Kinetic Consistency and Relevance in Belief Revision

Pavlos Peppas[1,2(✉)] and Mary-Anne Williams[2]

[1] Department of Business Administration, University of Patras, Patras, Greece
pavlos@upatras.gr
[2] QCIS, Faculty of Engineering and IT, University of Technology Sydney, Ultimo, Australia

Abstract. A critical aspect of rational belief revision that has been neglected by the classical AGM framework is what we call the *principle of kinetic consistency*. Loosely speaking, this principle dictates that the revision policies employed by a rational agent at different belief sets, are not independent, but ought to be related in a certain way. We formalise kinetic consistency axiomatically and semantically, and we establish a representation result explicitly connecting the two. We then combine the postulates for kinetic consistency, with Parikh's postulate for relevant change, and add them to the classical AGM postulates for revision; we call this augmented set the *extended AGM postulates*. We prove the consistency and demonstrate the scope of the extended AGM postulates by showing that a whole new class of concrete revision operators introduced hererin, called *PD operators*, satisfies all extended AGM postulates. PD operators are of interest in their own right as they are natural generalisations of Dalal's revision operator. We conclude the paper with some examples illustrating the strength of the extended AGM postulates, even for iterated revision scenarios.

1 Introduction

The classical AGM postulates for belief revision, named (K*1) - (K*8), [1], although immensely successful, [10], have left certain crucial aspects of the revision process unattended. One of them is the notion of *relevant change* discussed by Parikh in [9]. Parikh argues that during belief revision a rational agent does not change her entire belief corpus, but only the portion of it that is relevant to the new information. This intuition is formally captured by means of a new postulate called herein (wP).

Another aspect of rational belief revision, that surprisingly has received little attention in the literature, if any, is what we call the *principle of kinetic consistency*. Loosely speaking, this principle dictates that the revision policies of a rational agent over *different* theories, are not independent, but ought to be related in a certain way.

Our aim in this paper is, firstly, to formally capture the principle of kinetic consistency, and secondly, to investigate the implications of combining the AGM postulates for revision, together with the postulates for kinetic consistency and Parikh's postulate for relevant change; we call this combination the *extended AGM postulates*.

More precisely, in the first part of the paper we discuss the intuition behind the principle of kinetic consistency and we formulate two new postulates, named (KC1) - (KC2), to encode it. We also characterise kinetic consistency semantically and a representation result is provided connecting the new postulates with the semantic characterization.

We then proceed to investigate the extended AGM postulates; i.e., the postulates (K*1) - (K*8), (wP), and (KC1) - (KC2). A by-product of our investigation is the introduction of a whole new class of concrete revision operators, called *parametrised difference revision operators*, or *PD operators* for short, that are of interest in their own right. PD operators are natural generalisations of Dalal's revision operator, with a much greater range of applicability.

We prove that every PD operator satisfies all extended AGM postulates. This result not only establishes the consistency of the extended AGM postulates but also sheds light into the scope and nature of the revision functions satisfying these postulates. The latter is an important contribution since hitherto it was clear how many, and what kind, of classical AGM revision functions survive the addition of postulate (wP) (let alone the addition of (wP) *and* (KC1)-(KC2)).[1]

We conclude the paper by illustrating the strength of the extended AGM postulates in a number of examples that are out of reach for vanilla AGM. These include (some) *iterated revision* scenarios, [12].

The paper is structured as follows. In the next section we introduce the necessary notation and terminology. Section 3 gives a brief overview of the AGM framework. In Sect. 4 we introduce and formalise the kinetic consistency principle; we also provide a representation result connecting the new postulates and semantics. Section 5 provides a brief review of relevant change. In Sect. 6 we introduce the class of PD operators and we prove that they all satisfy the extended AGM postulates. Section 7 lists a number of examples that demonstrate the strength of the extended AGM postulates. The last section contains some concluding remarks.

2 Preliminaries

Throughout this article we shall be working with a propositional language L built over *finitely many* propositional variables. The finite, nonempty set of all propositional variables is denoted by P. A literal is a variable in P or the negation of a variable. For a variable $q \in P$ we shall often write $\bar{q}$ instead of $\neg q$. The set of all interpretations over P is denoted $\mathbb{M}$. Interpretations will also be called possible worlds. We will often identify a possible world with the set (or sequence) of literals it satisfies. Moreover, we will sometimes abuse notation and use a possible world as a sentence, namely, the conjunction of all the literals it satisfies; for example, for possible worlds w, r, we may write $\neg w$ or $w \vee r$.

For a set of sentences Γ of L, we denote by $Cn(\Gamma)$ the set of all logical consequences of Γ, i.e., $Cn(\Gamma) = \{\varphi \in L : \Gamma \models \varphi\}$. A theory K of L is any set of sentences of L closed under $\models$, i.e., $K = Cn(K)$. We shall denote the set of all theories of L by $\mathbb{T}$. The set of all *consistent* theories of L is denoted by $\mathbb{K}$. A theory K of L is complete iff for all sentences $\varphi \in L$, $\varphi \in K$ or $\neg\varphi \in K$.

For a set of sentences Γ of L, $[\Gamma]$ denotes the set of all possible worlds that satisfy Γ. Often we shall use the notation $[\varphi]$ for a sentence $\varphi \in L$, as an abbreviation of $[\{\varphi\}]$. For a theory K and a set of sentences Γ of L, we shall denote by $K + \Gamma$ the closure under $\models$

---

[1] In fact in [9], Parikh conjectured that *no* classical AGM revision function survives (wP). This was later refuted in [11].

of $K \cup \Gamma$, i.e., $K + \Gamma = Cn(K \cup \Gamma)$. For a sentence $\varphi \in L$, we shall often write $K + \varphi$ as an abbreviation of $K + \{\varphi\}$. For any two sentences φ, ψ, we shall write $\varphi \equiv \psi$ as an abbreviation of $Cn(\varphi) = Cn(\psi)$.

3   The AGM Framework

In the AGM framework, [1], an agent's belief set is modelled as a theory of L. Epistemic input is represented as a logical sentence of L, and the process of belief revision is modelled as a function $*$ mapping a theory K and a sentence φ to a new theory $K * \varphi$.

An *AGM revision function*, [1], or revision function for short, is any function $* :$ $\mathbb{T} \times L \mapsto \mathbb{T}$ that satisfies certain constraints known as the *AGM postulates for revision*. There are eight such postulates, numbered (K*1) - (K*8). They are widely considered to have captured the essence of the revision process:

$(K * 1)$ $K * \varphi$ is a theory of 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 $(K * \varphi) + \psi \subseteq K * (\varphi \wedge \psi)$.

In addition to the postulates, a constructive model for revision functions based on possible worlds was proposed in [4]. Katsuno and Mendelzon, [6], subsequently simplified this model by constraining it to propositional logic over finite variables. In particular, to every belief set K, Katsuno and Mendelzon assign a *total preorder* $\preccurlyeq_K$ over the set of possible worlds $\mathbb{M}$.[2] We recall that a preorder $\preccurlyeq_K$ over $\mathbb{M}$, is any binary relation in $\mathbb{M}$ that is reflexive and transitive. The preorder is total iff for all $w, w' \in \mathbb{M}$, $w \preccurlyeq_K w'$ or $w' \preccurlyeq w$ (we shall be using infix notation throughout this paper). As usual, $\prec_K$ denotes the strict part of $\preccurlyeq_K$; i.e., $w \prec_K w'$ iff $w \preccurlyeq_K w'$ and $w' \not\preccurlyeq_K w$. Moreover, we shall write $w \approx_K w'$ iff $w \preccurlyeq_K w'$ and $w' \preccurlyeq_K w$.

The preorder $\preccurlyeq_K$ is assumed to satisfy the following constraints:

(i)   If $w_1, w_2 \in [K]$, then $w_1 \approx_K w_2$.
(ii)   If $w_1 \in [K]$ and $w_2 \notin [K]$, then $w_1 \prec_K w_2$.

For any theory K and total preorder $\preccurlyeq_K$ over $\mathbb{M}$, we shall say that $\preccurlyeq_K$ is *faithful to K* iff it satisfies the constraints (i) - (ii) above.

When $\preccurlyeq$ appears without a subscript, it represents a *function*, mapping a theory K to a total preorder $\preccurlyeq_K$ over $\mathbb{M}$. If for all $K \in \mathbb{K}$, the preorder $\preccurlyeq_K$ is faithful to K, then the function $\preccurlyeq$ is called a *faithful assignment*.

---

[2] To be precise, Katsuno and Mendelzon represent an agent's beliefs by a *sentence* rather than a theory. Hence they assign preorders to sentences rather than to theories. We use theories in order to adhere more closely to the original AGM approach. Since we deal only with languages built over finitely many variables, the difference is immaterial.

Intuitively, $\leqslant_K$ represents comparative plausibility: $w \leqslant_K w'$ iff given the agent's current beliefs K, w is at least as plausible as w'. Based on this reading, Katsuno and Mendelzon define $K * \varphi$ as follows:

$(\leqslant*)$ $[K * \varphi] = \min([\varphi], \leqslant_K)$.

In the above definition, $min(S, \leqslant_K)$ is the set of *minimal* elements of the set S with respect to $\leqslant_K$; i.e., $min(S, \leqslant_K) = \{w \in S \; : \text{for all } w' \in S, \text{ if } w' \leqslant_K w, \text{ then } w \leqslant_K w'\}$. Katsuno and Mendelzon proved the following representation theorem:

Theorem 1 [6]. *A revision operator * satisfies postulates (K*1) - (K*8) iff there exists a faithful assignment $\leqslant$ such that $(\leqslant*)$ holds for every $K \in \mathbb{T}$ and $\varphi \in L$.*

For ease of presentation, in the rest of the paper we shall focus only on revision of *consistent* theories by *consistent* input. Hence from now, unless explicitly stated otherwise, we assume that the initial belief set K is a consistent theory, and that the epistemic input φ is a consistent sentence.

4   Kinetic Consistency

Consider a rational agent whose current belief set is K and who uses the revision function $*$ to respond to new information. We stress that $*$ is defined as a *binary function* in the AGM framework; $* : \mathbb{T} \times L \mapsto \mathbb{T}$. Hence, in view of Theorem 1, the agent is equipped with faithful preorders, not just for K, but *for every other theory as well*. Through these preorders, the agent is able to answer hypothetical questions like "would ψ be true after revision by ϕ, had the initial belief set been H instead of K?"

Should the faithful preorders assigned to different theories be related? Or does every collection $\{\leqslant_H\}_{H \in \mathbb{T}}$ of faithful preorders correspond to a rational revision policy? The AGM postulates support the latter view, since they place no constraints on the preorders assigned to different theories. This is too liberal for a wide range of applications where there needs to be some kind of *consistency* in the epistemic choices that a rational agent makes across different belief sets.

Consider for example two different consistent complete theories K and H. Then for some distinct worlds w_1, w_2, $[K] = \{w_1\}$ and $[H] = \{w_2\}$. Let r, r' be another two distinct worlds, different from w_1, w_2. Moreover assume that from the perspective of w_1 as well as from the perspective of w_2, r is at least as plausible as r'; in symbols, $r \leqslant_K r'$ and $r \leqslant_H r'$. Consider now the preorder $\leqslant_{K \cap H}$ assigned to the theory $K \cap H$. As far as the agent knows at $K \cap H$, the real world can be either w_1 or w_2. Since in both cases r is at least as plausible as r', we argue that it is unreasonable to reverse their plausibility at $K \cap H$. This is the main intuition for what we call the *principle of kinetic consistency*.

Generalising this simple intuition to arbitrary consistent theories K, H leads us to the following constraints:

(KS1)   If $r \leqslant_K r'$ and $r \leqslant_H r'$, then $r \leqslant_{K \cap H} r'$.
(KS2)   If $r <_K r'$ and $r <_H r'$, then $r <_{K \cap H} r'$.

The two constraints deal with different cases of the same intuitive idea: since the preorders assigned to K and H encode (part of) the revision strategy of a *single agent*, then if $\leqslant_K$ and $\leqslant_H$ agree on the relative plausibility of two worlds r, r', then $\leqslant_{K \cap H}$ should also agree. This is because $[K \cap H] = [K] \cup [H]$, and since the view that, say, r is more plausible than r', has prevailed among the K-worlds and moreover it has also prevailed among the H-worlds, it would be unreasonable to have a reversal of this perception when the K-worlds and the H-worlds are grouped together in $[K \cap H]$.

Depending on the class of scenarios under investigation, additional constraints between faithful preorders may be required. However, (KS1) - (KS2) are the *core* domain-independent constraints for kinetic consistency.

The postulates corresponding to the above constraints are listed below:

(KC1)   If $T * (\varphi \vee \psi) \cup H * (\varphi \vee \psi) \not\vdash \neg\varphi$, then $\neg\varphi \notin (T \cap H) * (\varphi \vee \psi)$.

(KC2)   If $\neg\psi \in T * (\varphi \vee \psi)$, $\neg\psi \in H * (\varphi \vee \psi)$, and $T * (\varphi \vee \psi) \cup H * (\varphi \vee \psi)$ is consistent, then $\neg\psi \in (T \cap H) * (\varphi \vee \psi)$.

Theorem 2. *Let $*$ be an AGM revision function and $\leqslant$ a faithful assignment corresponding to $*$ via ($\leqslant *$). Then $*$ satisfies (KC1) - (KC2) iff $\leqslant$ satisfies (KS1) - (KS2) respectively.*

Proof

($\Rightarrow$)

Assume that $*$ satisfies (KC1). We show that $\leqslant$ satisfies (KS1). Let T, H be any two theories of L and $w, w' \in \mathbb{M}$ any two worlds such that $w \leqslant_T w'$ and $w \leqslant_H w'$. Clearly then $w \in min([w \vee w'], \leqslant_T)$, and $w \in min([w \vee w'], \leqslant_H)$. Hence by ($\leqslant *$), $w \in [T * (w \vee w')] \cap [H * (w \vee w')]$, and therefore $T * (w \vee w') \cup H * (w \vee w') \not\vdash \neg w$. Consequently, from (KC1), $\neg w \notin (T \cap H) * (w \vee w')$. This again entails that $w \leqslant_{T \cap H} w'$. Hence (KS1) is satisfied.

Assume that $*$ satisfies (KC2). We show that $\leqslant$ satisfies (KS2). Let T, H be any two theories of L and $w, w' \in \mathbb{M}$ any two worlds such that $w <_T w'$ and $w <_H w'$. Then clearly, $w \neq w'$ and moreover, $min([w \vee w'], \leqslant_T) = min([w \vee w'], \leqslant_H) = \{w\}$. Hence, $T * (w \vee w')$ is consistent with $H * (w \vee w')$. Moreover, $\neg w' \in T * (w \vee w')$ and $\neg w' \in H * (w \vee w')$. Therefore by (KC2), $\neg w' \in (T \cap H) * (w \vee w')$. This again entails that $w <_{T \cap H} w'$. Hence (KS2) is satisfied.

($\Leftarrow$)

Assume that $\leqslant$ satisfies (KS1). We show that $*$ satisfies (KC1). Let T, H be any two theories of L and $\varphi, \psi \in L$ any two sentences such that $T*(\varphi\vee\psi)\cup H*(\varphi\vee\psi) \not\vdash \neg\varphi$. Then there is a world $w \in [T * (\varphi \vee \psi)] \cap [H * (\varphi \vee \psi)]$ such that $w \models \varphi$. From $w \in [T * (\varphi \vee \psi)]$ we derive that $w \leqslant_T w'$ for all $w' \in [\varphi \vee \psi]$. Similarly, from $w \in [H * (\varphi \vee \psi)]$ it follows that $w \leqslant_H w'$ for all $w' \in [\varphi \vee \psi]$. Hence from (KS1), $w \leqslant_{T \cap H} w'$ for all $w' \in [\varphi \vee \psi]$. This again entails that $w \in [(T \cap H) * (\varphi \vee \psi)]$, and since $w \models \varphi$, we derive that $\neg\varphi \notin (T \cap H) * (\varphi \vee \psi)$.

Assume that $\leqslant$ satisfies (KS2). We show that $*$ satisfies (KC2). Let T, H be any two theories of L and $\varphi, \psi \in L$ any two sentences such that $\neg\psi \in T*(\varphi\vee\psi)$, $\neg\psi \in H*(\varphi\vee\psi)$, and $T * (\varphi \vee \psi)$ is consistent with $H * (\varphi \vee \psi)$. If ψ is inconsistent, then by (K*1) we

immediately derive that $\neg\psi \in (T \cap H) * (\varphi \vee \psi)$ as desired. Assume therefore that ψ is consistent. Since $T * (\varphi \vee \psi)$ is consistent with $H * (\varphi \vee \psi)$, there exists a world $w \in [T * (\varphi \vee \psi)] \cap [H * (\varphi \vee \psi)]$. Moreover, from $\neg\psi \in T * (\varphi \vee \psi)$ and (K*2) we derive that $\varphi \in T * (\varphi \vee \psi)$. Consequently, $w \models \varphi$ and $w \models \neg\psi$. From $w \in [T * (\varphi \vee \psi)]$ and $\neg\psi \in T * (\varphi \vee \psi)$ it follows that $w \prec_T w'$ for all $w' \in [\psi]$. Similarly, from $w \in [H * (\varphi \vee \psi)]$ and $\neg\psi \in H * (\varphi \vee \psi)$ it follows that $w \prec_H w'$ for all $w' \in [\psi]$. Hence by (KS2), $w \prec_{T \cap H} w'$, for all $w' \in [\psi]$. Therefore, since $w \models \varphi$, $\neg\psi \in (T \cap H) * (\varphi \vee \psi)$. Hence (KC2) is satisfied. □

5   Relevant Change

As already mentioned, relevant change is also not properly addressed in the classical AGM framework. In this section we briefly review recent work on the subject; in the following section relevant change will be combined with kinetic consistency.

Relevant change was studied by Parikh in [9], where a new postulate for it was introduced, called (P). Postulate (P) was further analysed in [11] and two different interpretations of it were identified, called the *weak* and the *strong* version of (P). The weak version of postulate (P), which we denote (wP), is much more general and intuitive, and it is this version we shall use herein.

Before presenting (wP) we need some more notation: for any sentence x, L_x denotes the (unique) smallest language in which x can be expressed. Moreover, $\overline{L_x}$ denotes the complement language, that is the language built from the propositional variables that do not appear in L_x. With this additional notation we can now present (wP):

(wP)   If $K = Cn(\{x, y\})$, $L_x \cap L_y = \varnothing$, and $\varphi \in L_x$, then $(K * \varphi) \cap \overline{L_x} = K \cap \overline{L_x}$.

Postulate (wP) essentially says the following. Suppose that the initial belief set K is divided into two disjoint compartments x, y, in the sense that the minimal languages in which the two sentences x and y can be expressed, do not share any propositional variable. If the epistemic input φ happens to be expressible solely within the language of the first compartment, then the second compartment remains unaffected by the revision of K by φ, since arguably, it is not relevant to the epistemic input.

In [11], (wP) was characterised semantically in terms of constrains over faithful preorders. This semantic characterisation of (wP) will be used later in the proof of Theorem 4 and is therefore presented below. First however we need some additional terminology and notation.

The *difference* between two possible worlds w, r, denoted $Diff(w, r)$, is defined to be the set of variables over which the two worlds disagree. Formally, $Diff(w, r) = \{q \in P : w \models q \text{ and } r \models \neg q\} \cup \{q \in P : w \models \neg q \text{ and } r \models q\}$.

The definition of *Diff* can be extended to include the difference between a *theory K* and a world r, [11]. For this however we need some more notation: for any nonempty set of propositional variables $S \subseteq P$, by L^S we denote the propositional language built from the variables in S.

Consider now a consistent theory K, and let $Q = \{Q_1, \ldots, Q_n\}$ be a partition of P; i.e., $\bigcup Q = P$, $Q_i \neq \varnothing$, and $Q_i \cap Q_j = \varnothing$, for all $1 \leqslant i \neq j \leqslant n$. We say that $Q = \{Q_1, \ldots, Q_n\}$ is a *K-splitting* iff there exist sentences $\phi_1 \in L^{Q_1}, \ldots, \phi_n \in L^{Q_n}$, such that

$K = Cn(\{\phi_1, \ldots, \phi_n\})$. Parikh has shown in [9] that for every theory K there is a unique *finest K-splitting*, i.e., one which refines every other K-splitting.[3]

We can now define the difference between an arbitrary consistent theory K and a world r using the finest splitting of K, call it F, as follows: $Diff(K, r) = \bigcup\{F_i \in F :$ for some $\phi \in L^{F_i}$, $K \models \phi$ and $r \models \neg\phi\}$ (see [11] for a detailed discussion on this definition). With the extended definition of *Diff*, it was shown in [11] that (wP) can be semantically characterised by the following two constraints:

(Q1)   If $Diff(K, r) \subset Diff(K, r')$ and $Diff(r, r') \cap Diff(K, r) = \varnothing$, then $r \prec r'$.

(Q2)   If $Diff(K, r) = Diff(K, r')$ and $Diff(r, r') \cap Diff(K, r) = \varnothing$, then $r \approx r'$.

Theorem 3 [11]. *Let $*$ be a revision function satisfying (K*1) - (K*8), K a consistent theory, and $\leqslant_K$ a preorder faithful to K, that corresponds to $*$ at K by means of ($\leqslant*$). Then $*$ satisfies (wP) at K iff $\leqslant_K$ satisfies (Q1) - (Q2).*

It was furthermore shown in [11] that (wP) is consistent with the AGM postulates (K*1) - (K*8). Our next aim herein is to show that (KC1) - (KC2) can also be added consistently to (K*1) - (K*8) and (wP). This is the subject of the next section.

6   Parametrised Difference Operators

To prove the consistency of the extended AGM postulates it suffices to show that there is at least one concrete revision function that satisfies them all. A good candidate would be Dalal's revision operator, [2], since it is already known to satisfy (K*1) - (K*8) and (wP), [11].

Yet consistency would be all that such a result would demonstrate; the *scope* of the extended AGM postulates would still be undetermined. We will therefore prove something stronger. We shall introduce a whole new class of revision operators, called *parametrised difference operators*, or PD operators for short, and show that each one of them satisfies (K*1) - (K*8), (wP), and (KC1) - (KC2). By doing so, we would not only prove the consistency of the extended AGM postulates, but we will also shed light to the scope and nature of the revision functions satisfying these postulates. This is important, since up to now it is not known how restrictive the addition of (wP) to (K*1) - (K*8) might be.

The PD operators are natural generalizations of Dalal's operator, so we will look at Dalal's operator first.

Dalal defines his operator, which we denote $\square$, as the function induced by means of ($\leqslant*$), from the following preorders (one for each theory K):

$r \sqsubseteq_K r'$ *iff*   there is a $w \in [K]$ such that for all $w' \in [K]$, $|Diff(w, r)| \leqslant |Diff(w', r')|$

---

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

Let us examine the above definition in more details through an example. Consider a language L built from only three variables a, b, c, and suppose that K is the theory $K = Cn(\{a, b, c\})$. Then the preorder that Dalal attaches to K is the following:

$$
abc \quad \sqsubset_K \quad
\begin{array}{c} ab\bar{c} \\ \bar{a}bc \\ \bar{a}bc \end{array}
\quad \sqsubset_K \quad
\begin{array}{c} \bar{a}b\bar{c} \\ \bar{a}b\bar{c} \\ \bar{a}b c \end{array}
\quad \sqsubset_K \quad \bar{a}\bar{b}\bar{c}
$$

According to Dalal, the plausibility of a world r is determined by the number of propositional variables on which r differs from the initial world abc. A silent assumption in Dalal's approach is that *all variables have the same epistemic value*; hence for example, a change in the variable a is assumed to be as plausible (or implausible) as a change in variable b. In many scenarios though this is not true. Suppose for example that K describes our beliefs about a circuit consisting of a multiplier and two adders. Variable a represents the fact that "adder1 is working", variable b that "adder2 is working", and variable c that "the multiplier is working". Given that multipliers are less reliable than adders, if we observe that the circuit is malfunctioning, it is plausible to put the blame on the multiplier rather than on one of the adders (this is a modified version of an example in [3]). In other worlds, a change in a or b is less plausible than a change in c. We can represent this with a preorder $\trianglelefteq$ over variables, were variables appearing latter in the preorder are *more resistant to change* than the ones appearing earlier: $c \trianglelefteq a$, $c \trianglelefteq b$, $a \trianglelefteq b$, $b \trianglelefteq a$, $a \trianglelefteq a$, $b \trianglelefteq b$, and $c \trianglelefteq c$.

Given $\trianglelefteq$ we can now refine Dalal's preorder in a way that takes into account the difference in epistemic value between the three propositional variables. The resulting preorder is denoted $\sqsubseteq_K^{\trianglelefteq}$ (with $\sqsubset_K^{\trianglelefteq}$ denoting its strict part), and it is shown below:

$$
abc \quad \sqsubset_K^{\trianglelefteq} \quad ab\bar{c} \quad \sqsubset_K^{\trianglelefteq} \quad
\begin{array}{c} \bar{a}bc \\ \bar{a}bc \end{array}
\quad \sqsubset_K^{\trianglelefteq} \quad
\begin{array}{c} \bar{a}b\bar{c} \\ \bar{a}b\bar{c} \end{array}
\quad \sqsubset_K^{\trianglelefteq} \quad \bar{a}\bar{b}c \quad \sqsubset_K^{\trianglelefteq} \quad \bar{a}\bar{b}\bar{c}
$$

Observe that according to $\sqsubseteq_K^{\trianglelefteq}$, $ab\bar{c}$ is more plausible than $\bar{a}bc$, although both worlds differ in only one variables from the initial world abc. This is because, according to $\trianglelefteq$, a change in a induces greater epistemic loss than a change in c. For similar reasons, $ab\bar{c}$ is more plausible than $\bar{a}bc$, although both worlds differ in two variables from abc.

The example above illustrates the basic idea in our generalisation of Dalal's approach with PD preorders. Given a user-defined preorder $\trianglelefteq$, the comparative plausibility of any two worlds r and r' relative to an initial belief set K, is determined, firstly, by the number of switches in the sign of variables that are needed to take us from a K-world to r and to r' respectively. If the number of necessary switches for r is smaller than the number of necessary switches for r', then r is defined to be strictly more plausible that r'. This first step is identical to Dalal's definition. The differentiation appears when the two worlds r and r' require the *same* number of switches: Dalal defines them to be equally plausible, whereas we take into account $\trianglelefteq$ to order them. More precisely, we define r to be more plausible than r' if the set of variables that need to be switched to reach r from K, *lexicographically proceeds* (with respect to $\trianglelefteq$) the set of variables that need to be switched to reach r' from K. Below we present this more formally.

Let $\trianglelefteq$ be any total preorder over P (the set of propositional variables). For a set of propositional variables S and a variable $q \in P$, by S_q we denote the set $S_q = \{x \in S : x \trianglelefteq q\}$. We can now extend the definition of $\trianglelefteq$ to *sets* of propositional variables. In particular, for any two sets of propositional variables $S, S' \subseteq P$, we define $S \trianglelefteq S'$ iff one of the following three conditions holds:

(a) $|S| < |S'|$.
(b) $|S| = |S'|$, and for all $q \in P$, $|S_q| = |S'_q|$.
(c) $|S| = |S'|$, and for some $q \in P$, $|S_q| > |S'_q|$, and for all $p \vartriangleleft q$, $|S_p| = |S'_p|$

In the above definition, condition (b) states that S and S' are lexicographically indistinguishable with respect to $\trianglelefteq$, whereas (c) states that S lexicographically proceeds S' (wrt $\trianglelefteq$). It is not hard to verify that $\trianglelefteq$ is a total preorder over 2^P; moreover the empty set precedes every other set with respect to $\trianglelefteq$.

We can now define the *PD preorder* $\sqsubseteq_K^{\trianglelefteq}$ over $\mathbb{M}$, induced from $\trianglelefteq$ and associated to a theory K, as follows:

$r \sqsubseteq_K^{\trianglelefteq} r'$ iff there is a $w \in [K]$ such that for all $w' \in [K]$, $Diff(w, r) \trianglelefteq Diff(w', r')$.

Observe that when $\trianglelefteq = P \times P$, then $\sqsubseteq_K^{\trianglelefteq}$ reduces to Dalal's preorder $\sqsubseteq_K$. In what follows we will prove that, for any total preorder $\trianglelefteq$ over P, the revision function induced by the PD preorders $\{\sqsubseteq_K^{\trianglelefteq}\}_{K \in \mathbb{K}}$ satisfies all the extended AGM postulates. To this aim we recall the following lemma from [9]:

Lemma A [9]**.** *Let K be a theory and $\{Q_1, \ldots, Q_n\}$ a partition of P. If $\{Q_1, \ldots, Q_n\}$ is a K-splitting, then for any $r_1, \ldots, r_n \in [K]$, $Mix(r_1, \ldots, r_n; Q_1, \ldots, Q_n)$ belongs to $[K]$. Conversely, if $Mix(r_1, \ldots, r_n; Q_1, \ldots, Q_n)$ belongs to $[K]$ for all $r_1, \ldots, r_n \in [K]$, then $\{Q_1, \ldots, Q_n\}$ is a $K - splitting$.*

In the lemma above, $Mix(r_1, \ldots, r_n; Q_1, \ldots, Q_n)$ denotes the unique world r that agrees with r_1 on the variables in Q_1, with r_2 on the variables in Q_2, ..., and with r_n on the variables in Q_n.

We can now prove the theorem alluded earlier.

Theorem 4. *Let $\trianglelefteq$ be a total preorder over P and $*$ the revision function induced from the family of PD preorders $\{\sqsubseteq_K^{\trianglelefteq}\}_{K \in \mathbb{K}}$. Then $*$ satisfies (K*1) - (K*8), (wP), and (KC1) - (KC2).*

Proof. To prove that $*$ satisfies (K*1) - (K*8), it suffices to show that for any consistent theory K, $\sqsubseteq_K^{\trianglelefteq}$ is a total preorder and moreover that it is faithful to K.

For reflexivity, let r be any possible world. Define Q to be the set $Q = \{Diff(z, r) : z \in [K]\}$. Since K is consistent, $Q \neq \emptyset$. Let $Diff(w, r)$ be a minimal element of Q with respect to $\trianglelefteq$.[4] Clearly then, $w \in [K]$ and $Diff(w, r) \trianglelefteq Diff(w', r)$, for all $w' \in [K]$. Hence $r \sqsubseteq_K^{\trianglelefteq} r$ as desired.

For transitivity, let r, r', r'' be any three worlds such that $r \sqsubseteq_K^{\trianglelefteq} r' \sqsubseteq_K^{\trianglelefteq} r''$. From $r' \sqsubseteq_K^{\trianglelefteq} r''$ we derive that there is a world $w' \in [K]$ such that $Diff(w', r') \trianglelefteq Diff(w'', r'')$,

---

[4] Since $\trianglelefteq$ is a total preorder and $[K]$ is finite (because P is assumed to be finite), such a minimal element always exists.

for all $w'' \in [K]$. Moreover, from $r \sqsubseteq_K^\trianglelefteq r'$ it follows that there is a $w \in [K]$ such that $Diff(w, r) \trianglelefteq Diff(w', r')$. Hence, since $\trianglelefteq$ is transitive, $Diff(w, r) \trianglelefteq Diff(w'', r'')$, for all $w'' \in [K]$. Consequently, $r \sqsubseteq_K^\trianglelefteq r''$, and therefore $\sqsubseteq_K^\trianglelefteq$ is transitive.

For totality, let r, r' be any two worlds and assume that $r' \not\sqsubseteq_K^\trianglelefteq r$. Let Q be the set $Q = \{Diff(z, r) : z \in [K]\}$. Since K is consistent, $Q \neq \varnothing$. Let $Diff(w, r)$ be a minimal element of Q with respect to $\trianglelefteq$. Then $w \in [K]$ and $Diff(w, r) \trianglelefteq Diff(z, r)$, for all $z \in [K]$. Next consider any world $w' \in [K]$. From $r' \not\sqsubseteq_K^\trianglelefteq r$, it follows that there is a $z \in [K]$ such that $Diff(w', r') \not\trianglelefteq Diff(z, r)$, and consequently, since $\trianglelefteq$ is a *total* preorder, $Diff(z, r) \trianglelefteq Diff(w', r')$. On the other hand, by the definition of w, $Diff(w, r) \trianglelefteq Diff(z, r)$. Consequently, by the transitivity of $\trianglelefteq$, $Diff(w, r) \trianglelefteq Diff(w', r')$. Since w' was chosen arbitrarily, it follows that $r \sqsubseteq_K^\trianglelefteq r'$, and therefore $\sqsubseteq_K^\trianglelefteq$ is total.

Next we show that $\sqsubseteq_K^\trianglelefteq$ is faithful to K. Let r, r' be any two worlds. Moreover assume that $r \in [K]$. Clearly, $Diff(r, r) = \varnothing$ and consequently, $Diff(r, r) \trianglelefteq Diff(w', r')$ for all $w' \in [K]$; thus, since $r \in [K]$, $r \sqsubseteq_K^\trianglelefteq r'$. This clearly entails the first condition for faithfulness. For the second condition, assume that r, r' are two worlds such that $r \in [K]$ and $r' \notin [K]$. Then $Diff(r, r) = \varnothing$ and $Diff(w', r) \neq \varnothing$ for all $w' \in [K]$. Hence, for all $w' \in [K]$, $|Diff(r, r)| < |Diff(w'r')|$, which again entails $r \sqsubset_K^\trianglelefteq r'$ as desired.

We have thus shown that $\sqsubseteq_K^\trianglelefteq$ is a total preorder that is faithful to K. By Theorem 1 it then follows that the induced revision function $*$ satisfies (K*1) - (K*8).

For (KC1), consider two arbitrary consistent theories K, H. It suffices to show that (KS1) is satisfied. Let r, r' be any two worlds such that $r \sqsubseteq_K^\trianglelefteq r'$ and $r \sqsubseteq_H^\trianglelefteq r'$. From $r \sqsubseteq_K^\trianglelefteq r'$ it follows that there is a $w_1 \in [K]$ such that $Diff(w_1, r) \trianglelefteq Diff(w', r')$, for all $w' \in [K]$. Similarly, form $r \sqsubseteq_H^\trianglelefteq r'$ it follows that there is a $w_2 \in [H]$ such that $Diff(w_2, r) \trianglelefteq Diff(w'', r')$, for all $w'' \in [H]$. Since $\trianglelefteq$ is total, $Diff(w_1, r) \trianglelefteq Diff(w_2, r)$ or $Diff(w_2, r) \trianglelefteq Diff(w_1, r)$. We assume without loss of generality that $Diff(w_1, r) \trianglelefteq Diff(w_2, r)$. Then from the transitivity of $\trianglelefteq$ we derive that $Diff(w_1, r) \trianglelefteq Diff(z, r')$, for all $z \in [K] \cup [H]$. Given that $[K] \cup [H] = [K \cap H]$, we then derive that $r \sqsubseteq_{K \cap H}^\trianglelefteq r'$ as desired.

For (KC2), consider two arbitrary consistent theories K, H. It suffices to show that (KS2) is satisfied. This is done with a totally symmetric argument as the one above. In particular, let r, r' be two worlds such that $r \sqsubset_K^\trianglelefteq r'$ and $r \sqsubset_H^\trianglelefteq r'$. From $r \sqsubset_K^\trianglelefteq r'$ it follows that there is a $w_1 \in [K]$ such that $Diff(w_1, r) \triangleleft Diff(w', r')$, for all $w' \in [K]$. Similarly, form $r \sqsubset_H^\trianglelefteq r'$ it follows that there is a $w_2 \in [H]$ such that $Diff(w_2, r) \triangleleft Diff(w'', r')$, for all $w'' \in [H]$. Since $\trianglelefteq$ is total, $Diff(w_1, r) \trianglelefteq Diff(w_2, r)$ or $Diff(w_2, r) \trianglelefteq Diff(w_1, r)$. We assume without loss of generality that $Diff(w_1, r) \trianglelefteq Diff(w_2, r)$. Then from the transitivity of $\trianglelefteq$ we derive that $Diff(w_1, r) \triangleleft Diff(z, r')$, for all $z \in [K] \cup [H]$. Given that $[K] \cup [H] = [K \cap H]$, we then derive that $r \sqsubset_{K \cap H}^\trianglelefteq r'$ as desired.

Finally for (wP), we shall prove instead that conditions (Q1) and (Q2) are satisfied. The proof of (Q1) follows that same line of reasoning used in the proof of Theorem 7 in [11]; the proof of (Q2) is somewhat different.

Starting with condition (Q1), let K be a consistent theory and let r, r' be any two possible worlds such that $Diff(K, r) \subset Diff(K, r')$ and $Diff(r, r') \cap Diff(K, r) = \varnothing$. Then clearly, $P - Diff(K, r) \neq \varnothing$. Let u be a world in K that agrees with r on all variables in

$P - Diff(K, r)$. [5] Moreover, let z be a K-world that differs in the least number of variables from r when restricted to $Diff(K, r)$; i.e., $|Diff(z, r) \cap Diff(K, r)| \leqslant |Diff(z', r) \cap Diff(K, r)|$ for all $z' \in [K]$. Define w to be the world that agrees with z on the variables in $Diff(K, r)$ and agrees with u on the remaining variables. Clearly, $Diff(w, r) \subseteq Diff(K, r)$. Moreover, by the definition of $Diff$, $\{Diff(K, r), P - Diff(K, r)\}$ is a K-splitting, and consequently from Lemma A and the fact that $z, u \in [K]$, we derive that $w \in [K]$.

Consider now any world $w' \in [K]$. Since $Diff(K, r) \subset Diff(K, r')$, there is at least one sentence x, built entirely from variables in $P - Diff(K, r)$, such that $K \models x$ and $r' \models \neg x$. Hence from $w' \in [K]$, we derive that $Diff(w', r') \cap (P - Diff(K, r)) \neq \varnothing$ and consequently, $|Diff(w', r') \cap (P - Diff(K, r))| > 0$. Moreover, since r and r' agree on the variables in $Diff(K, r)$ it follows that $|Diff(w', r') \cap Diff(K, r)| = |Diff(w', r) \cap Diff(K, r)|$. Consequently, $|Diff(w', r')| = |Diff(w', r') \cap Diff(K, r)| + |Diff(w', r') \cap (P - Diff(K, r))| > |Diff(w', r') \cap Diff(K, r)| = |Diff(w', r) \cap Diff(K, r)| \geq |Diff(z, r) \cap Diff(K, r)| = |Diff(w, r) \cap Diff(K, r)| = |Diff(w, r)|$. Hence we have shown that $|Diff(w, r)| < |Diff(w', r')|$ for all $w' \in [K]$. This again entails that $r \sqsubset_K^{\trianglelefteq} r'$ as desired.

For (Q2), let K be a consistent theory and let r, r' be any two possible worlds such that $Diff(K, r) = Diff(K, r')$ and $Diff(r, r') \cap Diff(K, r) = \varnothing$. If $Diff(K, r) = P$ then $r = r'$ and (Q2) trivially holds. Moreover, if $Diff(K, r) = \varnothing$, then $r, r' \in [K]$ and therefore (Q2) follows from the fact that $\sqsubseteq_K^{\trianglelefteq}$ is faithful to K. Assume therefore that $\varnothing \neq Diff(K, r) \subset P$.

Let Q be the set, $Q = \{Diff(z, r) : z \in [K]\}$ and let $Diff(w, r)$ be a minimal element of Q with respect of $\trianglelefteq$. Then, $w \in [K]$ and for all $z \in [K]$, $Diff(w, r) \trianglelefteq Diff(z, r)$.

Next we show that $Diff(w, r) \subseteq Diff(K, r)$. Assume on the contrary that $Diff(w, r) \cap (P - Diff(K, r)) \neq \varnothing$. Since r does not differ from K in any variables in $P - Diff(K, r)$, we derive that there is a $u \in [K]$ that agrees with r on all variables in $P - Diff(K, r)$ (see Footnote 5). Define z to be the world that agrees with w on the variables in $Diff(K, r)$ and agrees with u on the remaining variables. Then $Diff(z, r) \subset Diff(w, r)$ and moreover, since $\{Diff(K, r), P - Diff(K, r)\}$ is a K-splitting, by Lemma A, $z \in [K]$. This however contradicts our assumption that $Diff(w, r)$ is $\trianglelefteq$-minimal in Q. Hence we have shown that $Diff(w, r) \subseteq Diff(K, r)$; i.e., w agrees with r over all variables in $P - Diff(K, r)$.

Now pick a world $w' \in [K]$ such that $Diff(w', r')$ is $\trianglelefteq$-minimal in the set $Q' = \{Diff(z', r') : z' \in [K]\}$. By a similar argument as the one above, we derive that $Diff(w', r') \subseteq Diff(K, r)$.

Next we show that $Diff(w, r) \trianglelefteq Diff(w', r')$. Assume towards contradiction that $Diff(w', r') \lhd Diff(w, r)$. Define z to be the world that agrees with w' over the variables in $Diff(K, r)$, and it agrees with w over all remaining variables. By Lemma A, $z \in [K]$. Moreover by construction, $Diff(z, r) \subseteq Diff(K, r)$. Hence, since r and r' agree over the variables in $Diff(K, r)$, and so do z and w', we derive that $Diff(z, r) = Diff(w', r')$. From $Diff(w', r') \lhd Diff(w, r)$ we then derive that $Diff(z, r) \lhd Diff(w, r)$. This of course contradicts our initial assumption that $Diff(w, r)$ is $\trianglelefteq$-minimal in $\{Diff(z', r) : z' \in [K]\}$.

---

[5] To see that such a world indeed exists, consider the sentence φ defined as the conjunction of all literals in r that are built from variables in $P - Diff(K, r)$. Clearly then, $r \models \psi$. Moreover, $\neg \psi \notin K$, for otherwise $Diff(K, r)$ would include variables from $P - Diff(K, r)$, which is of course a contradiction. Hence there is a $u \in [K]$ such that $u \models \psi$. By the construction of ψ it follows that u agrees with r on all variables outside $Diff(K, r)$.

Thus we have shown that $Diff(w, r) \trianglelefteq Diff(w', r')$. Since $Diff(w', r')$ is $\trianglelefteq$-minimal in $\{Diff(z', r') : z' \in [K]\}$ we derive that $Diff(w, r) \trianglelefteq Diff(z', r')$, for all $z' \in [K]$. Consequently, $r \sqsubseteq_K^{\trianglelefteq} r'$.

By a totally symmetric argument we also derive that $r' \sqsubseteq_K^{\trianglelefteq} r$, thus proving (Q2). □

7 Examples

The extended AGM postulates are clearly stronger than the original ones. In [9], Parikh has already demonstrated the benefits of adding (wP) to (K*1) - (K*8). In this section we provide further examples in support of the extended AGM postulates.

Yet instead of following the path set by Parikh, we choose a different direction which better illustrates the diversity of the implications of adding (wP) and (KC1) - (KC2) to the classical AGM postulates. In particular, the first two examples below, introduced in [5, 7] respectively, relate to *iterated revision* scenarios, [12], which are known to be out of reach for vanilla AGM. It is shown below that the extended AGM postulates, although not specifically designed to deal with iteration, are nevertheless strong enough to produce the desired conclusions in these examples.

Example 1 [7]. "Consider a circuit containing an adder and a multiplier. In this example, we have two atomic propositions, *adder_ok* and *multiplier_ok*, denoting respectively the fact that the adder and the multiplier are working. We have initially no information about this circuit ($\Psi \equiv \top$) and we learn that the adder and the multiplier are working ($\mu = adder\_ok \wedge multiplier\_ok$). Then someone tells us that the adder is not working ($\alpha = \neg adder\_ok$). There is, then, no reason to "forget" that the multiplier is working."

The extended AGM postulates turn the right results: since initially $\Psi = Cn(\varnothing)$, μ is consistent with Ψ, and therefore $\Psi * \mu = Cn(\{adder\_ok, multiplier\_ok\})$; (wP) then entails that $\Psi * \mu * \alpha = Cn(\{\neg adder\_ok, multiplier\_ok\})$ as desired.

Example 2 [5]. "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's 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."

Once again the extended AGM postulates deliver the anticipated results. Let us denote by a the proposition "the animal is red" and by b the proposition "the animal is a bird". Our initial belief set is $K = Cn(\{b\})$. Since a is consistent with K it follows that $K * a = Cn(\{a, b\})$. Hence (wP) entails that $K * a * \neg b = Cn(\{a, \neg b\})$ as desired.

It should be noted that, not only classical AGM, but even Darwiche and Pearl's iterated revision approach, [3], has trouble dealing such examples (see [5, 7, 8] for details).

In Examples 1 and 2, the addition of (wP) alone to (K*1) - (K*8), suffices to produce the desired results. The last example is new and requires the full strength of the extended AGM postulates.

Example 3. A circuit consists of a multiplier and two adders. Let us denote by m the proposition "the multiplier is working", and by a_1, a_2 the propositions "the first adder is working" and "the second adder is working" respectively. After performing some tests on the circuit, we discover that the multiplier or adder 1 is malfunctioning; in symbols, $\neg m \lor \neg a_1$. Suppose that our initial belief set is $Cn(\{m, a_1, a_2\})$. Since it is known that multipliers are less reliable than adders, we end up with the belief set $Cn(\{\overline{m}, a_1, a_2\})$. For the same reason, if our initial belief set was $Cn(\{m, a_1, \overline{a}_2\})$, then $\neg m \lor \neg a_1$ would have taken us to $Cn(\{\overline{m}, a_1, \overline{a}_2\})$. What would then be our response to $\neg m \lor \neg a_1$ had our initial belief set been $Cn(\{m, a_1\})$? Given our past preference to adder 1 over the multiplier (regardless of the status of adder 2), we argue that it is reasonable to once again put the blame on the multiplier. Moreover, since adder 2 is independent from the other two components, our beliefs about adder 2 should not be effected. Indeed from (KC2) (or rather (KS2)) and (wP) we derive that the resulting belief set is indeed $Cn(\{\overline{m}, a_1\})$ as desired.

8   Conclusion

There are three main contributions in this paper. Firstly, we identified and formalised an aspect of rational belief revision that has been neglected in the classical AGM framework. We call it the *principle of kinetic consistency*. We modelled this principle axiomatically and semantically and proved a representation result connecting the two.

Secondly, we investigated the model that results from combining the AGM postulates for revision (K*1) - (K*8), with the postulates for kinetic consistency (KC1) - (KC2), and Parikh's postulate (wP) for relevant change. The extended AGM postulates were shown to be consistent and their strength was demonstrated in a number of examples that are out of reach for vanilla AGM.

Thirdly, we introduced a whole new class of concrete revision operators, called *PD operators* that are natural generalisations of Dalal's revision operator. We proved that every PD operator satisfies all extended AGM postulates. This result not only established the consistency of the extended AGM postulates, but also sheds light into the nature and scope of the revision functions satisfying these postulates.

We conclude with a note on future work. Firstly, we intend to examine more thoroughly the relationship revealed by the examples of the previous section, between the extended AGM postulates and iterated revision. Secondly, in future work we will investigate the possibility for an axiomatic characterisation of PD operators. Theorem 4 shows that all PD operators satisfy the extended AGM postulates; the converse however may not be true. If so, new axioms need to be formulated to characterise parametrised difference revision.

References

1. Alchourron, C., Gardenfors, P., Makinson, D.: On the logic of theory change: partial meet functions for contraction and revision. J. Symbol. Logic **50**, 510–530 (1985)
2. Dalal, M.: Investigations into theory of knowledge base revision: preliminary report. In: Proceedings of 7th National Conference of the American Association for Artificial Intelligence (AAAI 1988), pp. 475–479 (1988)
3. Darwiche, A., Pearl, J.: On the logic of iterated belief revision. Artif. Intell. **89**, 1–29 (1997)
4. Grove, A.: Two modellings for theory change. J. Philos. Logic **17**, 157–170 (1988)
5. Jin, Y., Thielscher, M.: Iterated belief revision, revised. Artif. Intell. **171**, 1–18 (2007)
6. Katsuno, H., Mendelzon, A.: Propositional knowledge base revision and minimal change. Artif. Intell. **52**(3), 263–294 (1991)
7. Konieczny, S., Perez, R.P.: A framework for iterated revision. J. Appl. Non-Classical Logics **10**, 339–367 (2000)
8. Nayak, A., Pagnucco, M., Peppas, P.: Dynamic belief revision operators. Artif. Intell. **146**, 193–228 (2003)
9. Parikh, R.: Beliefs, belief revision, and splitting languages. In: Logic, Language, and Computation - CSLI Lecture Notes, vol. 2, pp. 266–278. CSLI Publications (1999)
10. Peppas, P.: Belief revision. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) Handbook of Knowledge Representation, pp. 317–359. Elsevier Science (2008)
11. Peppas, P., Williams, M.-A., Chopra, S., Foo, N.: Relevance in belief revision. Artif. Intell. **229**, 126–138 (2015)
12. Peppas, P.: A panorama of iterated revision. In: Ove Hansson, S. (ed.) David Makinson on Classical Methods for Non-Classical Problems. Outstanding Contributions to Logic, pp. 71–94. Springer, Netherlands (2014)