Incompatibilities between Iterated and Relevance-Sensitive Belief Revision

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

Mary-Anne Williams   MARY-ANNE.WILLIAMS@UTS.EDU.AU
Centre for Artificial Intelligence
Faculty of Engineering and Information Technology
University of Technology Sydney
NSW 2007, Australia

Abstract

The AGM paradigm for belief change, as originally introduced by Alchourrón, Gärdenfors and Makinson, lacks any guidelines for the process of iterated revision. One of the most influential work addressing this problem is Darwiche and Pearl's approach (DP approach, for short), which, despite its well-documented shortcomings, remains to this date the most dominant. In this article, we make further observations on the DP approach. In particular, we prove that the DP postulates are, in a strong sense, inconsistent with Parikh's relevance-sensitive axiom (P), extending previous initial conflicts. Immediate consequences of this result are that an entire class of intuitive revision operators, which includes Dalal's operator, violates the DP postulates, as well as that the Independence postulate and Spohn's conditionalization are inconsistent with axiom (P). The whole study, essentially, indicates that two fundamental aspects of the revision process, namely, iteration and relevance, are in deep conflict, and opens the discussion for a potential reconciliation towards a comprehensive formal framework for knowledge dynamics.

1. Introduction

Belief Revision captures the dynamics of knowledge (Gärdenfors, 1988; Fermé & Hansson, 2018). This research field was established by the seminal work of Alchourrón, Gärdenfors, and Makinson (1985), from which the so-called *AGM paradigm* arose. Within the AGM paradigm, a set of eight postulates, known as the *AGM postulates for revision*, characterizes *rational* belief revision, and the (revision) functions satisfying these postulates are known as *AGM revision functions.*[1]

One of the main shortcomings of the original AGM paradigm is its lack of any guidelines for *iterated* belief revision. Perhaps the most influential work addressing this problem, within the confines of the AGM paradigm, is Darwiche and Pearl's approach; for short,

---

1. The capitalized term *Belief Revision* shall be used to refer to the research area, whereas, the same term in lower-case letters, i.e., *belief revision*, shall be used to refer to the process.

DP approach (Darwiche & Pearl, 1994). In the DP approach, four new postulates were introduced, namely, the *DP postulates*, to regulate the process of iterated revision.

However, as originally introduced by Darwiche and Pearl (1994), not all the DP postulates were consistent with the AGM postulates for revision (Lehmann, 1995). The inconsistency was later resolved by the same authors (1997), by defining revision functions to operate on *belief states* (also referred to as *epistemic states*), rather than on belief sets. With the switch to belief states, Darwiche and Pearl modified the AGM postulates for revision accordingly; we shall refer to the functions satisfying these modified postulates as *modified AGM revision functions*.

Even so, the DP approach remains controversial. The criticism, typically, comes in the form of convincing *counter-examples* in which the DP postulates produce unintuitive results — see the work of Stalnaker (2009, pp. 205–206) for counter-examples on the first DP postulate, the works of Konieczny and Pérez (2000, p. 352) and Stalnaker (2009, pp. 206–208) for counter-examples on the second DP postulate, and the work of Hansson (2016, pp. 437–438) for counter-examples on the third and fourth DP postulates.[2] Another form of criticism constitutes the result by Peppas, Fotinopoulos, and Seremetaki (2008), that points out a first conflict between each one of the DP postulates and relevance-sensitive axiom (P), introduced by Parikh (1999) as an attempt to capture *relevant* change, supplementing the AGM postulates for revision.

In this article, we make further observations on the DP approach. In particular, the following contributions are made:

- We provide *fully-fledged incompatibility* results between *relevance-sensitive* and *iterated* belief revision, extending the result of Peppas et al. (2008) for any finite, non-empty set of propositional variables, as well as for any splittable or (non-trivially) confined theory of the language.[3] Hence, the possibility for modified AGM revision functions of *any type* to satisfy both axiom (P) and the DP postulates is excluded.

- The aforementioned incompatibility results imply significant corollaries for existing belief-change proposals. *Parametrized Difference* revision operators (Peppas & Williams, 2016, 2018), and their specialization *Dalal's* operator (Dalal, 1988), are incompatible with the DP approach. Moreover, the *Independence* postulate (Jin & Thielscher, 2007) (that remedies problems with the DP postulates) and *Spohn's conditionalization* (Spohn, 1988) (restricted to revision scenarios) are incompatible with axiom (P).

The remainder of this article is structured as follows. The next section introduces the basic terminology and notation. Following that, Section 3 provides a prerequisite, concrete background for presenting the outcomes of the study. Specifically, Subsection 3.1 presents a brief overview of the AGM paradigm, followed by Subsection 3.2 which introduces the problem of iterated belief revision. Subsections 3.3 and 3.4 present the DP approach and

---

2. The considerations on the second DP postulate by Lehmann (1995) and Jin and Thielscher (2007) are worth noting as well.

3. The main incompatibility result of Peppas et al. (2008) has been proved assuming a set of propositional variables with fixed cardinality (in particular, of only three variables), utilizing a particular theory of the language; hence, it is confined to *only one type* of modified AGM revision functions.

Spohn's conditionalization, respectively, and Subsection 3.5 introduces Parikh's notion of relevance. Against this background, Section 4 points out the main incompatibility results between relevance and iteration, Section 5 proceeds to their consequences, whereas, Section 6 conducts a comprehensive discussion on the established incompatibilities. The article closes with some concluding remarks indicating future research avenues.

2. Formal Preliminaries

This section fixes the basic notation and terminology that shall be used throughout the article.

2.1 Language

Throughout this article, we shall be working with a propositional language $\mathcal{L}$, built over *finitely many* propositional variables, using the standard Boolean connectives $\wedge$ (conjunction), $\vee$ (disjunction), $\rightarrow$ (implication), $\leftrightarrow$ (equivalence), $\neg$ (negation), the special symbol $\perp$ (arbitrary contradiction), and governed by *classical propositional logic*. The finite, non-empty set of all propositional variables is denoted by $\mathcal{P}$.

2.2 Sentences and Belief Sets

A sentence φ of $\mathcal{L}$ is *contingent* iff $\nvDash \varphi$ and $\nvDash \neg\varphi$. For a set of sentences Γ of $\mathcal{L}$, $Cn(\Gamma)$ denotes the set of all logical consequences of Γ; i.e.,

$$Cn(\Gamma) = \big\{\varphi \in \mathcal{L} : \Gamma \models \varphi\big\}.$$

For 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 two sentences φ, ψ of $\mathcal{L}$, we shall write $\varphi \equiv \psi$ iff $Cn(\varphi) = Cn(\psi)$.

An agent's set of beliefs will be modelled by a *theory*, also referred to as a *belief set*. A theory K of $\mathcal{L}$ is any set of sentences of $\mathcal{L}$ closed under logical consequence; thus,

$$K = Cn(K).$$

The set of all *consistent* theories of $\mathcal{L}$ is denoted by $\mathbb{K}$.

As we shall, subsequently, introduce formal properties of revision functions, let us first define the simpler operation of *expansion*. Accordingly, for a theory K and a sentence φ of $\mathcal{L}$, the expansion of K by φ, denoted by $K + \varphi$, is specified as follows:

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

2.3 Literals and Possible Worlds

A *literal* is a propositional variable $p \in \mathcal{P}$ or its negation. Sometimes, we shall treat a set of literals as a sentence, i.e., 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.

A *possible world* (abbrev. *world*) r is a consistent set of literals, such that, for any propositional variable $p \in \mathcal{P}$, either $p \in r$ or $\neg p \in r$. The set of all possible worlds is

denoted by $\mathbb{M}$. For a sentence or set of sentences φ of $\mathcal{L}$, $[\varphi]$ is the set of worlds at which φ is true. For a set of worlds $V \subseteq \mathbb{M}$, we denote by $th(V)$ the set of sentences (theory) satisfied by all worlds in V. When $V = \varnothing$, then vacuously $th(V) = \mathcal{L}$.

2.4 Preorders

A *preorder* over a set V is any reflexive, transitive binary relation on V. A preorder $\preceq$ is called *total* iff any two elements of V are comparable with respect to $\preceq$; i.e., for all $r, r' \in V$, $r \preceq r'$ or $r' \preceq r$. The strict part of the preorder $\preceq$ is denoted by $\prec$; i.e., $r \prec r'$ iff $r \preceq r'$ and $r' \not\preceq r$. The symmetric part of the preorder $\preceq$ is denoted by $\approx$; i.e., $r \approx r'$ iff $r \preceq r'$ and $r' \preceq r$.

For any $X \subseteq V$, by $min(X, \preceq)$ we denote the set of all $\preceq$-minimal elements of X; namely,

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

When the set X contains (natural) numbers, we shall simply write $min(X)$ to denote the minimum number in X.

2.5 Sublanguages

Let Q be a subset of the set of propositional variables $\mathcal{P}$. We denote by $\mathcal{L}^Q$ the *sublanguage* of $\mathcal{L}$ defined over Q. In the limiting case where Q is empty, we define $\mathcal{L}^Q$ to be the language generated by $\bot$, and the Boolean connectives.

For a sentence x of $\mathcal{L}$, we denote by $\mathcal{L}_x$ the unique *minimal* language within which x can be expressed. Essentially, $\mathcal{L}_x$ contains a sentence that is logically equivalent to x, and, moreover, no *proper* sublanguage of $\mathcal{L}_x$ contains such a sentence. If x is inconsistent or a tautology, we take $\mathcal{L}_x$ to be $\mathcal{L}^\varnothing$.

Lastly, $\overline{\mathcal{L}_x}$ denotes the complement language of $\mathcal{L}_x$; namely, the language built from the propositional variables that do not appear in $\mathcal{L}_x$.

3. Background

This section establishes a breeding ground for presenting the outcomes of the study. We note that, for ease of exposition (and only), throughout this article, we shall consider only *consistent* belief sets and *contingent* epistemic input.

3.1 The AGM Paradigm

Within the AGM paradigm, belief revision is modelled as a (binary) function $*$, mapping a consistent theory K and a contingent sentence φ to a consistent (revised) theory $K * \varphi$. The *AGM postulates for revision*, presented subsequently, circumscribe the territory of *rational* revision functions, the so-called *AGM revision functions*.

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

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

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

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

(**K ∗ 5**)   If φ is consistent, then $K * \varphi$ is also consistent.

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

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

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

It turns out that any AGM revision function can be *constructed* with the use of a special kind of total preorders over all possible worlds, called *faithful preorders* (Katsuno & Mendelzon, 1991).

Definition 1 (Faithful Assignment over Belief Sets, Katsuno & Mendelzon, 1991)**.** *A function that maps each belief set K to a total preorder $\preceq_K$ over the possible worlds of $\mathbb{M}$ is said to be a faithful assignment over belief sets precisely when the following two conditions hold, for any possible worlds $w_1, w_2 \in \mathbb{M}$:*[4]

- *If $w_1 \in [K]$, then $w_1 \preceq_K w_2$.*

- *If $w_1 \in [K]$ and $w_2 \notin [K]$, then $w_1 \prec_K w_2$.*

The total preorder $\preceq_K$ of Definition 1 is typically referred to as a preorder *faithful* to K. Intuitively, $w_1 \preceq_K w_2$ states that the world w_1 is at least as plausible (with respect to K) as the world w_2.

Katsuno and Mendelzon (1991) proceed, then, to the establishment of the next representation theorem.

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

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

3.2 The Problem of Iterated Belief Revision

Despite the immense success of the AGM paradigm, there is one aspect of the belief-revision process that was left unattended; that is, *iterated* belief revision. More precisely, while the AGM postulates for revision appear to be sound constraints for rational *one-step* transitions between belief sets K and $K * \varphi$, they say very little about transitions involving multiple

---

4. This definition of faithful assignment is slightly different from the one presented by Katsuno and Mendelzon (1991). Nevertheless, since we deal with a language built over finitely many propositional variables, the difference is immaterial.

steps; e.g., a transition from K to $K * \varphi$ to $(K * \varphi) * \psi$. In fact, the only constraint they impose relates to the special case of revising sequentially the initial belief set K by two sentences φ, ψ, such that $\neg\psi \notin K * \varphi$. For this special case, the AGM postulates for revision define $(K * \varphi) * \psi = K * (\varphi \wedge \psi)$.[5]

In all other cases, $(K * 1)$–$(K * 8)$ remain silent. This, often, leads to counter-intuitive results. The following example, borrowed from Darwiche and Pearl (1997), illustrates the problem.

Example 1 (Darwiche & Pearl, 1997)**.** "We are introduced to a lady X who sounds smart and looks rich, so we believe that X is smart and X is rich. Moreover, since we profess to no prejudice, we also maintain that X is smart, even if found to be poor and, conversely, X is rich, even if found to be not smart. Now, we obtain some evidence that X is, in fact, not smart, and we remain of course convinced that X is rich. Still, it would be strange for us to say, 'If the evidence turns out false, and X turns out smart after all, we would no longer believe that X is rich'. If we currently believe X is smart and rich, then evidence first refuting then supporting that X is smart should not, in any way, change our opinion about X being rich. Strangely, the AGM postulates (for revision) do permit such a change of opinion."

According to the above scenario, our initial belief set is the logical closure of $\{s, r\}$, where s is a propositional variable that stands for "X is smart", and r is a variable that stands for "X is rich". Moreover, it is assumed that $Cn(s, r) * \neg s = Cn(\neg s, r)$ and that $Cn(s, r) * \neg r = Cn(s, \neg r)$. Darwiche and Pearl argue that, in this case, it would be counter-intuitive to have $r \notin (Cn(s, r) * \neg s) * s$. Yet, as shown by Darwiche and Pearl (1997), such a possibility is not excluded from the AGM postulates for revision.

This, along with many other concrete examples in the literature of Belief Revision, has provided compelling evidence that the AGM postulates for revision need to be augmented with further constraints in order to properly regulate iterated belief revision.

More formally, in the AGM paradigm, the extra-logical structure of a faithful preorder serves as a guide to performing belief change. Notice, however, that this structure tells us what our new *belief set* should be. Nevertheless, in order to process the next epistemic input, we need not only a belief set, but also a *new* extra-logical structure associated with it (Figure 1). Presumably, this new structure would be a *rational* offspring of the old structure and the new input. Even if the former is not fully determined by the two latter, it should at least be *constrained*. The problem of iterated belief revision is the problem of formulating constraints that capture the dynamics of the extra-logical structures used to encode one-step revision policies.

Many researchers have raised to the challenge of addressing the problem of iterated belief revision, making important contributions — see, for example, the works of Darwiche and Pearl (1994, 1997), Nayak (1994), Konieczny and Pérez (2000), Hunter and Delgrande (2007), Booth and Meyer (2011), Shapiro, Pagnucco, Lespérance, and Levesque (2000), Hunter and Delgrande (2005), Delgrande and Schaub (2003), Delgrande, Dubois, and Lang (2006), Benferhat, Kaci, Le Berre, and Williams (2004), Areces and Becher (2001), Shapiro, Pagnucco, Lespérance, and Levesque (2011), Zhang (2004), Chopra, Ghose, Meyer, and

---

5. This follows immediately from postulates $(K * 3)$, $(K * 4)$, $(K * 7)$ and $(K * 8)$.
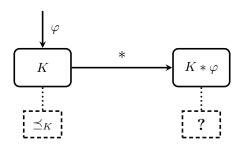
Figure 1: The problem of iterated belief revision.

Wong (2008), Spohn (1988). For a recent survey on the problem, the interested reader may refer to (Peppas, 2014).

In the course of this article, two well-established proposals shall be extensively studied. The first is *Darwiche and Pearl's approach* (alias, *DP approach*), a work that complies with the assumptions of the AGM paradigm (Darwiche & Pearl, 1994, 1997), and the second is *Spohn's conditionalization*, a *quantitative* approach which lies outside the confines of the AGM paradigm since, as we shall see, it allows some form of "enriched" epistemic input (Spohn, 1988).

3.3 The DP Approach

Despite the significant efforts of the research community, there is still no widely accepted formal model for iterated belief revision within the confines of the AGM paradigm. This subsection discusses the *DP approach*, originally proposed by Darwiche and Pearl (1994), which is perhaps one of the most influential proposals. The cornerstone of that work is the introduction of four additional postulates that regulate successive revisions, the so-called *DP postulates*, supplementing the AGM postulates for revision.

3.3.1 THE MODIFIED AGM POSTULATES FOR REVISION

The DP approach was later modified by its originators (1997), partially in response to a problem identified by Lehmann (1995).[6] In the modified version, Darwiche and Pearl defined revision functions to operate on *belief states*, rather than on belief sets.[7] Accordingly, we reserve the symbol $*$ denoting revision functions as originally defined in the AGM paradigm, and for the modified revision functions of the DP approach that operate on belief states, we shall use the symbol $\circ$. Hence, $\circ$ is a function that maps a belief state S and a sentence φ of $\mathcal{L}$ to a revised (new) belief state $S \circ \varphi$ (see Figure 2).

Darwiche and Pearl did not fully specify what *exactly* they consider a belief state to be. What was clear, however, from their exposition is that to each belief state is associated a belief set, as well as an ordering —such as a total preorder over all possible worlds (or any other equivalent structure)— that encodes preference information.

---

6. Lehmann showed that the second DP postulate, as originally stated, was inconsistent with the AGM postulates for revision.
7. Hence, they slightly deviated from the AGM paradigm.

Figure 2: Belief-state revision within the DP approach.

Darwiche and Pearl reformulated the AGM postulates for revision accordingly, in order to reflect the shift from belief sets to belief states:[8]

($\mathbf{S} \circ \mathbf{1}$) $S \circ \varphi$ is a belief state.

($\mathbf{S} \circ \mathbf{2}$) $\varphi \in B(S \circ \varphi)$.

($\mathbf{S} \circ \mathbf{3}$) $B(S \circ \varphi) \subseteq B(S) + \varphi$.

($\mathbf{S} \circ \mathbf{4}$) If $\neg\varphi \notin B(S)$, then $B(S) + \varphi \subseteq B(S \circ \varphi)$.

($\mathbf{S} \circ \mathbf{5}$) If φ is consistent, then $B(S \circ \varphi)$ is also consistent.

($\mathbf{S} \circ \mathbf{6}$) If $\varphi \equiv \psi$, then $B(S \circ \varphi) = B(S \circ \psi)$.

($\mathbf{S} \circ \mathbf{7}$) $B\big(S \circ (\varphi \wedge \psi)\big) \subseteq B(S \circ \varphi) + \psi$.

($\mathbf{S} \circ \mathbf{8}$) If $\neg\psi \notin B(S \circ \varphi)$, then $B(S \circ \varphi) + \psi \subseteq B\big(S \circ (\varphi \wedge \psi)\big)$.

We shall call a function $\circ$ *modified AGM revision function* iff it satisfies the modified AGM postulates for revision $(S\circ 1)$–$(S\circ 8)$, to differentiate it from an AGM revision function $*$ of the original AGM paradigm.

Following Katsuno and Mendelzon's approach, Darwiche and Pearl provided a representation result for the modified AGM postulates for revision.

Definition 2 (Faithful Assignment over Belief States, Darwiche & Pearl, 1997). *A function that maps each belief state S to a total preorder $\preceq_S$ over the possible worlds of $\mathbb{M}$ is said to be a faithful assignment over belief states precisely when the following two conditions hold, for any possible worlds $w_1, w_2 \in \mathbb{M}$:*

- *If $w_1 \in \big[B(S)\big]$, then $w_1 \preceq_S w_2$.*

- *If $w_1 \in \big[B(S)\big]$ and $w_2 \notin \big[B(S)\big]$, then $w_1 \prec_S w_2$.*

The belief set associated with the belief state S, denoted by $B(S)$, can be extracted by $\preceq_S$ itself, as the $\preceq_S$-minimal worlds determine $B(S)$; i.e., $B(S) = min(\mathbb{M}, \preceq_S)$. Note that Definition 2 implies that a belief state is a *richer* structure than a belief set, as two *different*

---

8. Postulates $(S \circ 1)$–$(S \circ 8)$, as well as the DP postulates that follow, have been rephrased in the AGM notation. Note that the original version of $(S \circ 6)$ refers to the equality of two belief states — see postulate (R*4) by Darwiche and Pearl (1997). We adopt the approach of Jin and Thielscher (2007) and reformulate this postulate for the sake of precision, as Darwiche and Pearl have not explicitly defined the notion of a belief state, let alone the equality of two belief states.

belief states may be associated with the *same* belief set; this feature of belief states is crucial for resolving the original inconsistency of the DP approach with the AGM postulates for revision, identified in (Lehmann, 1995).

The next theorem, which is analogous to Theorem 1, characterizes the class of (modified) revision functions induced from faithful assignments over belief states.

Theorem 2 (Darwiche & Pearl, 1997). *A modified revision function $\circ$ satisfies postulates $(S \circ 1)$–$(S \circ 8)$ iff there exists a (unique) faithful assignment over belief states that maps each belief state S to a total preorder $\preceq_S$ over $\mathbb{M}$, such that, for any $\varphi \in \mathcal{L}$:*

$$\textbf{(F}\circ\textbf{)} \quad B(S \circ \varphi) = th\Big(min([\varphi], \preceq_S)\Big).$$

Remark 1. A crucial difference between the original AGM paradigm and the DP approach is that, in the latter, a revision function $\circ$ can affect a belief set K in different ways for the *same* epistemic input φ. To see this, consider two distinct belief states S and S', which are associated with the *same* belief set K; i.e., $B(S) = B(S') = K$. Since Darwiche and Pearl define revision functions to operate on *belief states* (rather than belief sets), it is possible that $S \circ \varphi \neq S' \circ \varphi$. That is to say, a single agent revises the belief set K in different ways, depending on whether her belief state is S or S', even though she revises by the same epistemic input φ. This is not possible in the original AGM paradigm, where revision functions are defined as functions mapping belief sets and sentences to belief sets.

3.3.2 THE DP POSTULATES

To regulate successive revisions, Darwiche and Pearl introduced four additional postulates, known as the *DP postulates* and presented below.

(DP1) If $\varphi \models \psi$, then $B\big((S \circ \psi) \circ \varphi\big) = B(S \circ \varphi)$.

(DP2) If $\varphi \models \neg\psi$, then $B\big((S \circ \psi) \circ \varphi\big) = B(S \circ \varphi)$.

(DP3) If $\psi \in B(S \circ \varphi)$, then $\psi \in B\big((S \circ \psi) \circ \varphi\big)$.

(DP4) If $\neg\psi \notin B(S \circ \varphi)$, then $\neg\psi \notin B\big((S \circ \psi) \circ \varphi\big)$.

The analysis regarding (DP1)–(DP4) that follows is borrowed from Peppas (2014). Postulate (DP1) says that, if the subsequent evidence φ is logically stronger than the initial evidence ψ, then φ overrides whatever changes ψ may have made. (DP2) says that, if two contradictory pieces of evidence arrive sequentially one after the other, it is the latter that will prevail. Clearly, according to (DP1) and (DP2), the latter piece of evidence φ prevails in a very strong sense. (DP3) says that, if revising S by φ causes ψ to be accepted in the new belief state, then revising first by ψ and then by φ cannot possibly block the acceptance of ψ. Lastly, (DP4) says that, if the revision of S by φ does not cause the acceptance of $\neg\psi$, then this should still be the case if S is first revised by ψ before revised by φ.

Postulates (DP1)–(DP4) have a nice characterization in terms of total preorders over worlds. In particular, Darwiche and Pearl proved that there is a *one-to-one correspondence*

between (DP1)–(DP4) and the following constraints (R1)–(R4) on possible worlds, respectively, associating the preorder $\preceq_S$ of the initial belief state S with the preorder $\preceq_{S \circ \varphi}$ of the belief state that results from the revision of S by φ.

(R1) If $r, r' \in [\varphi]$, then $r \preceq_S r'$ iff $r \preceq_{S \circ \varphi} r'$.

(R2) If $r, r' \in [\neg\varphi]$, then $r \preceq_S r'$ iff $r \preceq_{S \circ \varphi} r'$.

(R3) If $r \in [\varphi]$ and $r' \in [\neg\varphi]$, then $r \prec_S r'$ implies $r \prec_{S \circ \varphi} r'$.

(R4) If $r \in [\varphi]$ and $r' \in [\neg\varphi]$, then $r \preceq_S r'$ implies $r \preceq_{S \circ \varphi} r'$.

When S is revised by φ, conditions (R1) and (R2) require not to change the relative plausibility ordering of φ-worlds and $\neg\varphi$-worlds, respectively. Conditions (R3) and (R4) require that, if a φ-world r is (strictly) more plausible than a $\neg\varphi$-world r', then r continues to be (strictly) more plausible than r' (Jin & Thielscher, 2007).

3.3.3 The Independence Postulate

Jin and Thielscher (2005, 2007) argue that (DP1)–(DP4) are too permissive in certain cases. Interestingly, the example they use to make their case is borrowed from the work of Darwiche and Pearl (1997).

Example 2 (Darwiche & Pearl, 1997). "We encounter a strange new animal and it appears to be a bird, so we believe the animal is a bird. As it comes closer to our hiding place, we see clearly that the animal is red, so we believe that it is a red bird. To remove further doubts about the animal birdhood, we call in a bird expert who takes it for examination, and concludes that it is not really a bird, but some sort of mammal. The question, now, is whether we should still believe that the animal is red."

Jin and Thielscher answer in the affirmative. Since the color of the animal is not related to its birdhood, giving up our belief that the animal is a bird, should not affect our belief that the animal is red. The DP postulates, however, do not rule out this possibility.

In particular, let us assume that our belief set, attached to the initial belief state S, is $B(S) = Cn(b)$, where b is a propositional variable standing for "the animal is a bird". Our language also contains the variable r that stands for "the animal is red". It turns out that there exists a modified AGM revision function $\circ$ that satisfies (DP1)–(DP4), such that $r \notin B\big((S \circ r) \circ \neg b\big)$; this revision function is known as *natural revision* (Boutilier, 1993).

Accordingly, Jin and Thielscher proposed a new postulate, which they call *Independence* postulate (Ind), that blocks unreasonable behaviours that are admitted by the DP postulates:[9]

(Ind) If $\neg\psi \notin B(S \circ \varphi)$, then $\psi \in B\big((S \circ \psi) \circ \varphi\big)$.

The characterization of the Independence postulate in terms of total preorders over worlds is as follows:

---

9. Postulate (Ind) was studied independently by Booth, Chopra, and Meyer (2005), Booth and Meyer (2006) as well.

(IndR)   If $r \in [\varphi]$ and $r' \in [\neg\varphi]$, then $r \preceq_S r'$ implies $r \prec_{S \circ \varphi} r'$.

It can be shown that postulate (Ind) is stronger than (DP3) and (DP4), as it implies both of them. Jin and Thielscher (2005, 2007) showed that postulate (Ind) is consistent with the modified AGM postulates for revision and the DP postulates, combined. They, also, suggest that any rational modified AGM revision function should satisfy postulates (DP1), (DP2) and (Ind).

3.4 Spohn's Conditionalization

Apart from the (qualitative) DP approach, the iterated-revision problem can be, also, addressed by the so-called Spohn's *conditionalization* (Spohn, 1988), which is the issue of the present subsection. Spohn's conditionalization is a process applied to an *ordinal conditional function* (OCF), a *quantitative* structure that Spohn uses to represent a belief state, related to a belief set.

Definition 3 (Ordinal Conditional Function, Spohn, 1988). *An ordinal conditional function (OCF) κ is a function from the set $\mathbb{M}$ of possible worlds to the class of ordinals, such that at least one world is assigned the ordinal 0.*

For mathematical simplicity, we take the range of an OCF κ to be the natural numbers (Spohn, 1991). For any world r, we call $\kappa(r)$ the *rank* of r. Intuitively, κ assigns a plausibility grading to worlds; the larger $\kappa(r)$ is for some world r, the less plausible r is.[10]

This plausibility grading can, easily, be extended to sentences, by requiring that the rank of a sentence be the lowest rank assigned to a world that satisfies the sentence.

Definition 4 (Sentence Plausibility Grading, Spohn, 1988). *For any contingent sentence $\varphi \in \mathcal{L}$, $\kappa(\varphi)$ denotes the lowest rank of a φ-world.[11] In symbols:*

$$\kappa(\varphi) = min\Big(\big\{\kappa(r) : r \in [\varphi]\big\}\Big).$$

The most plausible worlds (i.e., those whose κ-value is zero) define the belief set that κ is related to. In particular:

Definition 5 (Belief Set of an OCF, Spohn, 1988). *The belief set related to an OCF κ is the following theory:*

$$th\Big(\big\{r \in \mathbb{M} : \kappa(r) = 0\big\}\Big).$$

Given an OCF κ, the new information φ, and the *degree of firmness* $d \geqslant 0$ with which φ is incorporated into the *new* belief state, the revised (new) OCF, denoted $\kappa \bullet \langle \varphi, d \rangle$, is produced via Spohn's *conditionalization*, defined as follows:

---

10. An OCF is a richer structure than a total preorder over worlds, since the former tells us by *how much* one world is more plausible than another world, contrary to the latter which is a purely *qualitative* construction.

11. Notice that a sentence (rather than a world) is used as an argument in an OCF.

$$\kappa \bullet \langle \varphi, d \rangle (r) = \begin{cases} \kappa(r) - \kappa(\varphi) & \text{if } r \in [\varphi] \\ \\ \kappa(r) - \kappa(\neg\varphi) + d & \text{if } r \in [\neg\varphi] \end{cases}$$

In what follows, will shall confine ourselves to the *principal case of belief revision*; hence, we consider that $\kappa(\varphi) > 0$, $\kappa(\neg\varphi) = 0$ and $d > 0$. In this case, the agent initially believes $\neg\varphi$ with a degree of firmness $\kappa(\varphi)$, and, after revision, she ends up believing φ with a degree of firmness $\kappa \bullet \langle \varphi, d \rangle (\neg\varphi) = d$.

3.5 Parikh's Notion of Relevance

The classical AGM paradigm is liberal regarding not only iteration, but also *relevance*, as identified by Parikh (1999). Against this background, Parikh proposed a new axiom, named (P), to supplement the AGM postulates for revision.

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

In the above condition, $Cn_{\mathcal{L}_x}(x)$ denotes the logical closure of x *in* the sublanguage $\mathcal{L}_x$; i.e., $Cn_{\mathcal{L}_x}(x) = Cn(x) \cap \mathcal{L}_x$. Axiom (P) is essentially based on a *syntax-splitting* approach, and, as pointed out by Peppas, Williams, Chopra, and Foo (2015), is open to two different interpretations, called the *weak* and the *strong* version of (P). The weak version of (P), denoted by (P1) and presented below, is sufficient for the promised results.[12]

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

Condition (P1) is a quite intuitive constraint and says that the revision of the belief set K by the epistemic input φ will *not affect* any belief content outside the language $\mathcal{L}_x$.

Peppas et al. (2015) characterized condition (P1) in terms of faithful preorders; in order to present this characterization, we need some additional terminology.[13]

Definition 6 (Difference between Possible Worlds)**.** *The difference between two possible worlds w, r of $\mathbb{M}$, denoted by Diff (w, r), is the set of propositional variables over which the two worlds disagree. In symbols:*

$$\text{Diff } (w, r) = \Big\{ p \in \mathcal{P} : w \models p \text{ and } r \models \neg p \Big\} \cup \Big\{ p \in \mathcal{P} : w \models \neg p \text{ and } r \models p \Big\}.$$

Definition 6 can be extended in order to include the difference between a theory of $\mathcal{L}$ and a world of $\mathbb{M}$. For this, however, the subsequent definition is necessary.

Definition 7 (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*

---

12. For details on the strong version of (P), the reader is referred to (Peppas et al., 2015).
13. For a characterization of axiom (P) in the realm of all well-known constructive models of Belief Revision, refer to (Aravanis, Peppas, & Williams, 2019c).

$1 \leqslant i \neq j \leqslant n$. The set Q is a K-splitting iff there exist sentences $\varphi_1 \in \mathcal{L}^{Q_1}, \ldots, \varphi_n \in \mathcal{L}^{Q_n}$, such that $K = Cn(\varphi_1, \ldots, \varphi_n)$.

For every theory K of $\mathcal{L}$, there is a unique *finest K-splitting*, denoted by $\mathcal{F}$; i.e., one which refines every other K-splitting (Parikh, 1999).[14] That is to say, there is a unique way to think of theory K as being composed of disjoint compartments, referring to irrelevant subject matters.

Definition 8 (Difference between Theories and Possible Worlds, Peppas et al., 2015). *Let K be a theory of $\mathcal{L}$, and let $\mathcal{F} = \{F_1, \ldots, F_n\}$ be the finest K-splitting. Moreover, let r be a possible world of $\mathbb{M}$. The difference between K and r, denoted by $\text{Diff}(K, r)$, is the union of the elements F_i of $\mathcal{F}$, 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:*

$$\text{Diff}(K, r) = \bigcup \left\{ F_i \in \mathcal{F} : \text{for some } \varphi \in \mathcal{L}^{F_i}, \ K \models \varphi \text{ and } r \models \neg\varphi \right\}.$$

In the special case of a *complete* theory K (i.e., whenever $[K]$ is a singleton), Definition 8 collapses to Definition 6.

Example 3. Let $\mathcal{P} = \{a, b, c, d\}$, and $K = Cn(a \leftrightarrow b, c \leftrightarrow d)$. The finest K-splitting is $\mathcal{F} = \{\{a, b\}, \{c, d\}\}$. Now, let r be the world $r = \{\neg a, b, c, d\}$. Then, $\text{Diff}(K, r) = \{a, b\}$, since K and r disagree on at least one sentence φ that can be expressed in $\mathcal{L}^{\{a,b\}}$; for instance, for $\varphi = a \vee \neg b$, it holds that $K \models \varphi$ and $r \models \neg\varphi$.

Having defined the difference between theories and possible worlds, the faithful-preorders characterization of (P1) (in terms of belief sets) turns out to be the following two conditions on the relative plausibility of possible worlds.

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

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

Finally, we introduce the following useful definition:

Definition 9 (Splittable/Confined Theory). *Let K be a theory of $\mathcal{L}$. Iff, for some contingent sentences $x, y \in \mathcal{L}$, $K = Cn(x, y)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, we shall say that K is splittable. Iff, for some contingent sentence $x \in \mathcal{L}$, $K = Cn(x)$ and $\mathcal{L}_x \subset \mathcal{L}$, we shall say that K is confined (to $\mathcal{L}_x$).*

Note that, if theory K is confined to a sublanguage $\mathcal{L}'$ of $\mathcal{L}$, it also "splits" between $\mathcal{L}'$ and $\overline{\mathcal{L}'}$; the $\overline{\mathcal{L}'}$ part is trivial. Moreover, there is no contingent sentence $\xi \in \overline{\mathcal{L}'}$, such that $K \models \xi$. Notice, lastly, that a theory can be both splittable and confined.

Example 4. Suppose that $\mathcal{P} = \{a, b, c, d, e, f\}$, and let $K = Cn(a \leftrightarrow (b \vee c), d \leftrightarrow e)$. Then, the finest K-splitting is $\mathcal{F} = \{\{a, b, c\}, \{d, e\}, \{f\}\}$. Note that theory K is splittable, and, moreover, has no information about propositional variable f, since it is confined to the sublanguage $\mathcal{L}^{\{a,b,c,d,e\}}$.

---

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

4. Relevance-Sensitive Belief Revision and the DP Postulates

Despite relevance and iteration seem to be unrelated at first glance, dealing with different aspects of the revision process, it turns out that *each one* of the DP postulates is, in a *strong sense*, incompatible with Parikh's relevance-sensitive axiom (P).

In particular, as we show in this section, given *any splittable* theory K of $\mathcal{L}$, there exists a belief state S associated with K, and a sentence φ of $\mathcal{L}$, such that no modified AGM revision function $\circ$ satisfies (P) (at S and $S \circ \varphi$), *in addition* to any of (DP1)–(DP4). An *even stronger* incompatibility result concerning (DP2) is also established, covering *any belief state of any splittable or (non-trivially) confined* theory of the language.

For proving the above-mentioned incompatibilities, the *weak* version of axiom (P) is sufficient; i.e., condition (P1). Before proceeding to the formal results, however, a minor reformulation of this latter condition, and its corresponding faithful-preorders characterization, is required. Recall that Peppas *et al.* introduced (P1) and (Q1)–(Q2) in terms of belief sets. These conditions can quite straightforwardly be recast in terms of *belief states*. The only modifications required are to replace K with $B(S)$, $K * \varphi$ with $B(S \circ \varphi)$, $\prec_K$ with $\prec_S$, and $\approx_K$ with $\approx_S$. Then, we have the following conditions (for a belief state S):

(SP1) If $B(S) = Cn(x, y)$, $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$, and $\varphi \in \mathcal{L}_x$, then $B(S \circ \varphi) \cap \overline{\mathcal{L}_x} = B(S) \cap \overline{\mathcal{L}_x}$.

(SQ1) If $\mathit{Diff}\left(B(S), r\right) \subset \mathit{Diff}\left(B(S), r'\right)$ and $\mathit{Diff}\left(r, r'\right) \cap \mathit{Diff}\left(B(S), r\right) = \varnothing$, then $r \prec_S r'$.

(SQ2) If $\mathit{Diff}\left(B(S), r\right) = \mathit{Diff}\left(B(S), r'\right)$ and $\mathit{Diff}\left(r, r'\right) \cap \mathit{Diff}\left(B(S), r\right) = \varnothing$, then $r \approx_S r'$.

Theorem 3, then, establishes the first promised result.

Theorem 3. *Let K be any splittable theory of $\mathcal{L}$. There exists a belief state S, associated with K, and a sentence φ of $\mathcal{L}$, such that no modified AGM revision function $\circ$ satisfies (P) at S and $S \circ \varphi$, in addition to any of the postulates (DP1)–(DP4).*

Proof. Let K be any theory of $\mathcal{L}$, such that, for some contingent sentences $x, y \in \mathcal{L}$, $K = Cn(x, y)$ and $\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$. We will show that there exists a belief state S associated with K, such that, for some sentence φ of $\mathcal{L}$, the transition from $\preceq_S$ to $\preceq_{S \circ \varphi}$ cannot obey the dictates of (SQ1)–(SQ2), *in addition* to the dictates of each one of (R1)–(R4).

To this end, consider any belief state S, such that $B(S) = K$, and assume that (SQ1)–(SQ2) and (R1)–(R4) are satisfied. Moreover, suppose that F_x, F_y are two elements of the finest $B(S)$-splitting, such that F_x, F_y contain propositional variables of $\mathcal{L}_x$, $\mathcal{L}_y$, respectively.

- **Violation of (R1)**

    Let w be any $B(S)$-world. Consider a world r'', such that it disagrees with *all* $B(S)$-worlds on some propositional variables in F_x *and* in F_y (i.e., $r'' \cap \mathcal{L}^{F_x} \neq w' \cap \mathcal{L}^{F_x}$ and $r'' \cap \mathcal{L}^{F_y} \neq w' \cap \mathcal{L}^{F_y}$, for all $w' \in [B(S)]$), and it agrees with w on the remaining variables (if any, namely, in case $\mathcal{P} - (F_x \cup F_y) \neq \varnothing$). Then, $\mathit{Diff}\left(B(S), r''\right) =$

98

$F_x \cup F_y$, and thus $r'' \notin [B(S)]$.[15] Moreover, consider the two worlds r, r', such that $r = \left(r'' \cap \mathcal{L}^{F_x}\right) \cup \left(w \cap \mathcal{L}^{\mathcal{P}-F_x}\right)$ and $r' = \left(r'' \cap \mathcal{L}^{F_y}\right) \cup \left(w \cap \mathcal{L}^{\mathcal{P}-F_y}\right)$ (see Figure 3). Then, $\mathit{Diff}\left(B(S), r\right) = F_x$, $\mathit{Diff}\left(B(S), r'\right) = F_y$ (see Footnote 15), and thus $r, r' \notin [B(S)]$.

Therefore, we derive that $\mathit{Diff}\left(B(S), r\right) \subset \mathit{Diff}\left(B(S), r''\right)$ and $\mathit{Diff}\left(r, r''\right) \cap \mathit{Diff}\left(B(S), r\right) = \varnothing$, and $\mathit{Diff}\left(B(S), r'\right) \subset \mathit{Diff}\left(B(S), r''\right)$ and $\mathit{Diff}\left(r', r''\right) \cap \mathit{Diff}\left(B(S), r'\right) = \varnothing$. Then from (SQ1), it follows that $r \prec_S r''$ and $r' \prec_S r''$. Since $\mathit{Diff}\left(B(S), r\right) \cap \mathit{Diff}\left(B(S), r'\right) = \varnothing$, the antecedent of both (SQ1) and (SQ2) is not satisfied for worlds r and r'; hence, (SQ1)–(SQ2) place no constraints on the relative order of r and r', with respect to $\preceq_S$ (these two worlds are Diff-incomparable).

Now, consider the *particular* belief state —and call this S from now on— where it holds that $r \prec_S r' \prec_S r''$. Define φ to be the sentence $\varphi = r \vee r' \vee r''$. According to the definition of $\preceq_S$, there is *only one* $\preceq_S$-minimal φ-world, namely r; therefore from (F∘), it follows that $\left[B(S \circ \varphi)\right] = \{r\}$. Then, again by the choice of r', r'', we derive that $\mathit{Diff}\left(B(S \circ \varphi), r''\right) = \mathit{Diff}\left(r, r''\right) \subset \mathit{Diff}\left(B(S \circ \varphi), r'\right) = \mathit{Diff}\left(r, r'\right)$. Hence, no matter what the new preorder $\preceq_{S \circ \varphi}$ is, as long as it satisfies (SQ1), it holds that $r'' \prec_{S \circ \varphi} r'$.

However, given that $r', r'' \in [\varphi]$ and $r' \preceq_S r''$, (R1) entails that $r' \preceq_{S \circ \varphi} r''$. Contradiction.

- **Violation of (R2)**

  Consider the three worlds $r, r', r'' \notin [B(S)]$ of the previous case. As stated, the fact that $r \prec_S r''$ and $r' \prec_S r''$ holds for *any* belief state S associated with K. Define φ to be the sentence $\varphi = r$. Therefore from (F∘), it follows that $\left[B(S \circ \varphi)\right] = \{r\}$. Then, we derive that $\mathit{Diff}\left(B(S \circ \varphi), r''\right) = \mathit{Diff}\left(r, r''\right) \subset \mathit{Diff}\left(B(S \circ \varphi), r'\right) = \mathit{Diff}\left(r, r'\right)$. Hence, no matter what the new preorder $\preceq_{S \circ \varphi}$ is, as long as it satisfies (SQ1), it holds that $r'' \prec_{S \circ \varphi} r'$.

  However, given that $r', r'' \in [\neg \varphi]$ and $r' \preceq_S r''$, (R2) entails that $r' \preceq_{S \circ \varphi} r''$. Contradiction.

- **Violation of (R3) and (R4)**

  Assume the three worlds $r, r', r'' \notin [B(S)]$ of the previous case, and consider again the *particular* belief state S where $r \prec_S r' \prec_S r''$. Define φ to be the sentence $\varphi = r \vee r'$. Therefore from (F∘), it follows that $\left[B(S \circ \varphi)\right] = \{r\}$. Then, we derive that $\mathit{Diff}\left(B(S \circ \varphi), r''\right) = \mathit{Diff}\left(r, r''\right) \subset \mathit{Diff}\left(B(S \circ \varphi), r'\right) = \mathit{Diff}\left(r, r'\right)$. Hence, no matter what the new preorder $\preceq_{S \circ \varphi}$ is, as long as it satisfies (SQ1), it holds that $r'' \prec_{S \circ \varphi} r'$.

---

15. The world r'' indeed exists, since there exists at least one world outside $[B(S)]$, such that it entails $\neg x$ *and* $\neg y$. To see that $\mathit{Diff}\left(B(S), r''\right)$ contains the propositional variables of F_x, observe that, by the choice of r'', there is a sentence $\xi \in \mathcal{L}^{F_x}$ (i.e., $\xi = r'' \cap \mathcal{L}^{F_x}$), such that $K \models \neg \xi$ and $r'' \models \xi$. The same applies for F_y. To see that $\mathit{Diff}\left(B(S), r''\right)$ contains *only* the propositional variables of F_x and F_y, suppose, towards contradiction, that this is not the case (of course, in case $\mathcal{P} - (F_x \cup F_y) \neq \varnothing$). Then, there is a sentence $\zeta \in \mathcal{L}^{\mathcal{P}-(F_x \cup F_y)}$, such that $K \models \zeta$ and $r'' \models \neg \zeta$. This again entails that $w \models \zeta$, contradicting the fact that w and r'' agree on the variables outside $F_x \cup F_y$.

$w \in [B(S)]$

$r \notin [B(S)]$

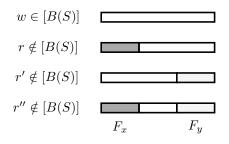$r' \notin [B(S)]$

$r'' \notin [B(S)]$

F_x F_y

Figure 3: Abstract visualization of the worlds appearing in the proof of Theorem 3. Same color indicates identical propositional variables between worlds.

S_1
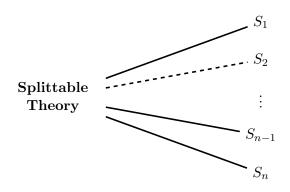
S_2

Splittable Theory

$\vdots$

S_{n-1}

S_n

Figure 4: All belief states associated with an arbitrary splittable theory. Dashed line indicates the belief state in which axiom (P) and postulates (DP1)–(DP4) clash.

However, given that $r' \in [\varphi]$, $r'' \in [\neg\varphi]$ and $r' \prec_S r''$, (R3) entails that $r' \prec_{S \circ \varphi} r''$, which is clearly a contradiction. In a similar vein, given that $r' \in [\varphi]$, $r'' \in [\neg\varphi]$ and $r' \preceq_S r''$, (R4) entails that $r' \preceq_{S \circ \varphi} r''$. A contradiction once again. ∎

The result of the aforementioned theorem is visualized in Figure 4.

A stronger incompatibility result concerning postulate (DP2) can be established as well. First, however, the following definition regarding confined theories is introduced.

Definition 10 (Non-Trivially Confined Theory). *We shall say that a theory K of $\mathcal{L}$ is non-trivially confined iff K is confined to a sublanguage $\mathcal{L}'$ of $\mathcal{L}$, and, moreover, $\overline{\mathcal{L}'}$ contains at least two propositional variables.*

Clearly, if a theory is non-trivially confined, then the set $\mathcal{P}$ contains *at least three* propositional variables.

Theorem 4. *Let K be any splittable or non-trivially confined theory of $\mathcal{L}$, and let S be any belief state, associated with K. There exists a sentence φ of $\mathcal{L}$, such that no modified AGM revision function $\circ$ satisfies (P) at S and $S \circ \varphi$, in addition to postulate (DP2).*

Proof. The proof is divided into two parts. The first part corresponds to the case of any *splittable* theory K, such that, for some contingent sentences $x, y \in \mathcal{L}$, $K = Cn(x, y)$ and

100

$\mathcal{L}_x \cap \mathcal{L}_y = \varnothing$. The second part corresponds to the case of any *non-trivially confined* theory K, such that, for some contingent sentence $x \in \mathcal{L}$, $K = Cn(x)$ and $\mathcal{L}_x \subset \mathcal{L}$.

Part I: Splittable Theory — $\mathbf{K = Cn(x,y)}$ and $\mathcal{L}_{\mathbf{x}} \cap \mathcal{L}_{\mathbf{y}} = \varnothing$

This part of the proof is identical to the proof of Theorem 3, concerning the violation of (R2), as the fact that $r \prec_S r''$ and $r' \prec_S r''$ holds for *any* belief state S associated with *any* splittable theory K.

Part II: Non-Trivially Confined Theory — $\mathbf{K = Cn(x)}$ and $\mathcal{L}_{\mathbf{x}} \subset \mathcal{L}$

We will show that, for *any* belief state S associated with K, there exists a sentence φ of $\mathcal{L}$, such that the transition from $\preceq_S$ to $\preceq_{S \circ \varphi}$ cannot obey the dictates of (SQ1)–(SQ2), *in addition* to the dictates of (R2).

To this end, consider *any* belief state S, such that $B(S) = K$, and assume that (SQ1)–(SQ2) and (R2) are satisfied. Since $B(S)$ is *non-trivially* confined to $\mathcal{L}_x$, it follows that the sublanguage $\overline{\mathcal{L}_x}$ contains *at least two* propositional variables, and, moreover, there exist *at least four* worlds outside $[B(S)]$, such that they agree on *all* propositional variables in $\mathcal{L}_x$.[16]

Let r be a world outside $[B(S)]$. Consider the world $r' \notin [B(S)]$, such that it agrees on all propositional variables with r, except from *one* variable of $\overline{\mathcal{L}_x}$. Moreover, consider the world $r'' \notin [B(S)]$, such that it agrees on all propositional variables with r', except from *one* variable of $\overline{\mathcal{L}_x}$, and $r'' \neq r$ (see Figure 5).[17] Then, by the choice of r, r', r'' (that they are outside $[B(S)]$ and agree on all propositional variables in $\mathcal{L}_x$) and since theory $B(S)$ knows nothing about $\overline{\mathcal{L}_x}$ (i.e., there is no contingent sentence $\xi \in \overline{\mathcal{L}_x}$, such that $B(S) \models \xi$), we have that $\varnothing \neq Diff\left(B(S), r\right) = Diff\left(B(S), r'\right) = Diff\left(B(S), r''\right) \subseteq \mathcal{L}_x$, $Diff\left(r, r'\right) \cap Diff\left(B(S), r\right) = \varnothing$ and $Diff\left(r', r''\right) \cap Diff\left(B(S), r'\right) = \varnothing$. Hence from (SQ2), we derive that $r \approx_S r' \approx_S r''$.

Now, define φ to be the sentence $\varphi = r$. Therefore from (F∘), it follows that $\left[B(S \circ \varphi)\right] = \{r\}$. Then, again by the choice of r, r', r'', we derive that $Diff\left(B(S \circ \varphi), r'\right) = Diff\left(r, r'\right) \subset Diff\left(B(S \circ \varphi), r''\right) = Diff\left(r, r''\right)$. Hence, no matter what the new preorder $\preceq_{S \circ \varphi}$ is, as long as it satisfies (SQ1), it holds that $r' \prec_{S \circ \varphi} r''$.

However, given that $r', r'' \in [\neg \varphi]$ and $r' \approx_S r''$, (R2) entails that $r' \approx_{S \circ \varphi} r''$. Contradiction. ∎

Likewise, the result of the aforementioned theorem is visualized in Figure 6.

Remark 2. Since a modified AGM revision function is constructed/defined by means of condition (F∘) for *every* belief state (see Theorem 2), Theorems 3 and 4 entail that there exists no modified AGM revision function *of any type*, satisfying Parikh's relevance-sensitive axiom (P) *in addition to any* of the DP postulates.

---

16. Let p, q be two propositional variables of $\overline{\mathcal{L}_x}$, and let r_1, r_2, r_3, r_4 be the aforementioned four worlds. Then, it is not hard to verify that $\{p, q\} \subseteq r_1$, $\{\neg p, q\} \subseteq r_2$, $\{p, \neg q\} \subseteq r_3$, and $\{\neg p, \neg q\} \subseteq r_4$.
17. To see that the worlds r' and r'' indeed exist, refer to the observation made in Footnote 16.
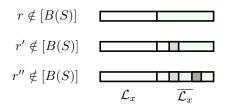
Figure 5: Abstract visualization of the worlds appearing in the proof of Theorem 4 (Part II). Same color indicates identical propositional variables between worlds.
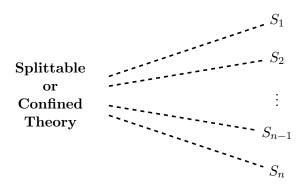

Figure 6: All belief states associated with an arbitrary splittable or (non-trivially) confined theory. Dashed lines indicate the belief states in which axiom (P) and postulate (DP2) clash.

5. Corollaries of the Main Incompatibility Results

The incompatibility results of the previous section imply *fully-fledged* inconsistencies between other existing belief-revision proposals.

5.1 Parametrized Difference Revision Operators

Parametrized Difference revision operators (abbrev. *PD operators*) constitute an important (proper) subclass of AGM revision functions, recently introduced by Peppas and Williams (2016, 2018) as a generalization of Dalal's Hamming-based revision operator (Dalal, 1988). They are well-suited for real-world implementations, due to their favourable properties, including low specification cost and high expressivity.

A central result of (Peppas & Williams, 2016) is that PD operators satisfy the weak version of axiom (P); i.e., condition (P1).[18] Thus, in view of Theorems 3 and 4, we obtain immediately the following two corollaries:[19]

---

18. Aravanis, Peppas, and Williams (2019a) showed that PD operators satisfy the strong version of (P) as well.

19. Although PD operators and Dalal's operator were introduced in terms of belief sets, they can *straightforwardly* be recast in terms of belief states, as it was done for conditions (SP1) and (SQ1)–(SQ2).

Corollary 1. *Parametrized Difference revision operators are inconsistent with the DP postulates.*

Corollary 2. *Dalal's revision operator is inconsistent with the DP postulates.*

PD operators, as well as Dalal's operator, are based on a simple and intuitive construction, and they are the *only* among a series of popular operators —including those by Borgida (1985), Winslett (1988), Satoh (1988) and Weber (1986)— that satisfy the full set of (modified) AGM postulates for revision. Obviously, having such well-behaved revision operators violating the DP postulates is a major drawback.

5.2 The Independence Postulate

Since postulate (Ind) implies both (DP3) and (DP4), the following corollary is an immediate consequence of Theorem 3.

Corollary 3. *The Independence postulate (Ind) is inconsistent with Parikh's axiom (P).*

Note that a weaker inconsistency between the Independence postulate and Parikh's axiom was identified by Peppas et al. (2008, Corollary 2).

5.3 Spohn's Conditionalization

Darwiche and Pearl (1997) showed that the DP postulates are satisfied by Spohn's proposal, restricted to the principal case of *revision* scenarios (where $\kappa(\varphi) > 0$, $\kappa(\neg\varphi) = 0$ and $d \geqslant 1$). Roughly speaking, this was achieved by constructing an OCF-based concrete revision operator, herein denoted by $\circledast$, and proving that $\circledast$ satisfies (DP1)–(DP4) (in addition to the modified AGM postulates for revision), given that each OCF κ can induce a total preorder $\preceq_\kappa$ over worlds, by letting

$$r \preceq_\kappa r' \quad \text{iff} \quad \kappa(r) \leqslant \kappa(r'), \quad \text{for all } r, r' \in \mathbb{M}.[20]$$

Therefore, in view of Theorems 3 and 4, we immediately obtain the following result, that indicates the inconsistency between Spohn's conditionalization (as expressed by means of the OCF-based revision operator $\circledast$) and Parikh's relevance-sensitive axiom (P).

Corollary 4. *Revision operator $\circledast$ does not satisfy axiom (P).*

6. Discussion on the Incompatibilities

Unfortunately, the obtained incompatibility results cannot be ignored, since both relevance and iteration play a central role in the revision process. Given the simplicity and intuitive appeal of Parikh's axiom, the conflicts raise further concerns about the validity of the DP postulates.

Part of the problem, we claim, is that the DP postulates try to do more than what they should. Consider, for instance, Example 1 of Subsection 3.2, which was presented as one of the motivations for (DP1)–(DP4). As argued by Darwiche and Pearl, it would be counterintuitive not to believe in r at the belief set $(Cn(s, r) * \neg s) * s$. While the AGM postulates

for revision alone cannot guarantee this, when augmented with axiom (P1), i.e., weak (P), they produce the desired result. In particular, since $\mathcal{L}_s \cap \mathcal{L}_r = \varnothing$, by (P1) it follows that $r \in Cn(s, r) * \neg s$. Invoking (P1) one more time leads us to $r \in \big(Cn(s, r) * \neg s\big) * s$.[21] The problem with Example 2 of Subsection 3.3 can, analogously, be resolved by appealing to relevance.

Hence, we observe that, although Examples 1 and 2 were introduced as typical scenarios for iterated belief revision, they are truly relevance-related scenarios. Concerning the former example, our belief that "X is rich" survives evidence refuting and then supporting that "X is smart", not because of some underlying principle of iterated revision, but because the two propositions are *irrelevant* to each other; thus, "X is rich" survives changes to our belief about X's economic situation, regardless of whether these changes occur in *one step* or in *multiple steps*. Concerning the latter example, the fact that "the animal is a bird" is not related to the fact that "the animal is red", and it is because of this unrelatedness that we retain the latter belief, despite changes to the former.

Analogous remarks can be made about several examples introduced as typical iterated-revision scenarios — e.g., in the seminal work of Darwiche and Pearl (1997) or in subsequent efforts such as those of Konieczny and Pérez (2000), Nayak, Pagnucco, and Peppas (2003), Jin and Thielscher (2007). In many cases, the involved scenarios are truly relevance-related, as they can be resolved by appealing to relevance.

Can axiom (P) solve all problem cases mentioned in the iterated-revision literature? Of course not. We do not claim that iterated revision reduces to relevance-sensitive revision; nor is axiom (P) a fully-fledged formalization of the role of relevance in belief change. It seems, yet, that the DP postulates incorporate aspects of the revision process that essentially relate to the notion of relevance, a fact that in turn leads to a clash with axiom (P). On that premise, a potential way to reconcile relevance and iteration would be to identify and address these two problems of belief revision *separately*.

7. Conclusion

Notwithstanding the extensive work on iterated belief revision, there is, still, no fully satisfactory solution. This article studies this fundamental aspect of the revision process, and points out important observations on the DP approach. In particular, at first, we generalized the conflicts between relevance-sensitive and iterated belief revision, proving that there is no modified AGM revision function of any type, satisfying axiom (P) in addition to the DP postulates. These fully-fledged incompatibilities imply significant corollaries; PD operators and Dalal's operator are incompatible with the DP approach, as well as the Independence postulate and Spohn's conditionalization are incompatible with axiom (P).

Arguably, iteration plays a central role in the revision process, and so does relevance. The established results show that the dominant models for these two aspects of the revision process are in deep conflict. Undoubtedly, however, any comprehensive formal framework of Belief Revision would need to consistently combine iterated and relevance-sensitive belief change. Finding ways to reconcile iteration and relevance —either by weakening the involved postulates or by following an entirely new approach— constitutes an imperative future task.

---

21. For coherence, we use the classical-AGM-paradigm notation of Example 1.

Acknowledgments

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. (2019a). An investigation of parametrized difference revision operators. *Annals of Mathematics and Artificial Intelligence*.

Aravanis, T., Peppas, P., & Williams, M.-A. (2019b). Observations on Darwiche and Pearl's approach for iterated belief revision. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI 2019)*, pp. 1509–1515, Macao, China.

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

Areces, C., & Becher, V. (2001). Iterable AGM functions. In Williams, M.-A., & Rott, H. (Eds.), *Frontiers in Belief Revision*, Vol. 22 of *Applied Logic Series*, pp. 165–196. Springer, Dordrecht.

Benferhat, S., Kaci, S., Le Berre, D., & Williams, M.-A. (2004). Weakening conflicting information for iterated revision and knowledge integration. *Artificial Intelligence*, *153*, 339–371.

Booth, R., Chopra, S., & Meyer, T. (2005). Restrained revision. In *6th Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC 2005)*, Edinburgh.

Booth, R., & Meyer, T. (2006). Admissible and restrained revision. *Journal Of Artificial Intelligence Research*, *26*, 127–151.

Booth, R., & Meyer, T. (2011). How to revise a total preorder. *Journal of Philosophical Logic*, *40*, 193–238.

Borgida, A. (1985). Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, *10*(4), 563–603.

Boutilier, C. (1993). Revision sequences and nested conditionals. In *Proceedings of the 13th International Joint Conference on Artifical Intelligence (IJCAI 1993)*, pp. 519–525. Morgan Kaufmann Publishers Inc.

Chopra, S., Ghose, A., Meyer, T., & Wong, K.-S. (2008). Iterated belief change and the recovery axiom. *Journal of Philosophical Logic*, *37*, 501–520.

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.

Darwiche, A., & Pearl, J. (1994). On the logic of iterated belief revision. In *Proceedings of the 5th Conference on Theoretical Aspects of Reasoning About Knowledge (TARK 1994)*, pp. 5–23, Pacific Grove, California. Morgan Kaufmann.

Darwiche, A., & Pearl, J. (1997). On the logic of iterated belief revision. *Artificial Intelligence*, *89*, 1–29.

Delgrande, J., Dubois, D., & Lang, J. (2006). Iterated revision as prioritized merging. In *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR 2006)*, pp. 210–220.

Delgrande, J. P., & Schaub, T. (2003). A consistency-based approach for belief change. *Artificial Intelligence*, *151*, 1–41.

Fermé, E., & Hansson, S. O. (2018). *Belief Change: Introduction and Overview*. Springer International Publishing.

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

Hansson, S. O. (2016). Iterated descriptor revision and the logic of Ramsey test conditionals. *Journal of Philosophical Logic*, *45*, 429–450.

Hunter, A., & Delgrande, J. P. (2005). Iterated belief change: A transition system approach. In *Proceedings of the 19th International Joint Conference on Artifical Intelligence (IJCAI 2005)*, pp. 460–465. Morgan Kaufmann Publishers Inc.

Hunter, A., & Delgrande, J. P. (2007). An action description language for iterated belief change. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence (IJCAI 2007)*, pp. 2498–2503. Morgan Kaufmann Publishers Inc.

Jin, Y., & Thielscher, M. (2005). Iterated belief revision, revised. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI 2005)*, pp. 478–483, Edinburgh.

Jin, Y., & Thielscher, M. (2007). Iterated belief revision, revised. *Artificial Intelligence*, *171*, 1–18.

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

Konieczny, S., & Pérez, R. P. (2000). A framework for iterated revision. *Journal of Applied Non-Classical Logics*, *10*, 339–367.

Lehmann, D. (1995). Belief revision, revised. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI 1995)*, pp. 1534–1540, Montreal, Quebec. Morgan Kaufmann Publishers Inc.

Nayak, A., Pagnucco, M., & Peppas, P. (2003). Dynamic belief revision operators. *Artificial Intelligence*, *146*(4), 193–228.

Nayak, A. C. (1994). Iterated belief change based on epistemic entrenchment. *Erkenntnis*, *41*, 353–390.

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. (2014). A panorama of iterated revision. In Hansson, S. O. (Ed.), *David Makinson on Classical Methods for Non-Classical Problems*, pp. 71–94. Springer Netherlands.

Peppas, P., Fotinopoulos, A. M., & Seremetaki, S. (2008). Conflicts between relevance-sensitive and iterated belief revision. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*, pp. 85–88, Patras, Greece. IOS Press.

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.

Satoh, K. (1988). Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pp. 455–462. Springer-Verlag, Tokyo.

Shapiro, S., Pagnucco, M., Lespérance, Y., & Levesque, H. J. (2000). Iterated belief change in the situation calculus. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR 2000)*, pp. 210–220.

Shapiro, S., Pagnucco, M., Lespérance, Y., & Levesque, H. J. (2011). Iterated belief change in the situation calculus. *Artificial Intelligence*, *175*, 165–192.

Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W. L., & Skyrms, B. (Eds.), *Causation in Decision, Belief Change, and Statistics*, Vol. 42 of *The University of Western Ontario Series in Philosophy of Science*, pp. 105–134. Springer Netherlands.

Spohn, W. (1991). A reason for explanation: Explanations provide stable reasons. In Spohn, W., Fraassen, B. C. V., & Skyrms, B. (Eds.), *Existence and Explanation*, Vol. 49 of *The University of Western Ontario Series in Philosophy of Science*, pp. 165–196. Springer, Dordrecht.

Stalnaker, R. (2009). Iterated belief revision. *Erkenntnis*, *70*(2), 189–209.

Weber, A. (1986). Updating propositional formulas. In *Proceedings of the 1st Conference on Expert Database Systems*, pp. 487–500. Columbia, SC: Institute of Information Management, Technology and Policy, University of South Carolina.

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.

Zhang, D. (2004). Properties of iterated multiple belief revision. In Lifschitz, V., & Niemelä, I. (Eds.), *LPNMR 2004: Logic Programming and Nonmonotonic Reasoning*, Vol. 2923 of *Lecture Notes in Computer Science*, pp. 314–325. Springer, Berlin, Heidelberg.