunction of the presented postulates is *consistent* with the KM postulates.

8.1 Semantic Characterization

Let $\trianglelefteq$ be a total preorder over the set $\mathcal{P}$ of all atoms. As the atoms of the language, essentially, represent conditions (facts) of a world state, we assume that the preorder $\trianglelefteq$ reflects the (prior) comparative plausibility of change of these world conditions; the more resistant to change a world condition is, the higher it appears in the preorder $\trianglelefteq$. That is to say, $a \trianglelefteq b$ asserts that a change of the world condition b is *less plausible* than a change of the world condition a.

For a set of atoms $\mathcal{S}$ and an atom q, by $\mathcal{S}_q$ we denote the set $\mathcal{S}_q = \{p \in \mathcal{S} : p \trianglelefteq q\}$. The definition of $\trianglelefteq$ can, then, be extended to *sets* of atoms.

Definition 9 (Total Preorder over Sets of Atoms, Peppas & Williams, 2016)**.** *Let $\trianglelefteq$ be a total preorder over the set of atoms $\mathcal{P}$. For any two sets of atoms $\mathcal{S}$, $\mathcal{S}'$, $\mathcal{S} \trianglelefteq \mathcal{S}'$ iff one of the following three conditions holds ($\triangleleft$ denotes the strict part of $\trianglelefteq$):*

- **(i)** $|\mathcal{S}| < |\mathcal{S}'|$.
- **(ii)** $|\mathcal{S}| = |\mathcal{S}'|$, *and for all* $q \in \mathcal{P}$, $|\mathcal{S}_q| = |\mathcal{S}'_q|$.
- **(iii)** $|\mathcal{S}| = |\mathcal{S}'|$, *and for some* $q \in \mathcal{P}$, $|\mathcal{S}_q| > |\mathcal{S}'_q|$, *and for all* $p \triangleleft q$, $|\mathcal{S}_p| = |\mathcal{S}'_p|$.

In the above definition, condition (ii) states that $\mathcal{S}$ and $\mathcal{S}'$ are *lexicographically indistinguishable* (with respect to $\trianglelefteq$), whereas, condition (iii) states that $\mathcal{S}$ *lexicographically precedes* $\mathcal{S}'$ (with respect to $\trianglelefteq$). It turns out that the extended $\trianglelefteq$ of Definition 9 is a total preorder over $2^{\mathcal{P}}$. The intended interpretation of the extended total preorder $\trianglelefteq$, defined over $2^{\mathcal{P}}$, is the same as that of a total preorder defined over $\mathcal{P}$; namely, $\mathcal{S} \trianglelefteq \mathcal{S}'$ states that a change of the world condition represented by *all* atoms of $\mathcal{S}'$ is less plausible than a change of the world condition represented by *all* atoms of $\mathcal{S}$.

Definition 10 (PD Preorder Associated with Worlds)**.** *Let $\trianglelefteq$ be a total preorder over the set of atoms $\mathcal{P}$. A PD preorder, associated with a world $w \in \mathbb{M}$ and denoted by $\preceq_w^{\trianglelefteq}$, is any binary relation over $\mathbb{M}$, such that, for any $r, r' \in \mathbb{M}$:*

$$\textbf{(PD)} \quad r \preceq_w^{\trianglelefteq} r' \quad \textit{iff} \quad \textit{Diff}\,(w, r) \trianglelefteq \textit{Diff}\,(w, r').$$

The results of Peppas and Williams (2016) entail that $\preceq_w^{\trianglelefteq}$ is a *total* preorder over $\mathbb{M}$, faithful to w.

17. Parametrized-difference belief revision was further investigated by Aravanis, Peppas, and Williams (2021).

In the special case where $\trianglelefteq = \mathcal{P} \times \mathcal{P}$ (i.e., all atoms are *equally* plausible), condition (PD) reduces to condition (H), presented below, which essentially produces a *Hamming-based* total preorder, identical to that proposed by Dalal (1988), and applied to belief update by Forbus (1989).

$$\textbf{(H)} \quad r \preceq_w^{\trianglelefteq} r' \quad \text{iff} \quad |Diff\,(w,r)| \leqslant |Diff\,(w,r')|, \quad \text{where} \quad \trianglelefteq = \mathcal{P} \times \mathcal{P}.$$

Let us, now, proceed to the definition of a PD update operator.

Definition 11 (PD Update Operator)**.** *Let $\trianglelefteq$ be a total preorder over the set of atoms $\mathcal{P}$. A PD update operator is the KM update operator induced from a family $\{\preceq_w^{\trianglelefteq}\}_{w \in \mathbb{M}}$ of PD preorders, by means of condition (F$\diamond$).*

Definition 11 implies that a *single* total preorder $\trianglelefteq$ over atoms induces a *unique* KM update operator, named, PD update operator, which, not only satisfies the KM postulates $(K \diamond 1)$–$(K \diamond 9)$, but also postulate $(K \diamond 10)$. Given that the specification of an arbitrary KM update operator requires a whole *family* of preorders over worlds (one for each world of $\mathbb{M}$), the *significantly lower* specification cost of PD update operators becomes apparent. Note, moreover, that, among all the concrete belief-update operators proposed in the literature (and surveyed by Herzig & Rifi, 1999), only those of Winslett and Forbus satisfy the full set of KM postulates $(K \diamond 1)$–$(K \diamond 9)$; Forbus' operator satisfies postulate $(K \diamond 10)$ as well, whereas, Winslett's operator does not.[18] Herein, we defined a *generalization* of Forbus' proposal, namely, PD update operators, which also satisfy all postulates $(K \diamond 1)$–$(K \diamond 10)$.

To illustrate the practical use of PD update operators, a concrete update-scenario, addressed by this type of operators, is presented subsequently.

Example 4. *Consider a chemical laboratory, in which two dangerous chemical reactions A and B take place. We denote by a the proposition "the chemical reaction A has been triggered", and by b the proposition "the chemical reaction B has been triggered". Initially, it is known that either both reactions have been triggered, or both reactions have not been triggered; hence, the initial state of the world can be described by a theory K, such that $K = Cn\big((a \wedge b) \vee (\neg a \wedge \neg b)\big)$. Furthermore, it is, also, known that chemical reaction B ignites much harder than chemical reaction A.*

Given the hazard involved with the chemical reactions A and B, we instruct a durable robot to go inside the chemical laboratory, and trigger at least one of the two reactions A and B; that is, if both reactions have not been triggered, the robot will try to trigger at least one of them, otherwise, the robot will do nothing. After the (successful) action of the robot, and given the different plausibility of ignition of the two reactions, it is plausible to assume that, in the new state of the world, reaction A —rather than reaction B— has been triggered.

The aforementioned scenario, essentially, constitutes an update-scenario which can easily be addressed with the use of a PD update operator. Specifically, consider the PD update operator $\diamond$, induced from the following total preorder $\trianglelefteq$ over the atoms a and b: $a \lhd b$. Then, the $\diamond$-update of the initial theory K by the sentence $\varphi = a \vee b$ leads to the modified theory

18. It is noteworthy that the KM update operator of Forbus produces the same intuitive results as Winslett's PMA, when applied to the book/magazine example of Subsection 3.3.

$K \diamond \varphi = Cn(a)$, which is a reasonable result representing the new state of the world, given the different plausibility of ignition of the chemical reactions A and B.[19]

It is important to mention that, in the above example, the use of a PD *revision* operator to modify theory K, in the light of the new information φ (which is *consistent* with K), would lead to the new theory $Cn(a, b)$ — since, in this case, revision reduces to *expansion*. This, however, is a counter-intuitive outcome, since, in the new state of the world, the chemical reaction B may have not been ignited; recall that reaction B ignites much harder than reaction A. We note, lastly, that one can devise a variety of real-world dynamically-changing domains that can be encoded with the use of PD update operators.

8.2 Axiomatic Characterization

In case of consistent *complete* theories, parametrized-difference belief revision has a nice axiomatic characterization (Peppas & Williams, 2018). Herein, we show that this characterization can be translated (with slight modifications) in the realm of belief update, in order to characterize PD update operators as well. Let us, first, fix the appropriate notation and terminology.

In this subsection, p, q, z, u denote literals, A, B denote non-empty consistent sets of literals, and K, H denote consistent complete theories. A set of literals shall, occasionally, be treated as a sentence, namely, the *conjunction* of all its literals, leaving it to the context to resolve any ambiguity; thus, for example, in the expression "$Q \cap \mathcal{P}$", Q is a set of literals, whereas, in "$\neg Q$", Q is a sentence of $\mathcal{L}$.

Definition 12 (Peppas & Williams, 2018)**.** *Let $\diamond$ be a KM update operator, and let K be a complete theory of $\mathcal{L}$. For two non-empty sets of literals A, B, we define:*

$$A \lll_K B \quad \textit{iff} \quad A, B \subseteq K \textit{ and } \neg\overline{A} \notin K \diamond (\overline{A} \vee \overline{B}).$$

Intuitively, given a theory K that contains both A and B, $A \lll_K B$ holds whenever it is at least as costly to change (the values of) all literals in B than it is to change (the values of) all literals in A.

Definition 13, below, is a strict version of Definition 12.

Definition 13 (Peppas & Williams, 2018)**.** *Let $\diamond$ be a KM update operator, and let K be a complete theory of $\mathcal{L}$. For two non-empty sets of literals A, B, we define:*

$$A \ll_K B \quad \textit{iff} \quad A \lll_K B \textit{ and } B \not\lll_K A,$$

or, equivalently,

$$A \ll_K B \quad \textit{iff} \quad A, B \subseteq K \textit{ and } \neg\overline{B} \in K \diamond (\overline{A} \vee \overline{B}).$$

19. To see how theory $K \diamond \varphi$ is produced, observe that $[K]$ contains exactly the two worlds $w_1 = \{a, b\}$ and $w_2 = \{\neg a, \neg b\}$. Then, the world w_1 satisfies φ, whereas, the PD preorder $\preceq^{\trianglelefteq}_{w_2}$ is such that: $\{\neg a, \neg b\} \prec^{\trianglelefteq}_{w_2} \{a, \neg b\} \prec^{\trianglelefteq}_{w_2} \{\neg a, b\} \prec^{\trianglelefteq}_{w_2} \{a, b\}$ (where $\prec^{\trianglelefteq}_{w_2}$ denotes the strict part of $\preceq^{\trianglelefteq}_{w_2}$). Hence, $min([\varphi], \preceq^{\trianglelefteq}_{w_2}) = \big\{\{a, \neg b\}\big\} = \big[Cn(w_2) \diamond \varphi\big]$. Therefore, we derive, from condition (F$\diamond$), that $[K \diamond \varphi] = \big\{\{a, b\}, \{a, \neg b\}\big\}$; i.e., $K \diamond \varphi = Cn(a)$.

Lastly, for two literals p, q, $p \lll_K q$ and $p \ll_K q$ denote the abbreviations of $\{p\} \lll_K \{q\}$ and $\{p\} \ll_K \{q\}$, respectively.

Consider, now, the following collection of postulates from Peppas and Williams (2018):

(D1) If $A \lll_K B$, then $|A| \leqslant |B|$.

(D2) If $A \lll_K B$, $p \lll_K q$, and $q \notin B$, then $A \wedge p \lll_K B \wedge q$.

(D3) If $A \lll_K B$, $p \ll_K q$, and $q \notin B$, then $A \wedge p \ll_K B \wedge q$.

(D4) If $A \ll_K B$, $p \in K$, $q \notin B$, and, for all $z \in B$, $z \lll_K q$, then $A \wedge p \ll_K B \wedge q$.

(D5) If $p \lll_K q$, $z \in \{p, \neg p\}$, $u \in \{q, \neg q\}$, and $z, u \in H$, then $z \lll_H u$.

A brief explanation of postulates (D1)–(D5) is presented along the following lines — for more details, the reader is referred to (Peppas & Williams, 2018). Postulate (D1) states that, if one needs to reverse all literals in A or all literals in B, then the update-process never picks the larger set. Postulate (D2) says that, if switching the literals in A is at least as easy as switching the literals in B, and switching p is at least as easy as switching q, then switching A and p together is at least as easy as switching B and q together (provided that q is not already in B). Postulate (D3) is, essentially, a strict version of (D2). Postulate (D4) states that, if reversing A is strictly easier than reversing B, and reversing q is at least as hard as reversing any literal $z \in B$, then, for any literal $p \in K$, changing A and p together is strictly easier than changing B and q together (provided that q is not already in B). Lastly, postulate (D5) says that, if, for a theory K, it is at least as easy to reverse p than it is to reverse q, then this relationship is preserved for any other theory H and any other two literals z, u that share the same atoms with p and q, respectively; for instance, if $p \lll_K q$ and $p, \neg q \in H$, then, according to (D5), $p \lll_H \neg q$.

In view of Remark 1 of Subsection 4.2, and since we are confined to consistent complete theories, the results of Peppas and Williams (2018) entail that, for any KM update operator $\diamond$, the following two statements are true:

- If $\diamond$ is a PD update operator, then $\diamond$ satisfies postulates (D1)–(D5).
- If $\diamond$ satisfies postulates (D1)–(D5), then there is a total preorder $\trianglelefteq$ over the atoms of $\mathcal{P}$, such that $\diamond$ is the KM update operator induced from the family $\{\preceq_w^{\trianglelefteq}\}_{w \in \mathbb{M}}$ of PD preorders, by means of condition (F$\diamond$).

Consequently, due to Peppas and Williams, we can formulate the following *representation theorem*, which provides the axiomatic characterization of the family of PD update operators.

Theorem 11 (Peppas & Williams, 2018)**.** *Let $\diamond$ be a KM update operator. Then, $\diamond$ is a PD update operator iff $\diamond$ satisfies postulates (D1)–(D5).*

By definition, the class of PD update operators is a *proper sub-class* of the class of KM update operators that are induced by total preorders over worlds, since PD update operators are those KM update operators identified by postulates (D1)–(D5).

8.3 PD Update Operators are Relevance-Sensitive

The next theorem shows that PD update operators respect postulates (R1) and (P2$\diamond$).

Theorem 12. *PD update operators satisfy postulates (R1) and (P2$\diamond$).*

Proof. Let $\trianglelefteq$ be a total preorder over $\mathcal{P}$, and let $\diamond$ be the PD update operator induced from the family $\{\preceq_w^{\trianglelefteq}\}_{w \in \mathbb{M}}$ of PD preorders, via condition (F$\diamond$).

For (R1), it suffices to show that $\{\preceq_w^{\trianglelefteq}\}_{w \in \mathbb{M}}$ satisfies condition (SR1). Let w, r, r' be any worlds of $\mathbb{M}$, such that $Diff(w, r) \subset Diff(w, r')$. Clearly then, $Diff(w, r) < Diff(w, r')$. Therefore, we derive, from condition (PD), that $r \prec_w^{\trianglelefteq} r'$, as desired.

For (P2$\diamond$), it suffices to show that condition (SP2$\diamond$) is satisfied. Let K, H be two theories of $\mathcal{L}$, such that, for some sentences $x, y, z \in \mathcal{L}$, $K = Cn(x, y)$, $H = Cn(x, z)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \mathcal{L}_x \cap \mathcal{L}_z = \varnothing$; thus, $\{u_x : u \in [K]\} = \{u_x : u \in [H]\}$. From the fact that the PD preorders assigned to worlds are all induced (via condition (PD)) from the *same* total preorder $\trianglelefteq$, it is not hard to verify that, for any world $w \in [K]$ and any world $w' \in [H]$, such that $w_x = w'_x$, it is true that $\preceq_w^{\trianglelefteq, x} = \preceq_{w'}^{\trianglelefteq, x}$ (where $\preceq_w^{\trianglelefteq, x}$, $\preceq_{w'}^{\trianglelefteq, x}$ denote the x-filterings of $\preceq_w^{\trianglelefteq}$, $\preceq_{w'}^{\trianglelefteq}$, respectively).[20] Therefore, it follows that, for every $w \in [K]$, there is a $w' \in [H]$, such that $\preceq_w^{\trianglelefteq, x} \subseteq \preceq_{w'}^{\trianglelefteq, x}$, as desired. ■

Since there exist total preorders (associated with possible worlds) that satisfy conditions (SR1) and (SP2$\diamond$), and, at the same time, violate condition (PD) (i.e., they are *not* PD preorders), it follows that there exist KM update operators that respect postulates (R1) and (P2$\diamond$), which are *not* PD update operators. Hence, the class of PD update operators is a *proper sub-class* of the class of KM update operators that satisfy (R1) and (P2$\diamond$).

Theorems 7 and 12 imply the following results.

Corollary 1. *PD update operators satisfy postulate (R2).*

Corollary 2. *The conjunction of postulates (R1), (R2) and (P2$\diamond$) is consistent with the KM postulates.*

Theorems 7 and 12, also, imply that all the relevance-sensitive postulates for belief update, presented herein, are entailed by postulates (D1)–(D5). Hence, (D1)–(D5) capture the notion of syntax-relevance for belief update, proposed in the present article.

Furthermore, as the KM update operator of Forbus is a specific PD update operator, the following corollary is, immediately, derived from Theorem 12 and Corollary 1.

Corollary 3. *Forbus' KM update operator satisfies postulates (R1), (R2) and (P2$\diamond$).*

Figure 3 depicts the whole class of KM update operators, induced by arbitrary partial preorders over worlds, its proper sub-classes circumscribed by the relevance-sensitive postulates (R1) (equivalently, (S1) or (P1$\diamond$)), (R2) (equivalently, (S2)), and (P2$\diamond$), as well as the place of PD update operators and Winslett's operator $\diamond_W$ relative to those classes. Figure 3, also, depicts the corresponding semantic characterizations of the aforementioned KM update operators. Recall that, by hypothesis, the relevance-sensitive postulates studied herein constrain KM update operators that are induced by *total* preorders over worlds.

20. This conclusion can, also, be reached through Theorem 7 by Aravanis et al. (2021), combined with the fact that, for a world $w \in \mathbb{M}$, the definitions of the two total preorders $\preceq_w$ and $\preceq_{Cn(w)}$ coincide.

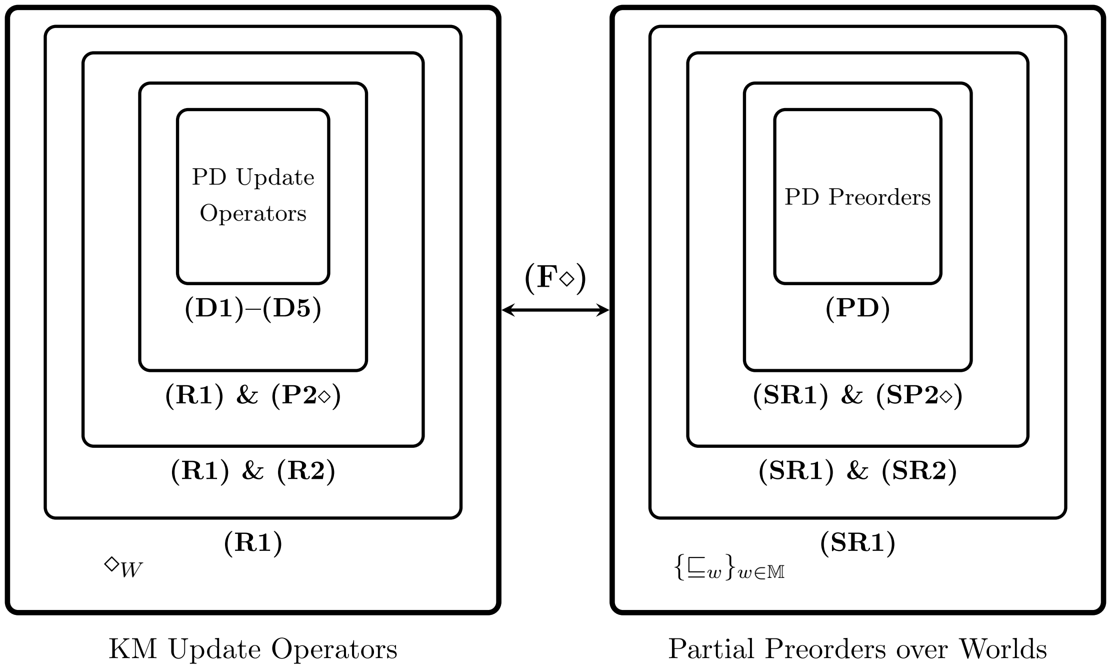

Figure 3: The class of KM update operators, induced by arbitrary partial preorders over worlds, its proper sub-classes circumscribed by the relevance-sensitive postulates (R1), (R2) and (P2$\diamond$), the place of PD update operators and Winslett's operator $\diamond_W$ relative to those classes, and the corresponding semantic characterizations; by hypothesis, the relevance-sensitive postulates studied herein constrain KM update operators that are induced by total preorders over worlds.

As stated in the Introduction, the only work on relevance-sensitive belief update that bears on the approach described herein is that by Perrussel et al. (2012). Recall, however, that the authors of that work define belief update in terms of *prime implicants* (PI). Although we are not aiming at focusing on PI-based update operators —as they significantly deviate in nature from KM update operators—, we note that Forbus' KM update operator, which, as stated in Corollary 3 respects the notion of relevance formalized in this article, is not equivalent to its PI-based counterpart, defined by Perrussel et al. (2012), which respects the notion of relevance formalized by its originators. This remark —which is exemplified through the following example— suggests, in turn, that PI-based relevance acts *differently* in relation to relevance presented herein.

Example 5. *Let us denote by $\diamond_F$ Forbus' KM update operator, and by $\diamond_F^{PI}$ its PI-based counterpart, defined by Perrussel et al. (2012). Assume that $\mathcal{P} = \{a, b, c, d, e\}$, and let $K = Cn\big((b \wedge c \wedge e) \vee (d \wedge e)\big)$ and $\varphi = (a \wedge b \wedge \neg c) \vee (\neg a \wedge \neg b \wedge \neg c)$. Then, we have that $[K \diamond_F \varphi] = \big\{\{a, b, \neg c, d, e\}, \{\neg a, \neg b, \neg c, d, e\}, \{a, b, \neg c, \neg d, e\}, \{\neg a, \neg b, \neg c, \neg d, e\}\big\}$, whereas, $[K \diamond_F^{PI} \varphi] = \big\{\{a, b, \neg c, d, e\}, \{\neg a, \neg b, \neg c, d, e\}, \{a, b, \neg c, \neg d, e\}\big\}$ — refer to Example 3 in (Perrussel et al., 2012) for details on how the set of worlds $[K \diamond_F^{PI} \varphi]$ is produced. Therefore, $[K \diamond_F \varphi] \neq [K \diamond_F^{PI} \varphi]$, thus, $K \diamond_F \varphi \neq K \diamond_F^{PI} \varphi$.*

9. Conclusion

In this article, we pointed out that the widely-accepted KM update operators are too liberal in their treatment of the notion of relevance. In response to this weakness, we showed that a recast of (the weak version of) Parikh's relevance-sensitive axiom (P), in the realm of belief update, suffices to exclude unreasonable update-policies, by strictly strengthening Katsuno and Mendelzon's framework. Axiom (P) for belief update was in detail investigated, both axiomatically and semantically. Specifically, it was proved that the weak version of (P) for belief update is equivalent to postulate (R1), whereas, the strong version of (P) strictly implies (but it is not equivalent to) postulate (R2). Both postulates (R1) and (R2) encode relevance at the possible-worlds level, according to which each possible world is locally modified, in the light of new information. Postulates (R1)–(R2) turned out to be equivalent to postulates (S1)–(S2), respectively, which encode relevance at the sentential level, in the sense that they consider sentences to be the building blocks of relevance. Furthermore, we concretely demonstrated that a slight adjustment of the weak version of axiom (P) for belief update, which incorporates causal information, can be regarded as (at least a partial) solution to the frame, ramification and qualification problems, encountered in dynamically-changing worlds. As a last contribution, a whole new family of concrete KM update operators, named PD update operators, was (axiomatically and semantically) introduced. PD update operators constitute a natural generalization of Forbus' KM update operator, are compactly-specified and relevance-sensitive; hence, they are ideal candidates for real-world implementations.

This article opens up many interesting avenues for future work. Firstly, the notion of relevance both at the possible-worlds and the sentential level —encoded, in the realm of belief update, in postulates (R1)–(R2) and (S1)–(S2), respectively— is very natural and, even, useful in the realm of belief revision as well; the study of this issue, and its interrelations with Parikh's axiom for belief revision, constitutes a promising research topic. Also of much interest would be a deeper investigation of how relevance-sensitive belief update can be applied to address the frame, ramification and qualification problems, in more practical terms. Lastly, an appealing direction of research would be the study of the relation between relevance-sensitive KM update operators and relevance-sensitive PI-based update operators; for instance, can we impose certain conditions so that these two types of update operators produce the same results?

Acknowledgments

This research is co-financed by Greece and the European Union (European Social Fund – ESF) through the Operational Programme "Human Resources Development, Education and Lifelong Learning" in the context of the project "Reinforcement of Postdoctoral Researchers – 2nd Cycle" (MIS-5033021), implemented by the State Scholarships Foundation (IKY).

The author is truly grateful to Pavlos Peppas for many fruitful comments on this work, as well as to Yannis Stamatiou and the anonymous reviewers for their detailed and constructive suggestions on a previous version of this article.

References

Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, *50*(2), 510–530.

Aravanis, T., Peppas, P., & Williams, M.-A. (2021). An investigation of parametrized difference revision operators. *Annals of Mathematics and Artificial Intelligence*, *89*, 7–28.

Aravanis, T. I. (2019). *Relevance and Knowledge Dynamics for Intelligent Agents*. Ph.D. thesis, Department of Business Administration, School of Economics & Business, University of Patras, Patras, Greece.

Aravanis, T. I., Peppas, P., & Williams, M.-A. (2019). Full characterization of Parikh's relevance-sensitive axiom for belief revision. *Journal of Artificial Intelligence Research*, *66*, 765–792.

Chopra, S., & Parikh, R. (2000). Relevance sensitive belief structures. *Annals of Mathematics and Artificial Intelligence*, *28*(1–4), 259–285.

Dalal, M. (1988). 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, California.

Delgrande, J., & Peppas, P. (2018). Incorporating relevance in epistemic states in belief revision. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning, KR 2018*, pp. 230–239.

Finger, J. J. (1987). *Exploiting constraints in design synthesis*. Ph.D. thesis, Stanford University, Stanford, United States.

Forbus, K. D. (1989). Introducing actions into qualitative simulation. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI 1989)*, pp. 1273–1278. Morgan Kaufmann Publishers Inc.

Gärdenfors, P. (1988). *Knowledge in Flux – Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Massachusetts.

Gärdenfors, P. (1990). Belief revision and relevance. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pp. 349–365.

Gärdenfors, P., & Makinson, D. (1988). Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning About Knowledge (TARK 1988)*, pp. 83–95, Pacific Grove, California. Morgan Kaufmann.

Hansson, S. O., & Wassermann, R. (2002). Local change. *Studia Logica: An International Journal for Symbolic Logic*, *70*(1), 49–76.

Herzig, A., & Rifi, O. (1999). Propositional belief base update and minimal change. *Artificial Intelligence*, *115*, 107–138.

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

Katsuno, H., & Mendelzon, A. (1992). On the difference between updating a knowledge base and revising it. In Gärdenfors, P. (Ed.), *Belief Revision*, pp. 183–203. Cambridge University Press.

Keller, A., & Winslett, M. (1985). On the use of an extended relational model to handle changing incomplete information. *IEEE Transactions on Software Engineering*, *11*, 620–633.

Kern-Isberner, G., & Brewka, G. (2017). Strong syntax splitting for iterated belief revision. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*, pp. 1131–1137.

Kourousias, G., & Makinson, D. (2007). Parallel interpolation, splitting, and relevance in belief change. *Journal of Symbolic Logic*, *72*(3), 994–1002.

Makinson, D. (2009). Propositional relevance through letter-sharing. *Journal of Applied Logic*, *7*, 377–387.

Makinson, D., & Kourousias, G. (2006). Respecting relevance in belief change. *Análisis Filosófico*, *26*(1), 53–61.

McCain, N., & Turner, H. (1995). A causal theory of ramifications and qualifications. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI 1995)*, pp. 1978–1984.

McCarthy, J. (1977). Epistemological problems of Artificial Intelligence. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI 1977)*, pp. 1038–1044.

McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of Artificial Intelligence. In Meltzer, B., & Michie, D. (Eds.), *Machine Intelligence 4*, pp. 463–502. Edinburgh University Press.

Pagnucco, M., & Peppas, P. (2001). Causality and minimal change demystified. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI 2001)*, pp. 125–135.

Parikh, R. (1999). 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.

Peppas, P. (1993). *Belief change and reasoning about action: An axiomatic approach to modelling dynamic worlds and the connection to the logic of theory change*. Ph.D. thesis, Basser Department of Computer Science, University of Sydney, Sydney.

Peppas, P. (2008). Belief revision. In van Harmelen, F., Lifschitz, V., & Porter, B. (Eds.), *Handbook of Knowledge Representation*, pp. 317–359. Elsevier Science.

Peppas, P., Chopra, S., & Foo, N. (2004). Distance semantics for relevance-sensitive belief revision. In *Proceedings of the 9th International Conference on the Principles of Knowledge Representation and Reasoning (KR 2004)*, pp. 319–328. AAAI Press.

Peppas, P., Foo, N., & Nayak, A. (2000). Measuring similarity in belief revision. *Journal of Logic and Computation*, *10*, 603–619.

Peppas, P., Nayak, A. C., Pagnucco, M., Foo, N., Kwok, R., & Prokopenko, M. (1996). Revision vs. update: Taking a closer look. In Wahlster, W. (Ed.), *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI 1996)*, pp. 95–99, Budapest, Hungary. John Wiley & Sons, Ltd.

Peppas, P., & Williams, M.-A. (2016). Kinetic consistency and relevance in belief revision. In *Proceedings of the 15th European Conference on Logics in Artificial Intelligence (JELIA 2016)*, pp. 401–414. Springer International Publishing.

Peppas, P., & Williams, M.-A. (2018). 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, California.

Peppas, P., Williams, M.-A., Chopra, S., & Foo, N. (2015). Relevance in belief revision. *Artificial Intelligence*, *229*, 126–138.

Perrussel, L., Marchi, J., Thévenin, J.-M., & Zhang, D. (2012). Relevant minimal change in belief update. In *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA 2012)*, pp. 333–345. Springer, Berlin, Heidelberg.

Rott, H. (2000). Two dogmas of belief revision. *Journal of Philosophy*, *97*(9), 503–522.

Thielscher, M. (1996). Causality and the qualification problem. In *Proceedings of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR 1996)*, pp. 51–62. San Francisco, CA: Morgan Kaufmann Publishers.

Thielscher, M. (2001). The qualification problem: A solution to the problem of anomalous models. *Artificial Intelligence*, *131*, 1–37.

Wassermann, R. (2001a). Local diagnosis. *Journal of Applied Non-Classical Logics*, *11*(1–2), 107–129.

Wassermann, R. (2001b). On structured belief bases. In Williams, M.-A., & Rott, H. (Eds.), *Frontiers in Belief Revision*, Applied Logic Series, pp. 349–367. Springer Netherlands.

Winslett, M. (1988). Reasoning about action using a possible models approach. In *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence (AAAI 1988)*, pp. 89–93. The AAAI Press, Menlo Park, California.