

Belief Revision Operators with Varying Attitudes Towards Initial Beliefs

Adrian Haret*, Stefan Woltran

¹Institute of Logic and Computation, TU Wien, Austria

{haret,woltran}@dbai.tuwien.ac.at

Abstract

Classical axiomatizations of belief revision include a postulate stating that if new information is consistent with initial beliefs, then revision amounts to simply adding the new information to the original knowledge base. This postulate assumes a conservative attitude towards initial beliefs, in the sense that an agent faced with the need to revise them will seek to preserve initial beliefs as much as possible. In this work we look at operators that can assume different attitudes towards original beliefs. We provide axiomatizations of these operators by varying the aforementioned postulate and obtain representation results that characterize the new types of operators using preorders on possible worlds. We also present concrete examples for each new type of operator, adapting notions from decision theory.

1 Introduction

Belief revision models rational changes of an agent’s epistemic state triggered by the availability of new, trusted information. In the standard logical approach, an agent’s epistemic state is represented by propositional formulas, while the standards of rationality a revision operator is expected to abide by are encoded as logical axioms [Alchourrón *et al.*, 1985; Gärdenfors, 1988; Katsuno and Mendelzon, 1992; Fermé and Hansson, 2018]. Notably, the classical set of revision postulates turn out to define a class of operators that can be looked at in two ways: on the one hand as change, guided by logical postulates, of propositional theories in response to new data; and on the other hand as choice functions over possible worlds exploiting plausibility rankings over such interpretations. This correspondence tells us that an agent faced with revision of its initial beliefs acts as if it chooses from a set of feasible possible worlds the ones it considers most plausible.

A distinguishing feature of revision operators, as typically axiomatized, is that they can be assumed to adopt a particular attitude towards initial beliefs, enforced through what are called the *Inclusion* and *Vacuity* postulates in the AGM formulation [Fermé and Hansson, 2018], or through a single postulate equivalent to their conjunction in the KM ax-

iomatization [Katsuno and Mendelzon, 1992]. This attitude articulates the policy by which the agent’s prior information behaves with respect to new data: thus, in the KM axiomatization, the postulate in question states that if new information μ is consistent with existing beliefs κ , then the result of revision is simply $\kappa \wedge \mu$. In other words, the agent retains its initial beliefs and simply supplements them with the new item of information, if it can do so in a consistent way. This is in line with a view of revision where the information κ with which the agent starts off represents the possible worlds the agent finds most plausible, information not to be given up unless challenged by conflicting new data, and spells out a conservative attitude towards initial beliefs, guided by the desire to preserve them as much as possible.

In the current work we view such a conservative attitude as one among many that an agent can have towards its initial beliefs. By varying the postulate responsible for enforcing this attitude we are able to axiomatize revision operators that embody a wider range of attitudes towards prior information, and characterize these operators in terms of the types of preorders they induce on the set of possible worlds. To illustrate these principles we provide concrete operators, constructed using two ingredients: a notion of *distance* between possible worlds and an *aggregation function* that ranks possible worlds depending on the initial beliefs. We also show, in each case, how these operators fit into the landscape of new postulates introduced. Without the theoretical apparatus of the new postulates, the concrete operators put forward would be merely classified as deviant, since they do not satisfy the traditional blend of *Inclusion* and *Vacuity*. But through the present analysis they can be viewed as encoding distinct and characterizable stances an agent can take towards its beliefs.

Alternatives to the classical revision postulates have been considered since as far back as the original publications in the field [Gärdenfors, 1988; Katsuno and Mendelzon, 1991; Herzog and Rifi, 1999]. However, we believe that a systematic analysis of the intuition underlying the KM postulate corresponding to *Inclusion* and *Vacuity*, as we perform here, sheds new light on familiar topics.

2 Preliminaries

We assume a finite set \mathcal{P} of propositional atoms, from which the set Prop of propositional formulas is generated using the usual propositional connectives. A propositional knowledge

*Contact Author

base κ is a finite set of propositional formulas, which we typically identify with the conjunction of its formulas $\bigwedge_{\mu \in \kappa} \mu$. The set of all propositional knowledge bases is 2^{Prop} . The universe \mathcal{U} is the set of all possible interpretations (also called *possible worlds*) for formulas in Prop. The models of a propositional formula μ are the interpretations which satisfy it, and we write $[\mu]$ (respectively, $[\kappa]$) for the set of models of μ (respectively, for $\bigcap_{\mu \in \kappa} [\mu]$). If there is no danger of ambiguity, we write models as words where the letters are the atoms assigned to true, e.g., $\{\{a, b\}, \{b, c\}\}$ is written as $\{ab, bc\}$. For $\mu_1, \mu_2 \in \text{Prop}$, we say that μ_1 entails μ_2 , written $\mu_1 \models \mu_2$, if $[\mu_1] \subseteq [\mu_2]$, and that they are equivalent, written $\mu_1 \equiv \mu_2$, if $[\mu_1] = [\mu_2]$. A formula μ (a knowledge base κ) is consistent if $[\mu] \neq \emptyset$ ($[\kappa] \neq \emptyset$), and *complete* if it has exactly one model. The set of consistent knowledge bases is $2_{\text{cons}}^{\text{Prop}}$. If κ is a propositional knowledge base, then the *dual* $\bar{\kappa}$ of κ is obtained by replacing every literal appearing in κ with its negation. If w is an interpretation, the dual interpretation \bar{w} is $\mathcal{P} \setminus w$. If W is a set of interpretations, its dual \bar{W} is defined as $\{\bar{w} \mid w \in W\}$.

Example 1. If $\mathcal{P} = \{a, b, c\}$ and $\kappa = \{a, a \rightarrow b\}$, then $\bar{\kappa} = \{\neg a, \neg a \rightarrow \neg b\}$. We have that $[\kappa] = \{ab, abc\}$, the dual of the interpretation ab is $\bar{ab} = c$ and $[\bar{\kappa}] = \{c, \emptyset\}$.

In Example 1 we obtain that $[\bar{\kappa}] = [\bar{[\kappa]}]$. Though we do not provide a formal proof, we mention here that this holds more generally, i.e., for any $\kappa \in 2^{\text{Prop}}$, it holds that $[\bar{\kappa}] = [\bar{[\kappa]}]$.

If \mathcal{M} is a set, then $\text{Bin}(\mathcal{M})$ is the set of binary relations on \mathcal{M} . We write $<$ for the strict part of \leq , i.e., $x < x'$ if $x \leq x'$ and $x' \not\leq x$; moreover, $x \approx x'$ if $x \leq x'$ and $x' \leq x$. The \leq -minimal elements of \mathcal{M} with respect to \leq are defined as $\min_{\leq} \mathcal{M} = \{x \in \mathcal{M} \mid \nexists x' \in \mathcal{M} \text{ such that } x' < x\}$. An *assignment* from \mathcal{M}_1 to \mathcal{M}_2 is a function $\alpha: \mathcal{M}_1 \rightarrow \text{Bin}(\mathcal{M}_2)$. We write \leq_{κ} instead of $\alpha(\kappa)$ if there is no danger of ambiguity. If W is a set of interpretations, we denote by φ_W a propositional formula such that $[\varphi_W] = W$.

3 Revision: Axioms and Characterizations

A propositional revision operator is a function $\circ: 2_{\text{cons}}^{\text{Prop}} \times \text{Prop} \rightarrow \text{Prop}$. The intention is that $\kappa \circ \mu$ encodes changes brought to existing held beliefs κ such that new, trusted information μ is accepted. A sensible revision operator is expected to resolve inconsistencies between κ and μ and to satisfy certain rationality criteria, presented below.

3.1 Basic Postulates

If $\kappa, \kappa_1, \kappa_2 \in 2_{\text{cons}}^{\text{Prop}}$ and $\mu, \mu_1, \mu_2 \in \text{Prop}$, we first single out a core set of axioms. Note that R₂ as we write it does not coincide with the second KM postulate. The second KM postulate shows up in the latter part of this section.

- (R₁) $\kappa \circ \mu \models \mu$.
- (R₂) If μ is consistent, then $\kappa \circ \mu$ is consistent.
- (R₃) If $\kappa_1 \equiv \kappa_2$ and $\mu_1 \equiv \mu_2$, then $\kappa_1 \circ \mu_1 \equiv \kappa_2 \circ \mu_2$.
- (R₄) $(\kappa \circ \mu_1) \wedge \mu_2 \models \kappa \circ (\mu_1 \wedge \mu_2)$.
- (R₅) If $(\kappa \circ \mu_1) \wedge \mu_2$ is cons., then $\kappa \circ (\mu_1 \wedge \mu_2) \models (\kappa \circ \mu_1) \wedge \mu_2$.

Postulates R_{1–5} are part of the standard set of KM postulates [Katsuno and Mendelzon, 1992], saying that a revision operator incorporates new information μ (R₁), returns a consistent output if μ is consistent (R₂), performs its task irrespectively of how beliefs are written down (R₃) and satisfies some coherence constraints when the revision formula is varied (R_{4–5}). A revision operator is *basic* if it satisfies R_{1–5}. Throughout the paper, we work only with basic operators.

3.2 Basic Assignments

Reflection on postulates R_{1–5} reveals that an operator \circ satisfying them chooses among models of μ and, in doing so, behaves as if it has preferences over units of information. Formally, this is cashed out by assigning to each consistent knowledge base κ in $2_{\text{cons}}^{\text{Prop}}$ a binary relation \leq_{κ} on interpretations in \mathcal{U} : to revise κ by μ , then, becomes equivalent to choosing the best models of μ in \leq_{κ} . And, in the same way that revision operators are expected to satisfy a set of basic properties (postulates R_{1–5}), rankings on \mathcal{U} must satisfy a set of properties, to be introduced in the following, that are conducive to rational choice.

For an assignment $\alpha: 2_{\text{cons}}^{\text{Prop}} \rightarrow \text{Bin}(\mathcal{U})$, a revision operator \circ and $\kappa, \kappa_1, \kappa_2 \in 2_{\text{cons}}^{\text{Prop}}$, $w_1, w_2 \in \mathcal{U}$, we first consider the following properties:

- (O₁) \leq_{κ} is reflexive.
- (O₂) \leq_{κ} is transitive.
- (O₃) If $\kappa_1 \equiv \kappa_2$, then $\leq_{\kappa_1} = \leq_{\kappa_2}$.
- (O₄) \leq_{κ} is total.
- (O₅) $[\kappa \circ \mu] = \min_{\leq_{\kappa}} [\mu]$, for any propositional formula μ .

An assignment is *basic* if it satisfies properties o_{1–4}. Notice that properties o_{1–2} imply that \leq_{κ} is a preorder on \mathcal{U} . Adding property o₄ makes \leq_{κ} total, and o₃ adds an independence of syntax aspect to the assignment. If, on top of this, \leq_{κ} satisfies o₅, we say that the assignment α *represents the revision operator* \circ (and that \circ is *represented* by α). The overloading of the term ‘basic’ is intentional: as Theorem 1 shows, properties o_{1–5} define a class of rankings on interpretations that fully characterize revision operators satisfying axioms R_{1–5}.

Theorem 1. A revision operator satisfies postulates R_{1–5} iff there exists an assignment $\alpha: 2_{\text{cons}}^{\text{Prop}} \rightarrow \mathcal{U}$ representing it such that, for any $\kappa \in 2_{\text{cons}}^{\text{Prop}}$, \leq_{κ} satisfies properties o_{1–5}.

Proof. (“ \Rightarrow ”) The proof amounts to checking, one by one, that axioms R_{1–5} are satisfied by operator \circ . The argument is not new, and can be extracted from classical works on the subject [Katsuno and Mendelzon, 1992].

(“ \Leftarrow ”) If w_1, \dots, w_n are interpretations, we write $\varphi_{1, \dots, n}$ for a propositional formula such that $[\varphi_{1, \dots, n}] = \{w_1, \dots, w_n\}$.

Assume, now, that we are given a revision operator \circ that satisfies axioms R_{1–5}. For any base κ and $w_1, w_2 \in \mathcal{U}$, we define the assignment $\alpha: 2_{\text{cons}}^{\text{Prop}} \rightarrow \mathcal{U}$ by saying that $w_1 \leq_{\kappa} w_2$ if $w_1 \in [\kappa \circ \varphi_{1,2}]$. We now show that \leq_{κ} , thus defined, satisfies properties o_{1–5}.

Property o₁ (reflexivity) follows using axioms R₁ and R₃. For property o₂ (transitivity), take w_1, w_2 and w_3 such that $w_1 \leq_{\kappa} w_2$ and $w_2 \leq_{\kappa} w_3$. First we show that $w_1 \in [\kappa \circ \varphi_{1,2,3}]$. Notice that, by axioms R₁ and R₂, we have that $\emptyset \subset [\kappa \circ$

$\varphi_{1,2,3}] \subseteq [\varphi_{1,2,3}]$. We now do a case analysis. *Case 1.* If $w_1 \in [\kappa \circ \varphi_{1,2,3}]$, the conclusion is immediate. *Case 2.* If $w_2 \in [\kappa \circ \varphi_{1,2,3}]$, then $(\kappa \circ \varphi_{1,2,3}) \wedge \varphi_{1,2}$ is consistent. Using axioms R_{4-5} and R_3 , this implies that $(\kappa \circ \varphi_{1,2,3}) \wedge \varphi_{1,2} \equiv \kappa \circ (\varphi_{1,2,3} \wedge \varphi_{1,2}) \equiv \kappa \circ \varphi_{1,2}$. Since $w_1 \in [\kappa \circ \varphi_{1,2}]$ (because $w_1 \leq_{\kappa} w_2$), we have that $w_1 \in [\kappa \circ \varphi_{1,2,3}]$. *Case 3.* If $w_3 \in [\kappa \circ \varphi_{1,2,3}]$, then $(\kappa \circ \varphi_{1,2,3}) \wedge \varphi_{2,3}$ is consistent. Using axioms R_{4-5} and R_3 again, we get that $w_2 \in [\kappa \circ \varphi_{1,2,3}]$. Reproducing the reasoning from Case 2, we conclude that $w_1 \in [\kappa \circ \varphi_{1,2,3}]$. Once we have that $w_1 \in [\kappa \circ \varphi_{1,2,3}]$, it follows that $(\kappa \circ \varphi_{1,2,3}) \wedge \varphi_{1,3}$ is consistent. So, by axioms R_{4-5} and R_3 , we have that: $(\kappa \circ \varphi_{1,2,3}) \wedge \varphi_{1,3} \equiv \kappa \circ (\varphi_{1,2,3} \wedge \varphi_{1,3}) \equiv \kappa \circ \varphi_{1,3}$. Since $w_1 \in [\kappa \circ \varphi_{1,2,3}]$ and $w_1 \in [\varphi_{1,3}]$, we conclude that $w_1 \in [\kappa \circ \varphi_{1,3}]$. This implies that $w_1 \leq_{\kappa} w_3$. Property \circ_3 follows from the fact that the definition of \leq_{κ} does not depend in any way on the syntax of κ . Property \circ_4 follows by axioms R_{1-2} . For property \circ_5 , we show the double inclusion. (“ \subseteq ”) Take $w_1 \in [\kappa \circ \mu]$, and $w_2 \in [\mu]$. We get that $w_1 \in [(\kappa \circ \mu) \wedge \varphi_{1,2}]$ and, by axioms R_4 and R_3 , it follows that $w_1 \in [\kappa \circ (\mu \wedge \varphi_{1,2})]$ and then that $w_1 \in [\kappa \circ \varphi_{1,2}]$. Thus, $w_1 \leq_{\kappa} w_2$ and hence, $w_1 \in \min_{\leq_{\kappa}} [\mu]$. (“ \supseteq ”) Take $w_1 \in \min_{\leq_{\kappa}} [\mu]$ and suppose that $w_1 \notin [\kappa \circ \mu]$. Since it follows that μ is consistent, we have, by axioms R_{1-2} , that there exists $w_2 \in [\kappa \circ \mu]$. By axiom R_5 , we have that $(\kappa \circ \mu) \wedge \varphi_{1,2} \models \kappa \circ (\mu \wedge \varphi_{1,2})$. By axiom R_3 , we have that $\kappa \circ (\mu \wedge \varphi_{1,2}) \equiv \kappa \circ \varphi_{1,2}$. Since $w_1 \notin [\kappa \circ \mu]$, it follows that $w_1 \notin [\kappa \circ \varphi_{1,2}]$, and hence $w_2 <_{\kappa} w_1$. But we also have that $w_1 \in \min_{\leq_{\kappa}} [\mu]$, implying that $w_1 \leq_{\kappa} w_2$, which is a contradiction. \square

Theorem 1 tells us that an agent revising beliefs along the lines of postulates R_{1-5} behaves as if it ranks interpretations in \mathcal{U} in a total preorder \leq_{κ} that depends on initial beliefs κ , and always picks the minimal models of μ according to \leq_{κ} . Such an agent, then, behaves like a rational agent, in the sense of rational choice theory [Sen, 1984; French, 1988], choosing the best elements from a given menu of options: the menu, here, would be the models of μ , i.e., the possible worlds the agent is allowed to believe in light of new information, while the best elements are decided with reference to \leq_{κ} . Thus, a revision operator can be seen as a *choice function* over sets of interpretations: accordingly, Theorem 1 aligns with standard choice theoretic results [Arrow, 1959; Sen, 1984; Moulin, 1991]. That a similar mathematical formalism underlies both belief revision and rational choice is, by itself, not a new insight, the topic having been studied under various guises [Rott, 2001; Bonanno, 2009]. The main difference is one of interpretation, with \leq_{κ} usually thought of in revision as a plausibility ranking, i.e., the agent’s assessment as to which possible worlds are more or less plausible.

Example 2. A doctor knows that the patient has been diagnosed with asthma (a), finds out that the patient is suffering from shortness of breath (b) and infers that chest pain (c) is also present: the two often go together in asthma, and such stereotypical examples are more salient in the doctor’s mind than less frequent, more exotic examples. In other words, the doctor has initial information $\kappa = a$ and a plausibility ranking \leq_{κ} over possible worlds (with respect to the alphabet $\mathcal{P} = \{a, b, c\}$) depicted in Figure 1. The doctor then revises by $\mu = b$ and settles on a possible state of affairs that




Figure 1: Revision scenario of Example 2, showing the relevant section of the preorder \leq_{κ} on the basis of which the revision result is constructed: only models of $\mu = b$ are depicted; the models of κ among this set are highlighted in grey.

is most plausible according to their plausibility ranking \leq_{κ} , i.e., $[\kappa \circ \mu] = \min_{\leq_{\kappa}} [\mu] = \{abc\}$. Note that the doctor, in this case, takes the situation represented by $a \wedge b \wedge c$ to be more likely than $a \wedge b \wedge \neg c$.

One way of thinking of postulates R_{1-5} is that they axiomatize total preorders on interpretations. These preorders nominally depend on κ , but nothing in postulates R_{1-5} touches on how models of κ should influence these preorders. In other words, there is as yet no information about the attitude of an agent towards its initial epistemic state, and postulates R_{1-5} are consistent with arbitrary attitudes towards κ . How should the models of κ stand in relation to all other interpretations? Example 2 offers a glimpse into one possible answer: the agent starts off with some information κ and differentiates among possible worlds consistent with κ : some of these worlds are more plausible than others, perhaps as a result of being more salient, or because of a systematic bias [Kahneman *et al.*, 1982]. Still, as a whole, models of κ are more plausible than any *other* interpretations consistent with the new information μ . In other words, the agent is biased towards the possible worlds consistent with κ . Are there, now, other ways of arranging the models of κ in \leq_{κ} , ways that span the space of possible such attitudes? We study this question through the lens of additional axioms.

3.3 Attitudes Towards Initial Beliefs

If $\kappa, \kappa' \in 2^{\text{Prop}}_{\text{cons}}$, $\mu \in \text{Prop}$, consider the following postulates:

- (R₆) If $\kappa \wedge \mu$ is consistent, then $\kappa \wedge \mu \models \kappa \circ \mu$.
- (R₇) If $\kappa \wedge \mu$ is consistent, then $\kappa \circ \mu \models \kappa \wedge \mu$.
- (R₈) If $\kappa \circ \mu \models \bar{\kappa}$, then $\kappa \circ \mu \equiv \mu$.
- (R₉) If $\mu \not\models \bar{\kappa}$, then $(\kappa \circ \mu) \wedge \bar{\kappa}$ is inconsistent.
- (R₁₀) If κ' is complete and $\kappa' \models \kappa \circ \mu$, then $\kappa' \models (\kappa \vee \kappa') \circ \mu$.

Each of these postulates encodes a particular type of attitude towards initial beliefs, and they are intended to be thought of in conjunction with the basic set of postulates R_{1-5} . Some clarification is in order. Postulate R_6 models an agent that incorporates all information in $\kappa \wedge \mu$, and possibly extends this to cover more ground. Postulate R_7 models an agent that reserves the right to drop information from κ if it so sees fit, even if that information is consistent with μ : we may imagine this is done on the basis of certain preferences over the information encoded by κ , i.e., the agent is partial towards certain parts of κ to the detriment of others. Taken together, postulates R_{6-7} imply that $\kappa \circ \mu$ is equivalent to $\kappa \wedge \mu$, when $\kappa \wedge \mu$ is consistent. This property models an agent who wants

to preserve as much of κ as it can, and does not have any bias towards either of the models of κ . Postulate R_6 can be equated with the *Inclusion* postulate in the AGM formulation and R_7 corresponds to *Vacuity* [Fermé and Hansson, 2018], while in the KM axiomatization R_6 and R_7 are packaged together in one postulate (i.e., KM postulate R_2) and presented alongside R_{1-5} as the default set of rational properties for revision [Katsuno and Mendelzon, 1992].

Postulates R_{8-9} focus on the dual knowledge base $\bar{\kappa}$ obtained by replacing every literal in κ with its negated version. If κ is a conjunction of literals, or if it is a complete (i.e., with exactly one model) formula, then $\bar{\kappa}$ will be a formula whose models are complements of the models of κ .

Example 3. If $\mathcal{P} = \{a, b, c\}$ and $\kappa = \{a \wedge b\}$ is a knowledge base over the alphabet \mathcal{P} , then $\bar{\kappa} = \{\neg a \wedge \neg b\}$, and we have that $[\kappa] = \{ab, abc\}$, while $[\bar{\kappa}] = \{\emptyset, c\}$

Thus, if κ is very specific (e.g., is a conjunction of literals), then $\bar{\kappa}$ can be thought of as the point of view opposite to that of κ , and situations can be imagined in which it is desirable to put bounds on the revision function in terms of how it treats information encoded by $\bar{\kappa}$. This is the case if the agent has, or is required to have, a definite opinion on every item from an agenda, as is typically the case in Judgment Aggregation [Endriss, 2016]; if κ is a ‘vivid’ knowledge base [Levesque, 1986]; or, if it encodes something like an agent’s preferred bundle from a set of available items. In all these cases κ can be required to be a conjunction of literals or a complete formula.

Postulate R_8 says that if κ undergoes revision by a formula μ embodying such an adverse perspective, then the agent must adopt μ : in other words, the agent has no room for maneuvering towards a more amenable middle ground. Such a revision policy makes more sense when considered alongside postulate R_9 , which specifies that if the agent has the option of believing states of affairs *not* compatible with $\bar{\kappa}$, it should wholeheartedly adopt those as the most plausible stance. Taken together, postulates R_{8-9} inform the agent to believe states of affairs compatible with $\bar{\kappa}$ only if it has no other choice in the matter: the models of $\bar{\kappa}$ should be part of a viewpoint one is willing to accept only as a last resort. Postulate R_{10} is best understood through an example.

Example 4. An agent intends to go to an art museum, the beach and a concert, i.e., $\kappa = \{a \wedge b \wedge c\}$, with $\mathcal{P} = \{a, b, c\}$. The agent then learns that it only has time for one of these activities and chooses the art museum, i.e., $\kappa \circ \mu \equiv a \wedge \neg b \wedge \neg c$. If the agent’s initial intentions were less specific, for instance that it would either go to all three places or only to the art museum (i.e., $\kappa = \{(a \wedge b \wedge c) \vee (a \wedge \neg b \wedge \neg c)\}$), then, faced with the same new information μ , $a \wedge \neg b \wedge \neg c$ should still feature as one of its most preferred options.

A clearer view of postulates R_{6-10} emerges when looking at how they situate the models of κ in a total preorder \leq_κ , for $\kappa, \kappa' \in 2^{\text{Prop}_{\text{cons}}}$, $w_1, w_2, w' \in \mathcal{U}$:

- (o₆) If $w_1 \in [\kappa]$, then $w_1 \leq_\kappa w_2$.
- (o₇) If $w_1 \in [\kappa]$ and $w_2 \notin [\kappa]$, then $w_1 <_\kappa w_2$.
- (o₈) If $w_1 \in [\bar{\kappa}]$, then $w_2 \leq_\kappa w_1$.

(o₉) If $w_1 \in [\bar{\kappa}]$ and $w_2 \notin [\bar{\kappa}]$, then $w_2 <_\kappa w_1$.

(o₁₀) If $w' \leq_\kappa w$ and $[\kappa'] = \{w'\}$, then $w' \leq_{\kappa \vee \kappa'} w$.

Properties o₆₋₁₀ turn out to characterize axioms R_{6-10} on the semantic level, as per the following representation result.

Theorem 2. If \circ is a basic revision operator and $\alpha: 2^{\text{Prop}_{\text{cons}}} \rightarrow \mathcal{U}$ is a basic assignment that represents it, then, for any $\kappa \in 2^{\text{Prop}_{\text{cons}}}$ and $\mu \in \text{Prop}$, the following equivalences hold:

- (1) \circ satisfies axiom R_6 iff \leq_κ satisfies property o₆;
- (2) \circ satisfies axiom R_7 iff \leq_κ satisfies property o₇;
- (3) \circ satisfies axiom R_8 iff \leq_κ satisfies property o₈;
- (4) \circ satisfies axiom R_9 iff \leq_κ satisfies property o₉;
- (5) \circ satisfies axiom R_{10} iff \leq_κ satisfies property o₁₀.

Proof. Recall that we denote by $\varphi_{1,2}$ a propositional formula such that $[\varphi_{1,2}] = \{w_1, w_2\}$. For equivalence 1, we show each direction in turn. (“ \Rightarrow ”) Take an assignment α satisfying property o₆, and the revision operator \circ represented by it. We assume that $\kappa \wedge \mu$ is consistent, and show that for any $w \in [\kappa \wedge \mu]$, it holds that $w \in [\kappa \circ \mu]$ as well. By property o₅, this is equivalent to showing that $w \in \min_{\leq_\kappa} [\mu]$. Take $w' \in [\mu]$. Since $w \in [\kappa]$, we can apply property o₆ to get that $w \leq_\kappa w'$. Hence $w \in \min_{\leq_\kappa} [\mu]$. (“ \Leftarrow ”) Take a basic revision operator \circ satisfying R_6 , and the assignment α which represents it. To show that \leq_κ satisfies property o₆, take two interpretations w_1 and w_2 such that $w_1 \in [\kappa]$. Then, by axiom R_6 , we have that $\kappa \wedge \varphi_{1,2} \models \kappa \circ \varphi_{1,2}$. By property o₅, it holds that $[\kappa \circ \varphi_{1,2}] = \min_{\leq_\kappa} [\varphi_{1,2}]$ and, since $w_1 \in [\kappa \wedge \varphi_{1,2}]$, it follows that $w_1 \in \min_{\leq_\kappa} [\varphi_{1,2}]$. Thus, $w_1 \leq_\kappa w_2$.

Equivalence 2 uses similar reasoning, and equivalences 3 and 4 are analogous to 1 and 2, respectively. For equivalence 5, assume first that axiom R_{10} holds, and take interpretations w and w' and a knowledge base κ' such that $w' \leq_\kappa w$ and $[\kappa'] = \{w'\}$. To show that $w' \leq_{\kappa \vee \kappa'} w$, we must show that $w' \in [(\kappa \vee \kappa') \circ \varphi_{w,w'}]$, where $\varphi_{w,w'}$ is a formula such that $[\varphi_{w,w'}] = \{w, w'\}$. This follows immediately by applying axiom R_{10} . Conversely, suppose $[\kappa'] = \{w'\}$, and take $w \in [\kappa \circ \mu]$. Then, we get that $w' \leq_\kappa w$, and we can apply property o₁₀ to derive the conclusion. \square

Theorem 2 is better understood through an illustration of how such preorders treat models of κ . Property o₆ says that models of κ are minimal elements in κ , i.e., the agent considers possible worlds satisfying its beliefs among the most plausible possible worlds, though possibly not uniquely so (Figure 2-(a)). Property o₇ states that there are no counter-models of κ more plausible than the models of κ , but the models of κ themselves may not be equally plausible (Figure 2-(b)). Properties o₈₋₉ say that models of the dual knowledge base κ are the least plausible interpretations in \leq_κ (Figure 2-(c,d)), while property o₁₀ says that if w' is more plausible than w when the initial beliefs are κ , then w' would still be more plausible than w if it were part of the initial beliefs (Figure 2-(e)).

Together, properties o₁₋₇ define what is more commonly known as a *faithful assignment*, placing all and only models of κ on the lowest level of \leq_κ . This corresponds to an agent that holds its initial beliefs to be the most plausible states of affairs [Katsuno and Mendelzon, 1992]. Consequently,




Figure 2: Schematic view of prototypical preorders satisfying each of the properties o_{6-10} ; models of κ are in the light gray area, models of $\bar{\kappa}$ are in the dark gray area.

Theorem 1 plus equivalences 1-2 from Theorem 2 make up the classical representation result for belief revision operators [Katsuno and Mendelzon, 1992]. Here we have opted for a more fine-grained approach to the placement of models of κ in \leq_{κ} , which allows a more diverse representation of the different types of attitudes an agent can have towards initial beliefs. Though operators that do not satisfy the classical KM postulate R_2 have been considered before [Ryan, 1996; Benferhat *et al.*, 2005], the idea that such deviations correspond to possible epistemic attitudes and can be axiomatized is, to the best of our knowledge, new.

3.4 Indifference to Already Held Beliefs

One particular consequence of weakening the KM axiom R_2 (axioms R_{6-7} in the current context) is that the following property is not guaranteed to hold anymore:

$$(R_{IDF}) \quad \kappa \circ \kappa \equiv \kappa.$$

This property, called here R_{IDF} (for *indifference to already held beliefs*), says that revising with information the agent already believes does not change the agent’s epistemic state. More generally, the KM standard set of postulates implies that revising by any formula μ such that $\kappa \models \mu$ results in κ . It quickly becomes apparent that axiom R_6 implies R_{IDF} ¹ but R_7 does not. Thus, if an agent is allowed to rank models of κ unequally, then R_{IDF} is not guaranteed to hold.

Example 5. If $\mathcal{P} = \{a, b\}$, take $\kappa = \{a \vee b\}$ and a revision operator that satisfies axiom R_7 , and which orders interpretations as follows: $a \approx_{\kappa} b <_{\kappa} ab <_{\kappa} \emptyset$. We get that $[\kappa \circ \kappa] = \{a, b\}$, i.e., $\kappa \circ \kappa \equiv (a \leftrightarrow \neg b) \not\equiv \kappa$.

This points to a more graded view of what it means to believe κ . Thus, an agent might have a certain threshold of plausibility, along the lines of what is known in epistemology as *the Lockean thesis* [Foley, 1993], according to which it calibrates its beliefs: anything above the threshold counts as part of the belief κ and anything below counts as disbelief. This fits with the idea that an agent might assign different degrees of plausibility to states of affairs consistent with its belief κ : indeed, this is the point of view we endorse here, in contrast to more standard approaches, which consider that an agent assigns equal degrees of plausibility to all items of its belief. Thus, incoming information that confirms an agent’s belief

¹The converse is not true: R_{IDF} enforces only that models of κ are equally plausible, but not where they are placed in \leq_{κ} .

might have the effect of *reinforcing* parts that are given more plausibility at the expense of parts that are given less, and this is the kind of situation we take to be modeled by Example 5.

What would be worrying would be a revision policy that makes an agent cycle between different viewpoints when confronted repeatedly with the same type of information: we will see that for revision operators satisfying R_7 this concern is unwarranted, but we must first introduce some new notation. We write κ^i for the knowledge base obtained by revising κ by itself an i number of times. Thus, $\kappa^0 = \kappa$ and $\kappa^{i+1} = \kappa^i \circ \kappa$. Consider now the following property:

(R_{STB}) There is $n \geq 1$ such that $\kappa^m \equiv \kappa^n$, for every $m \geq n$.

A revision operator \circ is *stable* if it satisfies property R_{STB} . Stability implies that repeated revision by κ ultimately settles (or stabilizes) on a set of models that does not change through subsequent revisions by κ . The following result proves relevant to the issue of stability.

Proposition 3. If a revision operator \circ satisfies axioms R_1 and R_7 , then $\kappa^{i+1} \models \kappa^i$.

Proof. By axiom R_1 , we have that $\kappa \circ \kappa \models \kappa$, and thus $\kappa^1 \models \kappa^0$. Applying axiom R_7 , we have that $(\kappa \circ \kappa) \circ \kappa \models (\kappa \circ \kappa) \wedge \kappa \models \kappa \circ \kappa$. Thus, $\kappa^2 \models \kappa^1$, and it is straightforward to see how this argument is iterated to get the conclusion. \square


If the operator \circ also satisfies axiom R_2 (which, here, says that if the revision formula is consistent, then the revision result is also consistent), it follows that if κ is consistent, then κ_i is consistent, for any $i \geq 0$. Thus, combining this fact and Proposition 3, we get that repeated revision by κ leads to a chain of ever more specific knowledge bases, i.e., $\emptyset \subset \dots \subseteq [\kappa^{i+1}] \subseteq [\kappa^i] \subseteq \dots \subseteq [\kappa^0]$. Since a knowledge base has a finite number of models, it falls out immediately from this that there must be a point at which further revision by κ does not change anything.

Corollary 4. A basic revision operator \circ satisfying axiom R_7 is stable.

Unfortunately, axioms R_{8-9} do not guarantee stability. Since these axioms require only that the agent places the models of $\bar{\kappa}$ as the least plausible interpretations, it becomes possible that an agent’s plausibility ranking does not hold on to a core set of interpretations through successive revisions by κ .

Example 6. If $\mathcal{P} = \{a, b\}$, take $\kappa = \{\neg b\}$ and a revision operator satisfying R_{8-9} , which orders interpretations as shown in Figure 3. We have that $[\kappa^0] = [\kappa] = \{\emptyset, a\}$, $[\kappa^1] = [\kappa \circ \kappa] = \{a\}$, and $[\kappa^2] = [\kappa^1 \circ \kappa] = \{\emptyset\}$. By R_3 , subsequent revisions by κ cycle between $\{a\}$ and $\{\emptyset\}$ ($[\kappa^3] = \{a\}$, $[\kappa^4] = \{\emptyset\}$), thus never settling on a stable answer.

The issue of stability suggests another dimension along which revision operators can be analyzed, with Corollary 4 and Example 6 showing that a revision operator does not satisfy it trivially. Example 6, in particular, shows that there is interplay between \leq_{κ} and $\leq_{\kappa'}$, if $\kappa' \models \kappa$, which is relevant to the question of whether an operator is stable. This interplay is reminiscent of topics like iterated revision and kinetic consistency [Darwiche and Pearl, 1997; Peppas and Williams, 2016], but pursuing it further would take us too far afield of the aims of the current work.


 Figure 3: Repeated revision by κ cycles between $\{a\}$ and $\{\emptyset\}$.

4 Concrete Operators with Varying Attitudes

We now ask: what is a natural way to construct operators with such varying attitudes towards initial information? Our answer builds on ideas found in belief merging [Konieczny *et al.*, 2004; Konieczny and Pérez, 2011], i.e., it defines rankings on interpretations by appeal to *two* main ingredients. The first is a *distance* $d: \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ between interpretations, defined such that $d(w_1, w_2) = 0$ iff $w_1 = w_2$ and $d(w_1, w_2) = d(w_2, w_1)$. Given $w \in \mathcal{U}$ and $\kappa \in 2^{\mathcal{P}_{\text{cons}}^{\text{prop}}}$ such that $[\kappa] = \{w_1, \dots, w_n\}$, the *vector of distances from w to κ* is $d(w, \kappa) = (d(w, w_1), \dots, d(w, w_n))$. For brevity we will omit commas and simply write $d(w, \kappa)$ as a string of numbers. We recall two prominent examples of distance: *drastic distance* d_D works by the all-or-nothing rule: $d_D(w_1, w_2) = 0$ if $w_1 = w_2$, and 1 otherwise; *Hamming distance* d_H , counts the number of atoms on which two interpretations differ. The second ingredient is an *aggregation function* (more precisely, a family of functions) $f: \mathbb{R}^n \rightarrow \mathbb{R}$, for $n \in \mathbb{N}$, mapping a distance vector $d(w, \kappa)$ to a number, and used to compare distance vectors. We write $\overrightarrow{d(w, \kappa)}$ and $\overleftarrow{d(w, \kappa)}$ for the vectors of distances from w to κ ordered in ascending order and descending order, respectively. The lexicographic order between two vectors is denoted by \leq_{lex} . The minimal and maximal elements of $d(w, \kappa)$ are $\min d(w, \kappa)$ and $\max d(w, \kappa)$, respectively, and $\sum d(w, \kappa) = \sum_{w_i \in [\kappa]} d(w, w_i)$. The *centrality of w with respect to κ* is $\text{cen}(w, \kappa) = \max d(w, \kappa) - \min d(w, \kappa)$. The *displacement of w with respect to κ* is $\text{dis}(w, \kappa) = \min d(w, \kappa) - \min d(w^*, \kappa)$, where w^* is an interpretation such that $\min d(w^*, \kappa)$ is minimal among all the interpretations w' for which $\text{cen}(w', \kappa) = \text{cen}(w, \kappa)$. Finally, the *agreeability index of w with respect to κ* is $\text{agr}(w, \kappa) = \min\{\min d(w, \kappa), \text{cen}(w, \kappa) + \text{dis}(w, \kappa)\}$, while the *disagreeability index of w with respect to κ* is $\text{dagr}(w, \kappa) = n - \text{agr}(w, \bar{\kappa})$, where $n = |\mathcal{P}|$.

Given a distance d between interpretations and aggregation function f , we write $\leq_{\kappa}^{d, f}$ for the ranking generated using d and f , and $\circ^{d, f}$ for the revision operator represented by the assignment generated using d and f , i.e., $[\kappa \circ^{d, f} \mu] = \min_{\leq_{\kappa}^{d, f}} [\mu]$. We then define the following types of rankings:

$$\begin{array}{ll}
 w_1 \leq_{\kappa}^{d, \min} w_2 & \text{if } \min d(w_1, \kappa) \leq \min d(w_2, \kappa), \\
 w_1 \leq_{\kappa}^{d, \text{lmin}} w_2 & \text{if } d(w_1, \kappa) \leq_{\text{lex}} d(w_2, \kappa), \\
 w_1 \leq_{\kappa}^{d, \text{agr}} w_2 & \text{if } \text{agr}(w_1, \kappa) \leq \text{agr}(w_2, \kappa), \\
 w_1 \leq_{\kappa}^{d, \text{max}} w_2 & \text{if } \max d(w_1, \kappa) \leq \max d(w_2, \kappa), \\
 w_1 \leq_{\kappa}^{d, \text{lmax}} w_2 & \text{if } \overleftarrow{d}(w_1, \kappa) \leq_{\text{lex}} \overleftarrow{d}(w_2, \kappa), \\
 w_1 \leq_{\kappa}^{d, \text{dagr}} w_2 & \text{if } \text{dagr}(w_1, \kappa) \leq \text{dagr}(w_2, \kappa), \\
 w_1 \leq_{\kappa}^{d, \text{sum}} w_2 & \text{if } \sum d(w_1, \kappa) \leq \sum d(w_2, \kappa).
 \end{array}$$

	\emptyset	a	b	abc	$\overrightarrow{d_H(w, \kappa)}$	$\overleftarrow{d_H(w, \kappa)}$	min	max	\sum
\emptyset	0	1	1	3	(0113)	(3110)	0	3	5
a	1	0	2	2	(0122)	(2210)	0	2	5
b	1	2	0	2	(0122)	(2210)	0	2	5
c	1	2	2	2	(1222)	(2221)	1	2	7
ab	2	1	1	1	(1112)	(2111)	1	2	5
ac	2	1	3	1	(1123)	(3211)	1	3	7
bc	2	3	1	1	(1123)	(3211)	1	3	7
abc	3	2	2	0	(0223)	(3220)	0	3	7

 Table 1: Table of Hamming distances for κ from Example 7

Example 7. If $\mathcal{P} = \{a, b, c\}$, take $\kappa = \{(\neg(a \wedge b) \wedge \neg c) \vee (a \wedge b \wedge c)\}$, for which we get that $[\kappa] = \{\emptyset, a, b, abc\}$. For the interpretation $w = \emptyset$, we get that $\overrightarrow{d_H(w, \kappa)} = (0113)$, $\overleftarrow{d_H(w, \kappa)} = (3110)$, $\min d_H(w, \kappa) = 0$, $\max d_H(w, \kappa) = 3$ and $\sum d_H(w, \kappa) = 5$. The distances and aggregated distances for each interpretation are depicted in Table 1. Notice how the models of κ are distributed when the interpretations are ranked according to the different aggregation functions used: we have $\emptyset \approx_{\kappa}^{\text{H}, \min} a$, since $\min d_H(\emptyset, \kappa) = \min d_H(a, \kappa) = 0$, but $\emptyset <_{\kappa}^{\text{H}, \text{lmin}} a$, since $(0113) \leq_{\text{lex}} (0122)$. Also, we have that $c <_{\kappa}^{\text{H}, \text{max}} abc$, $c <_{\kappa}^{\text{H}, \text{lmax}} abc$ and $ab <_{\kappa}^{\text{H}, \text{sum}} abc$, i.e., models of κ are not minimal in $\leq_{\kappa}^{\text{H}, \text{max}}$, $\leq_{\kappa}^{\text{H}, \text{lmax}}$ and $\leq_{\kappa}^{\text{H}, \text{sum}}$. In particular, $\leq_{\kappa}^{\text{H}, \text{max}}$ makes the models of $\bar{\kappa}$ (i.e., abc , bc , ac and \emptyset) the least plausible interpretations.

The agreement and disagreement operators ($\circ^{d, \text{agr}}$ and $\circ^{d, \text{dagr}}$) are simpler than they appear: the idea behind $\circ^{d, \text{agr}}$ is to allow interpretations other than the models of κ as the minimal elements of the preorder $\leq_{\kappa}^{d, \text{agr}}$ is 0 if it is either a model of κ , or it is equidistant from every model of κ (i.e., its centrality is 0) and it is the ‘closest’ interpretation to κ with this property. The disagreement operator $\circ^{d, \text{dagr}}$ works in similar fashion, by making models of $\bar{\kappa}$ and interpretations minimally equidistant to them the least plausible interpretations in $\leq_{\kappa}^{d, \text{dagr}}$.

Example 8. If $\mathcal{P} = \{a, b, c\}$, take κ such that $[\kappa] = \{a, b, c\}$, and notice that $d_H(\emptyset, \kappa) = (111)$ and $d_H(abc, \kappa) = (222)$, i.e., they are both equidistant to κ , hence their centrality is 0. However, \emptyset is closer to κ than abc (its displacement is 0, compared to abc ’s displacement of 1), and $\text{agr}(\emptyset, \kappa) = 0$. Thus, what $\leq_{\kappa}^{\text{H}, \text{agr}}$ does is to give a minimal score to models of κ and to the minimally equidistant interpretation \emptyset . By contrast, $\leq_{\kappa}^{\text{H}, \text{dagr}}$ gives a maximal score to the models of $\bar{\kappa}$ and to the maximally equidistant interpretation abc .

All operators proposed generate a total preorder \leq_{κ} over interpretations, but differ in how they arrange models of κ in \leq_{κ} : this corresponds to the different attitudes an agent can have towards κ prior to any revision. The operator $\circ^{\text{H}, \min}$, known as Dalal’s operator [Dalal, 1988], considers all models of κ as the most plausible elements in \leq_{κ} and is the only operator for which $\kappa \circ \mu$ is equivalent to $\kappa \wedge \mu$ when $\kappa \wedge \mu$ is consistent. Similarly, $\circ^{\text{H}, \text{lmin}}$ also ranks models of κ as more plausible than any other interpretation, but discriminates among models of κ . The operators $\leq_{\kappa}^{\text{H}, \text{max}}$ and $\leq_{\kappa}^{\text{H}, \text{lmax}}$ push away models of $\bar{\kappa}$, under the assumption that they are the most implausible possible worlds. Their difference between them is that $\leq_{\kappa}^{\text{H}, \text{max}}$ considers models of $\bar{\kappa}$ equally

	R ₆	R ₇	R ₈	R ₉	R ₁₀	R _{IDF}	R _{STB}
$\circ^{\text{H}, \text{min}}$	✓	✓	×	×	✓	✓	✓
$\circ^{\text{H}, \text{lmin}}$	×	✓	×	×	✓	×	✓
$\circ^{\text{H}, \text{agr}}$	✓	×	×	×	✓	✓	✓
$\circ^{\text{H}, \text{max}}$	×	×	✓	✓	×	×	✓
$\circ^{\text{H}, \text{lmax}}$	×	×	×	✓	×	×	✓
$\circ^{\text{H}, \text{dagr}}$	×	×	✓	×	×	✓	✓
$\circ^{\text{H}, \text{sum}}$	×	×	×	×	✓	×	✓
\circ^{dr}	✓	✓	×	×	✓	✓	✓
\circ^{fg}	✓	×	✓	×	✓	✓	✓

Table 2: Satisfaction of axioms

implausible, whereas $\leq_{\kappa}^{\text{H}, \text{lmax}}$ uses the more fine-grained lexicographic approach. The operator $\circ^{\text{H}, \text{agr}}$ makes models of κ the most plausible elements in \leq_{κ} but also allows other interpretations on that position, in particular certain interpretations that are equidistant to κ as per Example 8. The intuition is that an interpretation equally distanced from models of κ is like a compromise point of view, with good chances of being correct if it is close to κ . The operator $\circ^{\text{H}, \text{dagr}}$ is the dual of $\circ^{\text{H}, \text{agr}}$ and, finally, operator $\circ^{\text{H}, \text{sum}}$ evokes utilitarian approaches by choosing interpretations that minimize the sum of the distances to each model of κ .

Quick reflection shows that operators obtained with drastic distance d_{D} collapse into two main categories. To get a grasp on this fact, consider first the *drastic revision operator* \circ^{dr} defined, for $\kappa \in 2_{\text{cons}}^{\text{Prop}}$ and $\mu \in \text{Prop}$, as $\kappa \circ^{\text{dr}} \mu = \kappa \wedge \mu$, if $\kappa \wedge \mu$ is consistent, and μ otherwise, and the *forgetful revision operator* \circ^{fg} defined as $\kappa \circ^{\text{fg}} \mu = \mu$.

Proposition 5. For any knowledge base κ and formula μ , it holds that $\kappa \circ^{\text{D}, \text{min}} \mu \equiv \kappa \circ^{\text{D}, \text{lmin}} \mu \equiv \kappa \circ^{\text{D}, \text{lmax}} \mu \equiv \kappa \circ^{\text{D}, \text{sum}} \mu \equiv \kappa \circ^{\text{dr}} \mu$. Moreover, $\kappa \circ^{\text{D}, \text{agr}} \mu \equiv \kappa \circ^{\text{fg}} \mu$ and $\kappa \circ^{\text{D}, \text{max}} \mu \equiv \kappa \circ^{\text{D}, \text{dagr}} \mu \equiv \begin{cases} \kappa \circ^{\text{dr}} \mu, & \text{if } \kappa \text{ is complete,} \\ \kappa \circ^{\text{fg}} \mu, & \text{otherwise.} \end{cases}$

With Hamming distance the landscape is more diverse, as the different attitudes the operators assume towards models of κ lead to genuinely different revision strategies. Nonetheless, certain relationships between the operators still hold, with lexicographic operators being the most discriminating, in the sense that they pick formulas with fewer models.

Proposition 6. For any $\kappa \in 2_{\text{cons}}^{\text{Prop}}$ and $\mu \in \text{Prop}$, we have that $\kappa \circ^{\text{H}, \text{lmin}} \mu \models \kappa \circ^{\text{H}, \text{min}} \mu \models \kappa \circ^{\text{H}, \text{agr}} \mu$ and $\kappa \circ^{\text{H}, \text{lmax}} \mu \models \kappa \circ^{\text{H}, \text{max}} \mu \models \kappa \circ^{\text{H}, \text{dagr}} \mu$.

All operators generate total preorders over interpretations, so by Theorem 1 they all satisfy axioms R_{1–5}. Satisfaction with respect to the newly introduced postulates is clarified below.

Proposition 7. For $d \in \{\text{D}, \text{H}\}$ and $f \in \{\text{min}, \text{lmin}, \text{max}, \text{lmax}, \text{agr}, \text{dagr}, \text{sum}\}$, the operators $\circ^{d, f}$ satisfy postulates R_{6–10}, IDF, STB as per Table 2.

Proof. It is already known that $\leq_{\kappa}^{\text{H}, \text{min}}$ (known as Dalal’s operator [Dalal, 1988; Katsuno and Mendelzon, 1991]) satisfies axioms R_{5–6}. Operator $\circ^{\text{H}, \text{lmin}}$ satisfies R₇ (via property o₇) because if $w_1 \in [\kappa]$ and $w_2 \notin [\kappa]$, then $\min d_{\text{H}}(w_1, \kappa) = 0$, while $\min d_{\text{H}}(w_2, \kappa) > 0$; hence $d_{\text{H}}(w_1, \kappa) <_{\text{lex}} d_{\text{H}}(w_2, \kappa)$

and $w_1 <_{\kappa}^{\text{H}, \text{lmin}} w_2$. For $\circ^{\text{H}, \text{lmin}}$ and axiom R₆, take $[\kappa] = \{a, b, ab\}$ and $[\mu] = \{a, b, ab\}$. We get that $[\kappa \circ^{\text{H}, \text{lmin}} \mu] = \{ab\}$. The operator $\circ^{\text{H}, \text{agr}}$ satisfies axiom R₆ because it makes all models of κ , and potentially other interpretations as well (which is the reason why it does not satisfy axiom R₇), as the equally most plausible interpretations in $\leq_{\kappa}^{\text{H}, \text{agr}}$. Since all these operators place the models of κ on the lowest levels of \leq_{κ} , they all satisfy axiom R₁₀. Postulates R_{8–9} are not satisfied by $\circ^{\text{H}, \text{min}}$, $\circ^{\text{H}, \text{lmin}}$ and $\circ^{\text{H}, \text{agr}}$ because these operators do not make models of $\bar{\kappa}$ the least plausible interpretations in \leq_{κ} : if $\kappa = a \vee b$, then $\bar{\kappa}$ shares models with κ , yet these (along with all other models of $a \vee b$) are among the most plausible interpretations in $\leq_{\kappa}^{\text{H}, \text{min}}$, $\leq_{\kappa}^{\text{H}, \text{lmin}}$ and $\leq_{\kappa}^{\text{H}, \text{agr}}$. The case for $\circ^{\text{H}, \text{max}}$, $\circ^{\text{H}, \text{lmax}}$ and $\circ^{\text{H}, \text{dagr}}$ is analogous to the one for $\circ^{\text{H}, \text{min}}$, $\circ^{\text{H}, \text{lmin}}$ and $\circ^{\text{H}, \text{agr}}$, as they can be seen as duals of each other. For the operator $\circ^{\text{H}, \text{sum}}$, take $[\kappa] = \{a, b, c\}$ and $[\mu] = \{\emptyset, a, b, c\}$. We get that $[\kappa \circ^{\text{H}, \text{sum}} \mu] = \{\emptyset\}$, as \emptyset minimizes the sum of the Hamming distances to the models of κ : this is a counter-example to axioms R_{6–7}. For R_{8–9}, take $[\mu'] = \{\emptyset, ab, ac, bc\}$. For $\circ^{\text{H}, \text{sum}}$ and R₁₀, notice that adding w' to $[\kappa]$ creates a new column for w' in the table of distances, in which the distance corresponding to w' is 0, i.e., the score assigned to w' in $\leq_{\kappa \vee \kappa'}^{\text{H}, \text{sum}}$ does not increase with respect to $\leq_{\kappa}^{\text{H}, \text{sum}}$. For the operators defined with d_{D} , we use Proposition 5. In particular, operators $\circ^{\text{D}, \text{max}}$ and $\circ^{\text{D}, \text{dagr}}$ satisfy a postulate if and only if both \circ^{dr} and \circ^{fg} satisfy it. Satisfaction of R_{IDF} and R_{STB} is straightforward, keeping in mind how the various operators arrange the models of κ in the generated preorders. \square

5 Conclusion

We have looked at the classical revision axioms from the point of view of what they assume about an agent’s attitude towards its initial beliefs, and argued that this attitude is embedded in a specific axiom (the KM postulate corresponding to *Inclusion* and *Vacuity*). By varying this axiom we were able to put forward and characterize a wide range of revision operators, and refine previously entangled intuitions in the process. Analysis of the new operators also uncovered the principles of indifference to already held beliefs (R_{IDF}) and stability (R_{STB}). Further work is needed to link these notions to the other axioms, to map out their interplay and to provide them with semantic characterizations.

At the same time, the more fine grained view on the types of attitudes an agent can have towards its initial beliefs raises the question of what these attitudes are good for, i.e., whether they can be used for tasks such as learning [Kelly, 1998; Baltag *et al.*, 2011]. The idea here is to view revision as part of an ongoing process by which the agent continuously refines its representation of the outside world, with the aim of settling on stable, correct information. Such a task, we think, provides a natural benchmark for revision operators, and it has the potential to connect belief revision to other topics of importance to the field of AI.

Acknowledgments

This work has been supported by the Austrian Science Fund (FWF): P30168-N31, W1255-N23, Y698.

References

- [Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *J. Symb. Log.*, 50(2):510–530, 1985.
- [Arrow, 1959] Kenneth J. Arrow. Rational choice functions and orderings. *Economica*, 26(102):121–127, 1959.
- [Baltag *et al.*, 2011] Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. Belief revision as a truth-tracking process. In *Proc. of TARK 2011*, pages 187–190, 2011.
- [Benferhat *et al.*, 2005] Salem Benferhat, Sylvain Lagrue, and Odile Papini. Revision of Partially Ordered Information: Axiomatization, Semantics and Iteration. In *Proc. of IJCAI 2005*, pages 376–381, 2005.
- [Bonanno, 2009] Giacomo Bonanno. Rational choice and AGM belief revision. *Artif. Intell.*, 173(12-13):1194–1203, 2009.
- [Dalal, 1988] Mukesh Dalal. Investigations into a Theory of Knowledge Base Revision. In *Proc. of IJCAI 1988*, pages 475–479, 1988.
- [Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the Logic of Iterated Belief Revision. *Artif. Intell.*, 89(1-2):1–29, 1997.
- [Endriss, 2016] Ulle Endriss. Judgment Aggregation. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 399–426. Cambridge University Press, 2016.
- [Fermé and Hansson, 2018] Eduardo Fermé and Sven Ove Hansson. *Belief Change: Introduction and Overview*. Springer, 2018.
- [Foley, 1993] Richard Foley. *Working Without a Net: A Study of Egocentric Epistemology*. Oxford University Press, 1993.
- [French, 1988] Simon French. *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood Limited, 1988.
- [Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.
- [Herzig and Rifi, 1999] Andreas Herzig and Omar Rifi. Propositional Belief Base Update and Minimal Change. *Artif. Intell.*, 115(1):107–138, 1999.
- [Kahneman *et al.*, 1982] Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press, 1982.
- [Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. On the Difference between Updating a Knowledge Base and Revising It. In *Proc. of KR 1991*, pages 387–394, 1991.
- [Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.*, 52(3):263–294, 1992.
- [Kelly, 1998] Kevin T. Kelly. The Learning Power of Belief Revision. In *Proc. of TARK 1998*, pages 111–124, 1998.
- [Konieczny and Pérez, 2011] Sébastien Konieczny and Ramón Pino Pérez. Logic Based Merging. *J. Philosophical Logic*, 40(2):239–270, 2011.
- [Konieczny *et al.*, 2004] Sébastien Konieczny, Jérôme Lang, and Pierre Marquis. DA^2 merging operators. *Artif. Intell.*, 157(1-2):49–79, 2004.
- [Levesque, 1986] Hector J Levesque. Making believers out of computers. *Artif. Intell.*, 30(1):81–108, 1986.
- [Moulin, 1991] Hervé Moulin. *Axioms of Cooperative Decision Making*. Number 15 in Econometric Society Monographs. Cambridge University Press, 1991.
- [Peppas and Williams, 2016] Pavlos Peppas and Mary-Anne Williams. Kinetic Consistency and Relevance in Belief Revision. In *Proc. of JELIA 2016*, pages 401–414, 2016.
- [Rott, 2001] Hans Rott. *Change, choice and inference: A study of belief revision and nonmonotonic reasoning*. Number 42 in Oxford Logic Guides. Oxford University Press, 2001.
- [Ryan, 1996] Mark Ryan. Belief Revision and Ordered Theory Presentations. In *Logic, Action, and Information*, pages 129–151, 1996.
- [Sen, 1984] Amartya K. Sen. *Collective Choice and Social Welfare*. NY: North-Holland, 2nd edition, 1984.