BULLETIN OF THE
HELLENIC MATHEMATICAL SOCIETY
Volume 66, 2022 (1–10)

BELIEF REVISION AGAINST BACKGROUND KNOWLEDGE

THEOFANIS ARAVANIS

Abstract. In a cornerstone work, Katsuno and Mendelzon defined a revision-process according to which a rational agent always adheres to a given belief corpus that represents background knowledge. In this article, we provide an axiomatic characterization of this type of revision-process by formulating a natural generalization of the well-known postulates for rational belief revision, introduced by Alchourrón, Gärdenfors and Makinson. The proposed axiomatic characterization is shown to be highly relevant to a modification of a set of rationality constraints for multiple belief revision. Furthermore, we introduce a class of rational revision functions, along with a special type of non-prioritized revision, both of which have a smooth and reasonable connection with the aforementioned generalized revision-framework.

1. Introduction

Belief revision (or simply revision) is the process by which a rational agent changes her/his beliefs, in the light of new information [10]. A prominent approach that formalizes belief revision is that proposed by Alchourrón, Gärdenfors and Makinson in [1], now known as the *AGM paradigm*. Within the AGM paradigm, the agent's belief corpus is modelled as a logical theory K of an underlying logic language, also referred to as *belief set*, new information (alias, *epistemic input*) is represented as a logical sentence φ, and the revision of K by φ is modelled as a (revision) function $*$ that maps K and φ to a revised theory $K * \varphi$. A collection of eight postulates, formulated in [1] and called the *AGM postulates for revision*, *axiomatically* characterizes any *rational* revision operator, which is named *AGM revision function*. From a *semantic* point of view, it has been proven that any AGM revision function can be *constructed* (specified) by means of a special kind of total preorders over possible worlds, called *faithful preorders* [14].

Since the proposal of the AGM paradigm, several important aspects of the revision-process have been extensively studied in the literature, such as *relevance-sensitive* [11, 18, 22, 17, 5, 3] or *iterated* revision [8, 20], whereas, a variety of concrete "off-the-shelf" revision operators have been introduced [23, 7, 21, 4, 6, 2].

In this article, we study the case in which, during revision, a rational agent always adheres to a given (consistent) *background theory*, which essentially represents given *background knowledge* of a domain or application. On that premise, the agent's initial, as well as any revised, state of belief has to *include* (*entail*) all the information contained in the background theory. Undoubtedly, such a scenario is quite realistic, with the background theory occasionally taking the

Date: Received October 16, 2021. Accepted January 30, 2022.
2020 *Mathematics Subject Classification.* 03B42, 68T27.
Key words and phrases. Belief Revision, Background Knowledge, Multiple Belief Revision, Non-Prioritized Belief Revision, Knowledge Representation.
Acknowledgements. I am grateful to Pavlos Peppas for many fruitful discussions on this work, and to the anonymous referee for their valuable feedback.

form of a set of integrity constraints, a well-established scientific discipline, a corpus of moral statements, and so forth.

Already in [14], Katsuno and Mendelzon presented a *modification* of AGM revision functions to ensure that revision is implemented against background knowledge. Herein, we *axiomatically* characterize Katsuno and Mendelzon's modified revision operators, by formulating a natural *generalization* of the AGM paradigm. We show that this generalized AGM paradigm can be derived by Lindström's rationality postulates for *multiple* belief revision [15]. A natural class of AGM revision functions, as well as a special type of *non-prioritized* revision operator, are also introduced, which are shown to have a smooth and intuitive connection with the proposed generalized AGM paradigm.[1]

The remainder of this article is structured as follows. The next section fixes the basic notations and conventions that shall be used for our subsequent discussion. Thereafter, Section 3 introduces the AGM paradigm, followed by Section 4 which presents a natural generalization of the AGM paradigm that handles revision under background knowledge. Section 5 identifies the close relation between the proposed generalized revision-framework and multiple belief revision. Sections 6 and 7 introduce an interesting class of AGM revision functions and a special type of non-prioritized revision, respectively, and point out their relevance to revision under background knowledge. A brief conclusion closes the paper.

2. Basic Notations and Conventions

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

For a set of sentences Γ of $\mathcal{L}$, $Cn(\Gamma)$ denotes the set of all logical consequences of Γ; i.e., $Cn(\Gamma) = \{\varphi \in \mathcal{L} : \Gamma \models \varphi\}$. An agent's belief corpus shall be modelled by a *theory*, also referred to as a *belief set*. A theory K is any deductively closed set of sentences of $\mathcal{L}$; i.e., $K = Cn(K)$. The set of all theories is denoted by $\mathbb{K}$. For a set of sentences Γ of $\mathcal{L}$, $\bigwedge \Gamma$ denotes the sentence of $\mathcal{L}$ resulting from the *conjunction* of all the elements in Γ.

We will assume *consistent* background theories; such a theory shall be denoted by $\mathcal{T}$.[2] Given that the set of propositional variables $\mathcal{P}$ is finite, although a theory $\mathcal{T}$ is a infinite object, it is *finite modulo logical equivalence*; that is to say, there exists a consistent sentence of $\mathcal{L}$, denoted in what follows by τ, such that $\mathcal{T} = Cn(\{\tau\})$. Roughly speaking, the sentence τ answers the question "What is theory $\mathcal{T}$ actually about?".

A *literal* is a propositional variable $p \in \mathcal{P}$ or its negation. A *possible world* (or simply *world*) r is any 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 (set of sentences) φ of $\mathcal{L}$, $[\varphi]$ is the set of worlds at which φ is true.

A *preorder* over a set V is any reflexive, transitive binary relation in V. A preorder $\preceq$ is called *total* iff, for all $r, r' \in V$, $r \preceq r'$ or $r' \preceq r$. As usual, the strict part of $\preceq$ shall be denoted

[1]Multiple revision refers to the revision of a theory by a non-empty (possibly infinite) *set* of sentences [9], whereas, non-prioritized revision is a type of belief change, according to which the new information is *not always* accepted in the revised state of belief [12].

[2]The case of an inconsistent background theory is not interesting, since, in that case, the agent ought to adhere to absurd knowledge.

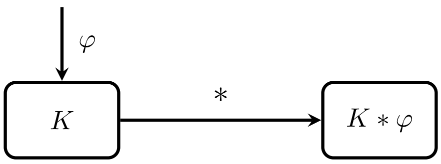

FIGURE 1. Belief revision within the AGM paradigm.

by $\prec$; namely, $r \prec r'$ iff $r \preceq r'$ and $r' \not\preceq r$. Also, $min(V, \preceq)$ denotes the set of all $\preceq$-minimal elements of V; i.e.,

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

3. The AGM Paradigm

Within the AGM paradigm [1], the process of belief revision is modelled as a binary function $*$ mapping a theory K and a sentence φ to a revised theory $K * \varphi$; i.e., $* : \mathbb{K} \times \mathcal{L} \mapsto \mathbb{K}$ (Figure 1). *Rational* revision functions, the so-called *AGM revision functions*, are those constrained by a set of eight postulates, called *AGM postulates for revision.* Before presenting these postulates, let us first define the notion of *expansion* of a belief set.

Definition 1 (Expansion)**.** For a theory K and a sentence φ of $\mathcal{L}$, the expansion of K by φ, denoted by $K + \varphi$, is defined to be the deductive closure of the set $K \cup \{\varphi\}$; i.e.,

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

On that premises, the AGM postulates for revision are the constraints $(K * 1)$–$(K * 8)$, listed below.[3]

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

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

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

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

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

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

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

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

Katsuno and Mendelzon proved that the revision functions that satisfy postulates $(K * 1)$–$(K * 8)$ are precisely those that are induced by means of a special type of total preorders over all possible worlds, called *faithful preorders* [14].

Definition 2 (Faithful Preorder, [14])**.** A total preorder $\preceq_K$ over $\mathbb{M}$ is faithful to a theory K iff the $\preceq_K$-minimal worlds are those satisfying K; i.e., $min(\mathbb{M}, \preceq_K) = [K]$.

[3]A detailed discussion on postulates $(K * 1)$–$(K * 8)$ can be found in [10, Section 3.3] or [19, Section 8.3.1].

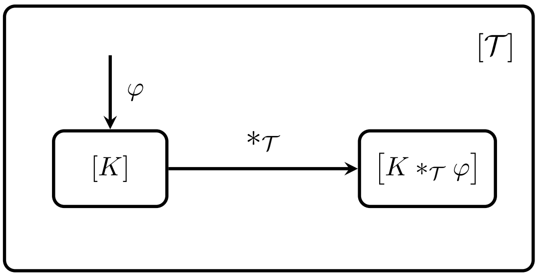

FIGURE 2. Belief revision against background theories.

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

Definition 3 (Faithful Assignment, [14])**.** A faithful assignment is a function that maps each theory K to a total preorder $\preceq_K$ over $\mathbb{M}$, that is faithful to K.

The following representation theorem precisely characterizes the class of AGM revision functions, in terms of faithful preorders.

Theorem 1 ([14])**.** *A revision function $*$ satisfies postulates $(K*1)$–$(K*8)$ iff there exists a faithful assignment that maps each theory K to a total preorder $\preceq_K$ over $\mathbb{M}$, such that, for any $\varphi \in \mathcal{L}$:*

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

Essentially, condition (F$*$) defines the revised belief set $K * \varphi$ as the theory corresponding to the most $\preceq_K$-plausible (i.e., $\preceq_K$-minimal) φ-worlds.

4. A Generalization of the AGM Paradigm

Let $\mathcal{T}$ be a (consistent) background theory, and let K be a belief set representing the belief corpus of a rational agent. Given that the agent *always* adheres to $\mathcal{T}$, it is true that $\mathcal{T} \subseteq K$; that is to say, all K-worlds are $\mathcal{T}$-worlds as well ($[K] \subseteq [\mathcal{T}]$). Katsuno and Mendelzon, in [14, Section 8], specified a revision function, herein denoted by $*_{\mathcal{T}}$, which revises K by an epistemic input $\varphi \in \mathcal{L}$, against (a *finite representation* of) theory $\mathcal{T}$.[4] As in the case of classical AGM revision functions, the revision function $*_{\mathcal{T}}$ maps a theory K and a sentence φ to a revised theory $K *_{\mathcal{T}} \varphi$, such that $\mathcal{T} \subseteq K *_{\mathcal{T}} \varphi$ (or equivalently $[K *_{\mathcal{T}} \varphi] \subseteq [\mathcal{T}]$); see Figure 2. As illustrated in Figure 2, the commitment to a background theory, during revision, essentially excludes possible worlds that are *not valid* relative to a specific domain or application, limiting accordingly the set $\mathbb{M}$ to some valid portion of it.

For each AGM revision function $*$, Katsuno and Mendelzon defined a revision function $*_{\mathcal{T}}$, through condition (KM) presented below.

$$\textbf{(KM)} \quad K *_{\mathcal{T}} \varphi = K * (\varphi \wedge \tau).$$

[4]Notice that Katsuno and Mendelzon, in [14], work with sentences, rather than theories. Since we have considered a *finite* set of propositional variables, the difference is immaterial — recall that, for any background theory $\mathcal{T}$, there exists a sentence $\tau \in \mathcal{L}$, such that $\mathcal{T} = Cn(\{\tau\})$.

In this section, our aim is to *axiomatically* characterize the revision functions defined via condition (KM). Accordingly, we shall show that a natural *generalization* of the AGM postulates for revision suffices to precisely characterize a revision function $*_\mathcal{T}$. Before presenting this characterization, let us define the *expansion* of a belief set, against a background theory.

Definition 4 (Expansion Against Background Theories)**.** Let $\mathcal{T}$ be a background theory. For a theory K and a sentence φ of $\mathcal{L}$, the expansion of K by φ, against $\mathcal{T}$, denoted by $K +_\mathcal{T} \varphi$, is defined to be the deductive closure of the set $K \cup \{\varphi\} \cup \mathcal{T}$; i.e.,

$$K +_\mathcal{T} \varphi = Cn\Big(K \cup \{\varphi\} \cup \mathcal{T}\Big).$$

Analogously to the classical expansion of Definition 1, the expansion of a belief set, against a background theory, is defined to be the smallest deductively closed set containing the initial theory, the new information, as well as the background theory. It is noteworthy that $K +_\mathcal{T} \varphi = K + (\varphi \wedge \tau)$.

Now, we are ready to present the alluded generalization of the AGM postulates for revision, which is encoded in the following collection of postulates.

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

($\mathbf{K} *_\mathcal{T} \mathbf{2}$) $\mathcal{T} \cup \{\varphi\} \subseteq K *_\mathcal{T} \varphi$.

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

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

($\mathbf{K} *_\mathcal{T} \mathbf{5}$) $K *_\mathcal{T} \varphi$ is inconsistent iff $\mathcal{T} \cup \{\varphi\}$ is inconsistent.

($\mathbf{K} *_\mathcal{T} \mathbf{6}$) If $\mathcal{T} \models \varphi \leftrightarrow \psi$, then $K *_\mathcal{T} \varphi = K *_\mathcal{T} \psi$.

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

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

Each one of postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ generalizes each one of postulates $(K * 1)$–$(K * 8)$, respectively; observe that, in the special case where the background theory $\mathcal{T}$ contains *only tautological sentences*, postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ reduce to $(K * 1)$–$(K * 8)$, respectively.

Following the vein of [19], some comments on $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ are in order. Postulate $(K *_\mathcal{T} 1)$ simply requires that the output of the revision-process, against a background theory $\mathcal{T}$, indeed be a belief set. Postulate $(K *_\mathcal{T} 2)$ ensures that the epistemic input φ, along with the background theory $\mathcal{T}$, are believed in the revised belief set $K *_\mathcal{T} \varphi$. Postulates $(K *_\mathcal{T} 3)$ and $(K *_\mathcal{T} 4)$, viewed together, state that, whenever the new information φ does not contradict the initial belief set K (which of course contains $\mathcal{T}$), there is no reason to remove any of the original beliefs at all; the new belief set $K *_\mathcal{T} \varphi$ is equal to the expansion of K by φ, against $\mathcal{T}$; i.e., $K *_\mathcal{T} \varphi = K +_\mathcal{T} \varphi$. Postulate $(K *_\mathcal{T} 5)$ says that the only case where $K *_\mathcal{T} \varphi$ is inconsistent is when theory $\mathcal{T} \cup \{\varphi\}$ is inconsistent. Postulate $(K *_\mathcal{T} 6)$ ensures that, when the logical equivalence $\varphi \leftrightarrow \psi$ is included in $\mathcal{T}$ (that is to say, the "content" of φ, ψ is identical according to $\mathcal{T}$), then the $*_\mathcal{T}$-revision of K by φ or by ψ, against $\mathcal{T}$, leads to identical results — notice that $\mathcal{T} \models \varphi \leftrightarrow \psi$ iff $Cn\big(\mathcal{T} \cup \{\varphi\}\big) = Cn\big(\mathcal{T} \cup \{\psi\}\big)$. Lastly, postulates $(K *_\mathcal{T} 7)$ and $(K *_\mathcal{T} 8)$ are essentially a generalization of $(K *_\mathcal{T} 3)$ and $(K *_\mathcal{T} 4)$, respectively. Viewed together, $(K *_\mathcal{T} 7)$ and $(K *_\mathcal{T} 8)$ say that, for any two sentences φ and ψ of $\mathcal{L}$, if in revising the initial belief set

K by φ, against $\mathcal{T}$, one is lucky enough to reach a belief set $K *_\mathcal{T} \varphi$ that is *consistent* with ψ, then, to produce $K *_\mathcal{T} (\varphi \wedge \psi)$, all that one needs to do is to expand $K *_\mathcal{T} \varphi$ by ψ, against $\mathcal{T}$.

On that premises, Theorem 2 is a representation theorem which precisely characterizes the revision functions that satisfy postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$, in terms of AGM revision functions.

Theorem 2. *Let $\mathcal{T}$ be a background theory, such that $\mathcal{T} = Cn\big(\{\tau\}\big)$. Moreover, let $*$ be an AGM revision function, K be a theory, and φ be a sentence of $\mathcal{L}$. A revision function $*_\mathcal{T}$ satisfies postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ iff $K *_\mathcal{T} \varphi = K * (\varphi \wedge \tau)$.*

Proof. The left-to-right implication follows straightforwardly, in view of the fact that, for a background theory $\mathcal{T}$ and a sentence χ of $\mathcal{L}$, $Cn\big(\mathcal{T} \cup \{\chi\}\big) = Cn\big(\{\chi \wedge \tau\}\big)$. As far as the right-to-left implication is concerned, it suffices to show that postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ are equivalent to postulates $(K*1)$–$(K*8)$, respectively, if the sentences φ and ψ in $(K*1)$–$(K*8)$ are replaced with the sentences $\varphi \wedge \tau$ and $\psi \wedge \tau$, respectively. This follows also straightforwardly, in view of the fact that $Cn\big(\mathcal{T} \cup \{\chi\}\big) = Cn\big(\{\chi \wedge \tau\}\big)$. □

Theorems 1 and 2 imply the following corollary.

Corollary 1. *Let $\mathcal{T}$ be a background theory, such that $\mathcal{T} = Cn\big(\{\tau\}\big)$. A revision function $*_\mathcal{T}$ satisfies postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ iff there exists a faithful assignment that maps each theory K to a total preorder $\preceq_K$ over $\mathbb{M}$, such that, for any $\varphi \in \mathcal{L}$:*

$$\textbf{(F}*_\mathcal{T}\textbf{)} \quad [K *_\mathcal{T} \varphi] = min\Big(\, [\varphi \wedge \tau], \preceq_K \Big).$$

An important remark is in order. In many practical applications, the set of possible worlds $[\varphi \wedge \tau]$, appearing in condition $(\mathrm{F}*_\mathcal{T})$, is a *proper* subset of the set of possible worlds $[\varphi]$, appearing in condition $(\mathrm{F}*)$. This fact potentially reduces the resources required for determining the revised state of belief, since the search-space for specifying $\preceq_K$-minimal possible worlds *decreases.*

5. Relation to Multiple Belief Revision

In this section, we show that postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$ have a strong connection with Lindström's postulates for *multiple* belief revision [15]. We recall that multiple belief revision is the revision of a theory by a non-empty (possibly infinite) *set* of sentences [9]. This kind of revision is modelled as a binary function $\oplus$ mapping a theory K and a set of sentences Γ to a revised theory $K \oplus \Gamma$; i.e., $\oplus : \mathbb{K} \times 2^{\mathcal{L}} \mapsto \mathbb{K}$. We present the alluded Lindström's postulates as presented by Peppas in [19, Section 8.5.1] (Γ, Δ denote non-empty sets of sentences).[5]

[5]There are some differences between the definition of multiple belief revision presented by Peppas in [19] and that presented by Lindström in [15], which however are only superficial.

($\mathbf{K} \oplus \mathbf{1}$) $K \oplus \Gamma$ is a theory.

($\mathbf{K} \oplus \mathbf{2}$) $\Gamma \subseteq K \oplus \Gamma$.

($\mathbf{K} \oplus \mathbf{3}$) $K \oplus \Gamma \subseteq K + \Gamma$.

($\mathbf{K} \oplus \mathbf{4}$) If $K \cup \Gamma$ is consistent, then $K + \Gamma \subseteq K \oplus \Gamma$.

($\mathbf{K} \oplus \mathbf{5}$) $K \oplus \Gamma$ is inconsistent iff Γ is inconsistent.

($\mathbf{K} \oplus \mathbf{6}$) If $Cn(\Gamma) = Cn(\Delta)$, then $K \oplus \Gamma = K \oplus \Delta$.

($\mathbf{K} \oplus \mathbf{7}$) $K \oplus (\Gamma \cup \Delta) \subseteq (K \oplus \Gamma) + \Delta$.

($\mathbf{K} \oplus \mathbf{8}$) If $(K \oplus \Gamma) + \Delta$ is consistent, then $(K \oplus \Gamma) + \Delta \subseteq K \oplus (\Gamma \cup \Delta)$.

It follows from postulates ($K \oplus 1$)–($K \oplus 8$) that, whenever the set of sentences Γ is *finite*, the multiple revision of K by Γ is identical to the (classical) revision of K by the *conjunction* of all the members of Γ; in symbols, $K \oplus \Gamma = K * \big(\bigwedge \Gamma\big)$. It is not hard to verify that, in case Γ is considered to be the set of sentences $\{\varphi, \tau\}$, we end up with the next interesting equality, which involves all three revision operators $*_{\mathcal{T}}$, $\oplus$ and $*$:

$$K *_{\mathcal{T}} \varphi = K * (\varphi \wedge \tau) = K \oplus \{\varphi, \tau\}.$$

6. A Natural Class of AGM Revision Functions

This section is devoted to the identification of an interesting *proper* sub-class of the whole class of AGM revision functions, which, in the next section, will be related to the revision functions satisfying postulates ($K *_{\mathcal{T}} 1$)–($K *_{\mathcal{T}} 8$).

To that end, consider an agent whose belief corpus is represented by a theory K. Moreover, suppose that her/his revision-policy is encoded in an AGM revision function $*$, at which a faithful preorder $\preceq_K$ over $\mathbb{M}$ corresponds (via (F$*$)). Provided that the agent needs always to adhere to a background theory $\mathcal{T}$, one might reasonably postulate that all $\mathcal{T}$-worlds are *strictly more plausible* than all non-$\mathcal{T}$-worlds (with respect $\preceq_K$). This type of faithful preorders —which essentially represents a particular type of revision-policies— is identified by condition (TR), presented below, whereas Figure 3 depicts a rough layout of possible worlds whose comparative plausibility is constrained by (TR).

(TR) If $r \in [\mathcal{T}]$ and $r' \notin [\mathcal{T}]$, then $r \prec_K r'$.

Postulate (T), shown subsequently, constrains the behaviour of AGM revision functions, and is the *axiomatic* counterpart of condition (TR).

(T) If φ is consistent with $\mathcal{T}$, then $K * \varphi = K * (\varphi \wedge \tau)$.

The following representation theorem establishes the connection between (T) and (TR).

Theorem 3. *Let $\mathcal{T}$ be a background theory, such that $\mathcal{T} = Cn\big(\{\tau\}\big)$. Moreover, let $*$ be an AGM revision function, K be a theory, and $\preceq_K$ be the faithful preorder corresponding to $*$, by means of (F$*$). Then, $*$ satisfies (T) iff $\preceq_K$ satisfies (TR).*

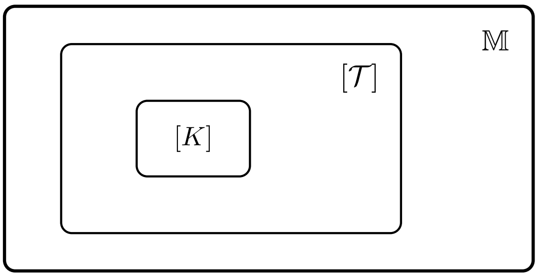

Figure 3. A rough layout of possible worlds whose comparative plausibility (encoded in a faithful preorder $\preceq_K$) is constrained by condition (TR); it is assumed that the comparative plausibility of worlds is inversely proportional to the Euclidean distance from their geometric center of the rectangle representing $[K]$.

Proof. For the left-to-right implication, assume that $*$ satisfies (T). We show that $\preceq_K$ satisfies (TR). Consider any two worlds $r, r' \in \mathbb{M}$, such that $r \in [\mathcal{T}]$ and $r' \notin [\mathcal{T}]$. Let φ be a sentence of $\mathcal{L}$, such that $[\varphi] = \{r, r'\}$. Since φ is consistent with $\mathcal{T}$, we derive, from condition (T), that $K * \varphi = K * (\varphi \wedge \tau)$; thus, all the worlds in $[K * \varphi]$ are φ-worlds that entail τ. Therefore, $r \in [K * \varphi]$ and $r' \notin [K * \varphi]$. Hence, from condition (F$*$), we derive that $r \prec_K r'$, as desired.

For the right-to-left implication, assume that $\preceq_K$ satisfies (TR). We show that $*$ satisfies (T). Let φ be a sentence of $\mathcal{L}$ which is consistent with $\mathcal{T}$. Since, from condition (TR), all $\mathcal{T}$-worlds are *strictly more plausible* than all non-$\mathcal{T}$-worlds (with respect to $\preceq_K$), and given that φ is consistent with $\mathcal{T}$, all the $\preceq_K$-minimal φ-worlds are also $\mathcal{T}$-worlds, or equivalently τ-worlds. Therefore, we derive, from condition (F$*$), that $min([\varphi], \preceq_K) = min([\varphi \wedge \tau], \preceq_K)$; hence, $K * \varphi = K * (\varphi \wedge \tau)$, as desired. □

We close this section with the following remark, which follows directly from condition (KM) of Section 4, and relates a revision function $*_\mathcal{T}$ (that respects postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$) to an AGM revision function $*$ that respects postulate (T), in case the epistemic input is *consistent* with the underlying background theory.

Remark 1. *Let $\mathcal{T}$ be a background theory. Moreover, let $*$ be an AGM revision function that satisfies postulate (T), $*_\mathcal{T}$ be the revision function induced from $*$, via condition (KM), K be a theory, and φ be a sentence of $\mathcal{L}$. If φ is consistent with $\mathcal{T}$, then $K *_\mathcal{T} \varphi = K * \varphi$.*

7. A Non-Prioritized Revision Operator

In this section, we introduce a special type of *non-prioritized* revision operator defined through a revision function $*_\mathcal{T}$ that satisfies postulates $(K *_\mathcal{T} 1)$–$(K *_\mathcal{T} 8)$. We recall that non-prioritized belief revision is a special type of belief change, according to which the new information is *not always* accepted in the revised state of belief, contrary to what postulate $(K * 2)$ enforces. Depending on a number of parameters, a non-prioritized revision operator may fully accept, partially accept, or even totally reject the new information; for a survey on non-prioritized belief revision, the interested reader is referred to [12].

The operator that we define herein is inspired by two well-known non-prioritized revision schemes, that is, *credibility-limited revision* of Hansson *et al.* [13], and *screened revision* of Makinson [16]. According to credibility-limited revision, the epistemic input is accepted in the revised state of belief only if it belongs to a given set of (a-priori) *credible* sentences. In screened revision, the intersection of a given set of sentences A with the agent's initial belief set K determines a set of *core* beliefs $A \cap K$; if the epistemic input is inconsistent with the set of core beliefs $A \cap K$, then it is rejected as implausible, otherwise, K is revised by φ so that none of the core beliefs are withdrawn [19].

On these premises, let $\mathcal{T}$ be a background theory, K be a theory, and φ be a sentence of $\mathcal{L}$. Then, the alluded non-prioritized revision operator, induced by a revision function $*_{\mathcal{T}}$ and denoted by $\otimes$, is defined as follows:

$$K \otimes \varphi = \begin{cases} K *_{\mathcal{T}} \varphi & \text{if } \varphi \text{ is consistent with } \mathcal{T} \\ \\ K & \text{otherwise} \end{cases}$$

Analogously to credibility-limited and screened revision, the revision operator $\otimes$, first checks whether the epistemic input φ should be accepted in the revised state of belief, and, if φ is acceptable (that is, φ is consistent with $\mathcal{T}$), theory K is $*_{\mathcal{T}}$-revised by φ; otherwise, the revised state of belief is identical to the initial state of belief.[6] Evidently, in view of the operator $\otimes$, φ is either fully accepted or totally rejected, in the revised state of belief.

In view of Remark 1, the above non-prioritized revision operator $\otimes$ can *equivalently* be defined as follows, by means of an AGM revision function $*$ which, in addition to $(K{*}1)$–$(K{*}8)$, satisfies postulate (T):

$$K \otimes \varphi = \begin{cases} K * \varphi & \text{if } \varphi \text{ is consistent with } \mathcal{T} \text{ and } * \text{ satisfies (T)} \\ \\ K & \text{otherwise.} \end{cases}$$

In this latter case, in which by definition the comparative plausibility of worlds corresponding to $*$ is constrained by condition (TR), the boundary between the $\mathcal{T}$-worlds and the non-$\mathcal{T}$-worlds can be regarded as a natural *threshold* of credibility of the new information φ (see Figure 3). Intuitively, the set of $\mathcal{T}$-worlds is the set of the admissible worlds; any world outside $[\mathcal{T}]$ is so implausible that it should *not* be accepted as a possible state of affairs. Consequently, any sentence φ that takes us to the "forbidden land" of the non-admissible worlds should be rejected; otherwise, it is business as usual, and the revised state of belief is $K{*}\varphi$, which is defined through condition (F$*$).

8. Conclusion

In this article, a generalized AGM paradigm that encodes rational belief revision, against background theories (knowledge), was provided. The proposed formal framework was shown to be highly relevant to a modification of the classical AGM paradigm for multiple belief revision. A natural class of AGM revision functions, along with an interesting non-prioritized revision operator, were also introduced; both these proposals were shown to have a smooth and reasonable connection with the aforementioned generalized revision-framework. An interesting avenue for future research is the study of belief contraction against background knowledge, in the context of which information is withdrawn from an initial belief corpus with the proviso that no part of the background knowledge is removed.

[6]Recall that, in case φ is inconsistent with $\mathcal{T}$, $K *_{\mathcal{T}} \varphi$ is inconsistent as well, due to postulate $(K *_{\mathcal{T}} 5)$.

References

[1] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[2] Theofanis Aravanis. On uniform belief revision. *Journal of Logic and Computation*, 30:1357–1376, 2020.

[3] Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams. A study of possible-worlds semantics of relevance-sensitive belief revision. *Journal of Logic and Computation*, 30:1125–1142, 2020.

[4] Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams. An investigation of parametrized difference revision operators. *Annals of Mathematics and Artificial Intelligence*, 89:7–28, 2021.

[5] Theofanis I. Aravanis, Pavlos Peppas, and Mary-Anne Williams. Full characterization of Parikh's relevance-sensitive axiom for belief revision. *Journal of Artificial Intelligence Research*, 66:765–792, 2019.

[6] Carlos Areces and Verónica Becher. Iterable AGM functions. In Mary-Anne Williams and Hans Rott, editors, *Frontiers in Belief Revision*, Lecture Notes in Computer Science, pages 261–277. Kluwer Academic Publishers, 2001.

[7] Mukesh Dalal. 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)*, pages 475–479. The AAAI Press, Menlo Park, California, 1988.

[8] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[9] André Fuhrmann and Sven O. Hansson. A survey of multiple contractions. *Journal of Logic, Language and Information*, 3:39–75, 1994.

[10] Peter Gärdenfors. *Knowledge in Flux – Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge, Massachusetts, 1988.

[11] Peter Gärdenfors. Belief revision and relevance. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pages 349–365, 1990.

[12] Sven O. Hansson. A survey of non-prioritized belief revision. *Erkenntnis*, 50(4):413–427, 1999.

[13] Sven O. Hansson, Eduardo L. Fermé, John Cantwell, and Marcelo A. Falappa. Credibility limited revision. *Journal of Symbolic Logic*, 66(4):1581–1596, 2001.

[14] Hirofumi Katsuno and Alberto Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

[15] Sten Lindström. A semantic approach to nonmonotonic reasoning: Inference operations and choice. *Uppsala Prints and Preprints in Philosophy, Department of Philosophy, University of Uppsala*, 6, 1991.

[16] David Makinson. Screened revision. *Theoria*, 63:14–23, 1997.

[17] David Makinson. Propositional relevance through letter-sharing. *Journal of Applied Logic*, 7:377–387, 2009.

[18] Rohit Parikh. Beliefs, belief revision, and splitting languages. In Lawrence S. Moss, Jonathan Ginzburg, and Maarten de Rijke, editors, *Logic, Language and Computation*, volume 2, pages 266–278. CSLI Publications, 1999.

[19] Pavlos Peppas. Belief revision. In Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, 2008.

[20] Pavlos Peppas. A panorama of iterated revision. In Sven O. Hansson, editor, *David Makinson on Classical Methods for Non-Classical Problems*, pages 71–94. Springer Netherlands, 2014.

[21] Pavlos Peppas and Mary-Anne Williams. Parametrised difference revision. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018)*, pages 277–286, 2018.

[22] Pavlos Peppas, Mary-Anne Williams, Samir Chopra, and Norman Foo. Relevance in belief revision. *Artificial Intelligence*, 229:126–138, 2015.

[23] Ken Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 455–462. Springer-Verlag, Tokyo, 1988.

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

Email address: taravanis@upatras.gr