Belief Revision Operators with Varying Attitudes Towards Initial Beliefs

Adrian Haret and **Stefan Woltran**
Institute of Logic and Computation, TU Wien, Austria

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 initial knowledge base. This postulate assumes a conservative attitude towards initial beliefs, in the sense that an agent faced with the task of revising 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, and make the case that these operators can be put to use when doing revision in fragments of propositional logic. We provide axiomatizations of these operators by varying the aforementioned postulate and obtain representation results which characterize the new types of operators in terms of preorders on possible worlds. We also present concrete examples for each new type of operator, using notions inspired from decision theory.

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 that a revision operator is expected to abide by are encoded as logical axioms (Alchourrón *et al.* 1985; Gärdenfors 1988; Katsuno and Mendelzon 1992b; Fermé and Hansson 2018). Remarkably, 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 that 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 axiomatization (Katsuno and Mendelzon 1992b). This attitude spells out how 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 a *comparison 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.

Moreover, we argue that our insights can be put to use in the current research stream of *fragment-based* revision (Creignou *et al.* 2014; Delgrande and Peppas 2015; Zhuang *et al.* 2017; Delgrande *et al.* 2018), which seeks to understand belief revision in more applied formalisms. The main motivation behind this line of research is that there is much to be gained computationally if we assume that an agent states its beliefs in a simpler language, e.g., a restricted fragment of propositional logic, and that by revising the agent is able to stay within this fragment. However, it turns out that such a mild assumption (i.e., that revision returns a knowledge base expressible in a certain fragment) clashes systematically with commonly accepted properties, including the combined version of the aforementioned *In-*

clusion and *Vacuity* postulates: their conjunction puts certain demands on the underlying language (e.g., that the conjunction of κ and μ is always expressible in it), which are not met in all scenarios that interest us. Thus, there is an unexpected payoff in looking, as we do here, at weaker versions of the standard postulates and their semantic characterizations.

Considering alternatives to the classical revision postulates has a distinguished history, going as far back as the original publications in the field (Gärdenfors 1988; Katsuno and Mendelzon 1992a; Hansson 1999; Herzig and Rifi 1999; Velázquez-Quesada 2017). However, we believe that a systematic analysis of the intuition underlying KM postulate R_2, as we perform here, sheds new light on familiar topics.

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 base κ is a finite set of propositional formulas, which we typically identify with the conjunction of its formulas. Thus, we think of a knowledge base κ as a single formula, i.e., $\bigwedge_{\mu\in\kappa}\mu$. 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 $\{ab, bc\}$; hence, for instance, $[a\vee b] = \{a, b, ab\}$. If μ_1 and μ_2 are propositional formulas, 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 $2^{\mathsf{Prop}}_{\mathsf{cons}}$. If κ is a propositional knowledge base, then *the dual $\overline{\kappa}$ of κ* is obtained by replacing every literal l appearing in κ with its negation. If w is an interpretation, the dual interpretation $\overline{w}$ is $\mathcal{P}\setminus w$. If W is a set of interpretations, its dual $\overline{W}$ is defined as $\{\overline{w}\mid w\in W\}$.

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

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

If $\mathcal{M}$ is a set, then $Bin(\mathcal{M})$ 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 $\min_{\leq}\mathcal{M}=\{x\in\mathcal{M}\mid \nexists x'\in\mathcal{M}$ such that $x'<x\}$. An *assignment from $\mathcal{M}_1$ to $\mathcal{M}_2$* is a function $\alpha\colon\mathcal{M}_1\to Bin(\mathcal{M}_2)$. We write $\leq_\kappa$ 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 $[\varphi_W]=W$. If w_1 and w_2 are two interpretations, $\varphi_{1,2}$ is a propositional formula such that $[\varphi_{1,2}]=\{w_1,w_2\}$. A renaming ρ is a permutation on the set $\mathcal{P}$ which, applied to a formula μ, yields a formula $\rho(\mu)$ whose atoms are replaced according to ρ.

Example 2. If ρ is a renaming such that $\rho(a)=b$, $\rho(b)=c$ and $\rho(c)=a$, then $\rho(\neg a\wedge(b\to c))=\neg b\wedge(c\to a)$.

Revision: axioms and characterizations

A propositional revision operator is a function $\circ\colon 2^{\mathsf{Prop}}_{\mathsf{cons}}\times\mathsf{Prop}\to\mathsf{Prop}$. 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 any inconsistencies between κ and μ and to satisfy other rationality criteria, presented below.

Basic postulates. If $\kappa,\kappa_1,\kappa_2\in 2^{\mathsf{Prop}}_{\mathsf{cons}}$ and $\mu,\mu_1,\mu_2\in\mathsf{Prop}$, we first single out the following core set of axioms:

(R_1) $\kappa\circ\mu\models\mu$.

(R_2) If μ is consistent, then $\kappa\circ\mu$ is consistent.

(R_3) If $\kappa_1\equiv\kappa_2$ and $\mu_1\equiv\mu_2$, then $\kappa_1\circ\mu_1\equiv\kappa_2\circ\mu_2$.

(R_4) $(\kappa\circ\mu_1)\wedge\mu_2\models\kappa\circ(\mu_1\wedge\mu_2)$.

(R_5) If $(\kappa\circ\mu_1)\wedge\mu_2$ is consistent, then $\kappa\circ(\mu_1\wedge\mu_2)\models(\kappa\circ\mu_1)\wedge\mu_2$.

(R_N) $\rho(\kappa\circ\mu)\equiv\rho(\kappa)\circ\rho(\mu)$.

Postulates R_{1-5} form part of the standard set of KM postulates (Katsuno and Mendelzon 1992b), saying that a revision operator incorporates new information μ (R_1), returns a consistent output if μ is consistent (R_2), performs its task irrespective of how beliefs are written down (R_3) and satisfies some coherence constraints when the revision formula is varied (R_{4-5}).[1] We have found it suitable to add here a neutrality axiom R_N, requiring that a revision operator does not favor propositional atoms based solely on their names. This idea is expressed by requiring the revision output to be invariant under a renaming ρ of atoms, and, in conjunction with the irrelevance of syntax postulate R_3, it is perhaps natural to expect it from any revision operator. The inspiration for postulate R_N is the social choice literature (see (Rothe 2015), Chapter 4) and, though it has appeared in belief change before, under various guises (Herzig and Rifi 1999; Marquis and Schwind 2014; Haret *et al.* 2016), neutrality usually goes unstated in standard presentations of revision.

A revision operator is *basic* if it satisfies postulates R_{1-5} and *neutral* if it satisfies postulate R_N. We will typically assume that all revision operators we work with are basic.

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 $2^{\mathsf{Prop}}_{\mathsf{cons}}$ a binary relation $\leq_\kappa$ on interpretations in $\mathcal{U}$: to revise κ by μ, then, becomes equivalent to choosing the best models of μ in $\leq_\kappa$. 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

[1] Note that R_2 as we show it does not coincide with KM postulate R_2. The KM postulate shows up in the latter part of this section.

satisfy a set of properties, to be introduced in the following, that are conducive to rational choice.

First of all, let us extend our notion of renamings to cover interpretations. Thus, if w is an interpretation and ρ is a renaming of atoms, then $\rho(w)$ is an interpretation obtained by replacing every atom p in w with $\rho(p)$. If $\mathcal{M}$ is a set of interpretations, then $\rho(\mathcal{M}) = \{\rho(w) \mid w \in \mathcal{M}\}$.

Example 3. Take $\mathcal{M} \subseteq \mathcal{U}$ such that $\mathcal{M} = \{a, ab, bc\}$ and a renaming ρ such that $\rho(a) = b$, $\rho(b) = c$ and $\rho(c) = a$. We get $\rho(\mathcal{M}) = \{\rho(a), \rho(ab), \rho(bc)\} = \{b, bc, ac\}$.

For an assignment from consistent knowledge bases in $2^{\mathsf{Prop}}_{\mathsf{cons}}$ to binary relations on $\mathcal{U}$, a revision operator $\circ$ and $\kappa, \kappa_1, \kappa_2 \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$, $w_1, w_2 \in \mathcal{U}$, we look at these properties:

(o_1) $\leq_\kappa$ is reflexive.

(o_2) $\leq_\kappa$ is transitive.

(o_3) If $\kappa_1 \equiv \kappa_2$, then $\leq_{\kappa_1} = \leq_{\kappa_2}$.

(o_4) $\leq_\kappa$ is total.

(o_5) $[\kappa \circ \mu] = \min_{\leq_\kappa}[\mu]$, for any propositional formula μ.

(o_N) $w_1 \leq_\kappa w_2$ iff $\rho(w_1) \leq_{\rho(\kappa)} \rho(w_2)$.

An assignment is *basic* if it satisfies properties o_{1-4}. Notice that properties o_{1-2} imply that $\leq_\kappa$ is a preorder on $\mathcal{U}$. Adding property o_4 makes $\leq_\kappa$ total, and o_3 adds an independence of syntax aspect to the assignment. If, on top of this, $\leq_\kappa$ satisfies o_5, we say that the assignment α *represents the revision operator* $\circ$ (and that $\circ$ *is represented by* α). We call an assignment *neutral* if it satisfies property o_N. The overloading of notation ('basic', 'neutral') is intentional: properties $\mathsf{o}_{1-5,\,\mathsf{N}}$ define a class of rankings on interpretations that fully characterize revision operators satisfying axioms $\mathsf{R}_{1-5,\,\mathsf{N}}$. The following results make this precise.

Theorem 1. A revision operator satisfies postulates R_{1-5} iff there exists an assignment from $2^{\mathsf{Prop}}_{\mathsf{cons}}$ to $\mathcal{U}$ representing it such that, for any $\kappa \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$, the preorder $\leq_\kappa$ satisfies properties o_{1-5}.

This result follows from existing work on revision and can be extracted, for example, from the proof of the representation result in (Katsuno and Mendelzon 1992b). Keeping in mind that for any $\mu \in \mathsf{Prop}$ and renaming ρ it holds that $[\rho(\mu)] = \rho([\mu])$, we also get the following result.

Theorem 2. If $\circ$ is a basic revision operator and α is an assignment representing it, then $\circ$ satisfies axiom R_N iff α satisfies property o_N.

Proof. Recall, first, that we denote by $\varphi_{1,2}$ a propositional formula such that $[\varphi_{1,2}] = \{w_1, w_2\}$. ("$\Rightarrow$") Take a basic revision operator $\circ$ that satisfies axiom R_N, and the assignment which represents it. Take w_1 and w_2 and suppose, first, that $w_1 \leq_\kappa w_2$. Then $w_1 \in [\kappa \circ \varphi_{1,2}]$, and hence $\rho(w_1) \in \rho([\kappa \circ \varphi_{1,2}])$. We get that $\rho(w_1) \in [\rho(\kappa \circ \varphi_{1,2})]$ and by postulate R_N it follows that $\rho(w_1) \in [\rho(\kappa) \circ \rho(\varphi_{1,2})]$. This implies that $\rho(w_1) \in \min_{\leq_{\rho(\kappa)}}[\rho(\varphi_{1,2})]$, which implies that $\rho(w_1) \in \min_{\leq_{\rho(\kappa)}} \rho([\varphi_{1,2}])$. Thus, $\rho(w_1) \leq_{\rho(\kappa)} \rho(w_2)$. Conversely, suppose that $\rho(w_1) \leq_{\rho(\kappa)} \rho(w_2)$. This implies that $\rho(w_1) \in \min_{\leq_{\rho(\kappa)}} \rho([\varphi_{1,2}])$. By postulate R_N, we get that $w_1 \leq_\kappa w_2$. ("$\Leftarrow$") Take an assignment that satisfies property o_N and the revision operator $\circ$ represented by it. We have that $[\rho(\kappa \circ \mu)] = \rho([\kappa \circ \mu]) = \rho(\min_{\leq_\kappa}[\mu])$, and $[\rho(\kappa) \circ \rho(\mu)] = \min_{\leq_{\rho(\kappa)}}[\rho(\mu)] = \min_{\leq_{\rho(\kappa)}} \rho([\mu])$. We show that $[\rho(\kappa \circ \mu)] = [\rho(\kappa) \circ \rho(\mu)]$ by double inclusion. Take, first, $\rho(w_1) \in \rho(\min_{\leq_\kappa}[\mu])$, for $w_1 \in \min_{\leq_\kappa}[\mu]$, and $\rho(w_2) \in \rho([\mu])$, for $w_2 \in [\mu]$. Then $w_1 \leq_\kappa w_2$ and by o_N we get that $\rho(w_1) \leq_{\rho(\kappa)} \rho(w_2)$, which implies that $\rho(w_1) \in \min_{\leq_{\rho(\kappa)}} \rho([\mu])$. This shows that $[\rho(\kappa \circ \mu)] \subseteq [\rho(\kappa) \circ \rho(\mu)]$. Next, take $\rho(w_1) \in \min_{\leq_{\rho(\kappa)}} \rho([\mu])$, for $w_1 \in [\mu]$, and $\rho(w_2) \in \rho([\mu])$, for $w_2 \in [\mu]$. We get that $\rho(w_1) \leq_{\rho(\kappa)} \rho(w_2)$, which by o_N implies that $w_1 \leq_\kappa w_2$. Thus, $w_1 \in \min_{\leq_\kappa}[\mu]$ and hence $\rho(w_1) \in \rho(\min_{\leq_\kappa}[\mu])$. $\square$

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 $\leq_\kappa$ that depends on initial beliefs κ, and always picks the minimal models of μ according to $\leq_\kappa$. 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 $\leq_\kappa$. Thus, a revision operator can be seen as a *choice function* over sets of interpretations: for instance, postulates R_{5-6} are equivalent to what is known as the *Weak Axiom of Revealed Preference (WARP)* and, taken together, postulates R_{1-5} characterize choice functions *rationalizable* by total preorders. 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, in itself, not a new insight, the topic having been studied before under various guises (Rott 2001; Bonanno 2009).

Theorem 2 adds invariance of $\leq_\kappa$ under renamings. The ranking $\leq_\kappa$ is usually thought of as a plausibility ranking, i.e., the assessment of an agent believing κ as to which possible worlds are more or less plausible.

Example 4. 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). In other words, the doctor has initial information $\kappa = a$ and a plausibility ranking $\leq_\kappa$ over possible worlds, depicted in Figure 1. The doctor then revises by $\mu = b$ and settles on a possible state of affairs that is most plausible according to their plausibility ranking $\leq_\kappa$, i.e., $[\kappa \circ \mu] = \min_{\leq_\kappa}[\mu] = \{abc\}$. Note that the doctor, in this case, believes that the situation represented by $a \wedge b \wedge c$ is more likely than $a \wedge b \wedge \neg c$.

One way of thinking of postulates $\mathsf{R}_{1-5,\,\mathsf{N}}$ is that they axiomatize neutral total preorders on interpretations. These preorders nominally depend on κ, but nothing in postulates $\mathsf{R}_{1-5,\,\mathsf{N}}$ 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 $\mathsf{R}_{1-5,\,\mathsf{N}}$ are consistent with arbitrary attitudes towards κ. How should the models of κ stand in

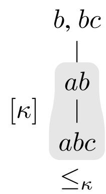

Figure 1: Revision scenario of Example 4, showing the preorder $\leq_\kappa$ on the basis of which the revision result is constructed: only models of $\mu = b$ are depicted; the models of κ among this set are highlighted in grey.

relation to all other interpretations? In Example 4 we catch a glimpse of one possible answer: the agent there starts off with some information κ and differentiates among the possible worlds consistent with κ: some of these interpretations are, *a priori*, more plausible than others. Still, as a whole, models of κ are more plausible than any *other* interpretations consistent with the new information μ—the agent is biased towards the possible worlds consistent with κ, an attitude which fits with the idea of κ being the agent's *belief*. Are there, now, other ways of arranging the models of κ in $\leq_\kappa$, ways that span the space of possible such attitudes? We study this question through the lens of additional axioms.

Attitudes towards initial beliefs. Let $\kappa, \kappa' \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$, $\mu \in \mathsf{Prop}$, and consider the following postulates:

(R_6) If $\kappa \wedge \mu$ is consistent, then $\kappa \wedge \mu \models \kappa \circ \mu$.

(R_7) If $\kappa \wedge \mu$ is consistent, then $\kappa \circ \mu \models \kappa \wedge \mu$.

(R_8) If $\mu \models \overline{\kappa}$, then $\kappa \circ \mu \equiv \mu$.

(R_9) If $\mu \not\models \overline{\kappa}$, then $(\kappa \circ \mu) \wedge \overline{\kappa}$ is inconsistent.

(R_{10}) 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 itself 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 1992b).

Postulates R_{8-9} focus on the dual knowledge base $\overline{\kappa}$ 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 $\overline{\kappa}$ will be a formula whose models are complements of the models of κ.

Example 5. If $\mathcal{P} = \{a, b, c\}$ and $\kappa = \{a \wedge b\}$ is a knowledge base over the alphabet $\mathcal{P}$, then $\overline{\kappa} = \{\neg a \wedge \neg b\}$, and we have that $[\kappa] = \{ab, abc\}$, while $[\overline{\kappa}] = \{\emptyset, c\}$

Thus, if κ is reasonably specific (e.g., is a conjunction of literals), then $\overline{\kappa}$ can be thought of as a point of view opposite to that of κ. Of course, this analogy breaks down if κ is a knowledge base such as $a \vee b$, or $a \vee \neg a$, where κ and $\overline{\kappa}$ share models. Our point here is simply that there are situations in which it makes sense to view κ and $\overline{\kappa}$ as embodying opposing stances, and in which one would like to place bounds on the revision function in terms of how it treats information encoded by the opposing point of view $\overline{\kappa}$: 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), and when κ will be a complete formula; or if κ is a 'vivid' knowledge base (Levesque 1986), or encodes something like an agent's preferred bundle from a set of available items, in which case κ can be required to be a conjunction of literals.

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 $\overline{\kappa}$, 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 $\overline{\kappa}$ only if it has no other choice in the matter: the models of $\overline{\kappa}$ 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 6. An agent intends to go to an art museum, the beach and a concert, i.e., $\kappa = \{a \wedge b \wedge c\}$. 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 place the models of κ in a total preorder $\leq_\kappa$ ($\kappa, \kappa' \in 2^{\mathsf{Prop}}_{\mathsf{cons}}, w_1, w_2, w' \in \mathcal{U}$):

(o_6) If $w_1 \in [\kappa]$, then $w_1 \leq_\kappa w_2$.

(o_7) If $w_1 \in [\kappa]$ and $w_2 \notin [\kappa]$, then $w_1 <_\kappa w_2$.

(o_8) If $w_1 \in [\overline{\kappa}]$, then $w_2 \leq_\kappa w_1$.

(o_9) If $w_1 \in [\overline{\kappa}]$ and $w_2 \notin [\overline{\kappa}]$, then $w_2 <_\kappa w_1$.

(o_{10}) If $w' \leq_\kappa w$ and $[\kappa'] = \{w'\}$, then $w' \leq_{\kappa \vee \kappa'} w$.

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

Theorem 3. If $\circ$ is a revision operator satisfying postulates R_{1-5} and α is an assignment from $2^{\mathsf{Prop}}_{\mathsf{cons}}$ to $\mathcal{U}$ which satisfies properties o_{1-5} then, for any $\kappa \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$ and $\mu \in \mathsf{Prop}$, the following equivalences hold:

(1) $\circ$ satisfies axiom R_6 iff $\leq_\kappa$ satisfies property o_6;
(2) $\circ$ satisfies axiom R_7 iff $\leq_\kappa$ satisfies property o_7;
(3) $\circ$ satisfies axiom R_8 iff $\leq_\kappa$ satisfies property o_8;
(4) $\circ$ satisfies axiom R_9 iff $\leq_\kappa$ satisfies property o_9;
(5) $\circ$ satisfies axiom R_{10} iff $\leq_\kappa$ satisfies property o_{10}.

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, first, an assignment α satisfying property o_6, and the revision operator $\circ$ represented by it. Let us 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 o_5, this is equivalent to showing that $w \in \min_{\leq_\kappa}[\mu]$. Take an arbitrary interpretation $w' \in [\mu]$. Since $w \in [\kappa]$, we can apply property o_6 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 $\leq_\kappa$ satisfies property o_6, 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 o_5, 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, $w_1 \leq_\kappa w_2$.

For equivalence 2, we show again each direction in turn. ("$\Rightarrow$") Take an assignment α satisfying property o_7 and the revision operator $\circ$ represented by it. Let us assume that $\kappa \wedge \mu$ is consistent. Take $w \in [\kappa \circ \mu]$, and suppose $w \notin [\kappa \wedge \mu]$. By property o_5, we have that $[\kappa \circ \mu] = \min_{\leq_\kappa}[\mu]$, and hence $w \in [\mu]$. Thus, the fact that $w \notin [\kappa \wedge \mu]$ implies that $w \notin [\kappa]$. But, by assumption, it holds that $[\kappa \wedge \mu] \neq \emptyset$. Thus, there exists $w' \in [\kappa \wedge \mu]$ and, by property o_7, it follows that $w' <_\kappa w$. But we also have that $w \in \min_{\leq_\kappa}[\mu]$, which implies that it cannot be the case that $w' <_\kappa w$, which is a contradiction. ("$\Leftarrow$") Take a basic revision operator $\circ$ satisfying R_7 and the preorder $\leq_\kappa$ that represents it. To show that $\leq_\kappa$ satisfies property o_7, take $w_1 \in [\kappa]$ and $w_2 \notin [\kappa]$. We then have that $\kappa \wedge \varphi_{1,2}$ is consistent and hence, by axiom R_7, that $\kappa \circ \varphi_{1,2} \models \kappa \wedge \varphi_{1,2}$. Since $[\kappa \circ \varphi_{1,2}]$ is, by axioms R_{1-2}, a non-empty subset of $[\varphi_{1,2}] = \{w_1, w_2\}$, we have that at least one of w_1 and w_2 is in $[\kappa \circ \varphi_{1,2}]$. Notice, now, that we cannot have $w_2 \in [\kappa \circ \varphi_{1,2}]$, since it would follow that $w_2 \in [\kappa \wedge \varphi_{1,2}]$ and $w_2 \in [\kappa]$, which is a contradiction. Thus, $[\kappa \circ \varphi_{1,2}] = \{w_1\}$. Since $\leq_\kappa$ represents $\circ$, we have by property o_5 that $[\kappa \circ \varphi_{1,2}] = \min_{\leq_\kappa}[\varphi_{1,2}]$. It follows that $w_1 <_\kappa w_2$.

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 o_{10} to derive the conclusion. $\square$

Theorem 3 is better understood through an illustration of how such preorders treat models of κ. Property o_6 says that

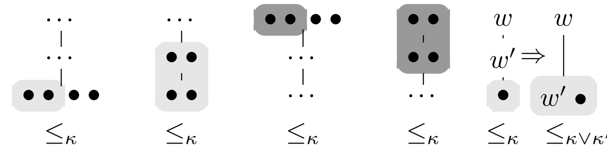

(a) o_6 (b) o_7 (c) o_8 (d) o_9 (e) o_{10}

Figure 2: Schematic view of prototypical preorders satisfying each of the properties o_{6-9}; models of κ are in the light gray area, models of $\overline{\kappa}$ are in the dark gray area.

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 o_7 states that there are no countermodels of κ more plausible than the models of κ, but the models of κ themselves may not be equally plausible (Figure 2-(b)). Properties o_{8-9} say that models of the dual knowledge base κ are the least plausible interpretations in $\leq_\kappa$ (Figure 2-(c,d)), while property o_{10} 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 o_{1-7} define what is more commonly known as a *faithful assignment*, placing all and only models of κ on the lowest level of $\leq_\kappa$. This corresponds to an agent that holds its initial beliefs to be the most plausible states of affairs (Katsuno and Mendelzon 1992b). Consequently, Theorem 1 plus equivalences 1-2 from Theorem 3 make up the classical representation result for belief revision operators (Katsuno and Mendelzon 1992b). Here we have opted for a more fine-grained approach to the placement of models of κ in $\leq_\kappa$, 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.

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:

($\mathsf{R}_{\mathsf{IDF}}$) $\kappa \circ \kappa \equiv \kappa$.

This property, called here $\mathsf{R}_{\mathsf{IDF}}$ (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 $\mathsf{R}_{\mathsf{IDF}}$,[2] but R_7 does not. Thus, if an agent is allowed to rank models of κ unequally, then $\mathsf{R}_{\mathsf{IDF}}$ is not guaranteed to hold.

Example 7. Consider the knowledge base $\kappa = \{a \vee b\}$ and a revision operator that satisfies axiom R_7, and which orders

[2]The converse is not true: $\mathsf{R}_{\mathsf{IDF}}$ enforces only that models of κ are equally plausible, but not where they are placed in $\leq_\kappa$.

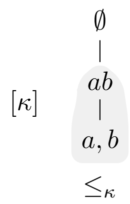

Figure 3: $[\kappa] = \{a, b, ab\}$, but $[\kappa \circ \kappa] = \{a, b\}$

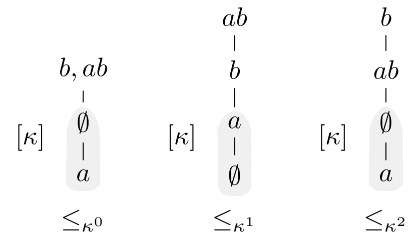

Figure 4: Repeated revision by κ ends up jumping around from $\{a\}$ to $\{\emptyset\}$.

the models of κ as in Figure 3. We get that $[\kappa \circ \kappa] = \{a, b\}$, i.e., $\kappa \circ \kappa \equiv (a \leftrightarrow \neg b)$.

What this fact points to is 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 then leaves open the possibility that the agent assigns different degrees of plausibility to states of affairs it counts as part of its belief κ: indeed, this is the point of view we endorse here. This is in contrast to more standard approaches, which consider that an agent assigns equal degrees of plausibility to all items of its belief. Thus, new information which confirms an agent's belief 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 phenomenon we take to be modeled by Example 7.

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 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:

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

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

Proposition 4. 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. □

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 4, we get that repeated revision by κ leads to a chain of ever more specific knowledge bases, i.e., $\emptyset \subset \cdots \subseteq [\kappa^{i+1}] \subseteq [\kappa^i] \subseteq \cdots \subseteq [\kappa^0]$. 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 5. 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 $\overline{\kappa}$ 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 8. Take a knowledge base $\kappa = \{\neg b\}$ and a revision operator satisfying R_{8-9} which orders interpretations as shown in Figure 4. We have that $[\kappa^0] = [\kappa] = \{\emptyset, a\}$, and $[\kappa^1] = [\kappa \circ \kappa] = \{\emptyset\}$ and $[\kappa^2] = [\kappa^1 \circ \kappa] = \{a\}$. By axiom R_3, we get that subsequent revisions by κ jump around between $\{a\}$ and $\{\emptyset\}$, i.e., $[\kappa^3] = \{\emptyset\}$, $[\kappa^4] = \{a\}$, and so on, therefore never settling on a stable answer.

The issue of stability suggests another dimension along which revision operators can be analyzed, with Proposition 5 and Example 8 showing that a revision operator does not satisfy it trivially. Example 8, in particular, shows that there is interplay between the preorders $\leq_\kappa$ and $\leq_{\kappa'}$, where $\kappa' \models \kappa$, which is relevant to the question of whether a revision operator is stable or not. This interplay is reminiscent of issues surrounding iterated revision and kinetic consistency (Darwiche and Pearl 1997; Peppas and Williams 2016) and pursuing it further is worthwhile, though it would take us too far afield of the aims of the current work.

Concrete propositional revision operators

Having characterized revision operators in terms of assignments on interpretations, we ask ourselves what is a natural way to construct such assignments. The usual way of exploiting the insight afforded by Theorem 1 is to use some type of distance d between interpretations, interpreted as a measure of plausibility of one interpretation relative to the other. The distance d is then used to compare interpretations with respect to how plausible (or close) they are with respect to the initial beliefs κ.

Our answer extends this method in a way reminiscent of techniques used to construct belief merging operators (Konieczny *et al.* 2004; Konieczny and Pérez 2011), i.e., by employing *two* main ingredients. The first ingredient is a *distance* between interpretations, defined as a function $d \colon \mathcal{U} \times \mathcal{U} \to \mathbb{R}_+$ 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 $\kappa \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$ such that $[\kappa] = \{w_1, \ldots, w_n\}$, the *vector of distances from w to κ* is $d(w, \kappa) = (d(w, w_1), \ldots, d(w, w_n))$. If there is no danger of ambiguity we omit commas and simply write $d(w, \kappa)$ as a string of numbers.

Example 9. If $w = a$, $[\kappa] = \{a, b, ab\}$, d is a distance function such that $d(a, a) = 0$, $d(a, b) = 2$ and $d(a, ab) = 1$, then $d(w, \kappa) = (0, 2, 1)$, written as $d(w, \kappa) = (021)$.

We mention here two prominent examples of distance. The first one, called *the drastic distance* d_{D} works by the all-or-nothing rule: $d_{\mathsf{D}}(w_1, w_2) = 0$ if $w_1 = w_2$, and 1 otherwise. The second one is *the Hamming distance* d_{H}, which counts the number of atoms on which two interpretations differ.

The second ingredient is a *comparison function*, which is a family of functions $f\colon \mathbb{R}^n \to \mathbb{R}$ 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 $\leq_{\mathsf{lex}}$. Minimal and maximal elements of $d(w, \kappa)$ are denoted by $\min d(w, \kappa)$ and $\max d(w, \kappa)$, respectively. We define $\sum d(w, \kappa) = \sum_{w_i \in [\kappa]} d(w, w_i)$. The *centrality of w with respect to κ* is $\mathrm{cen}(w, \kappa) = \max d(w, \kappa) - \min d(w, \kappa)$. The *displacement of w with respect to κ* is $\mathrm{dis}(w, \kappa) = \min d(w, \kappa) - \min d(w^*, \kappa)$, where w^* is an interpretation such that $\min d(w^*, \kappa)$ is minimal among all the interpretations w' for which $\mathrm{cen}(w', \kappa) = \mathrm{cen}(w, \kappa)$. Finally, the *agreeability index of w with respect to κ* is defined as $\mathrm{agr}(w, \kappa) = \min\{\min d(w, \kappa),\ \mathrm{cen}(w, \kappa) + \mathrm{dis}(w, \kappa)\}$, while the *disagreeability index of w with respect to κ* is $\mathrm{dagr}(w, \kappa) = n - \mathrm{agr}(w, \overline{\kappa})$, where $n = |\mathcal{P}|$.

Let us now put the two ingredients together. Given a distance d and comparison 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., defined by taking $[\kappa \circ^{d,\, f} \mu] = \min_{\leq_\kappa^{d,\, f}}[\mu]$. Assuming some distance d between interpretations, we will look at the types of rankings defined in Table 1.

$$
\begin{array}{lll}
w_1 \leq_\kappa^{d,\, \min} w_2 & \text{iff} & \min d(w_1, \kappa) \leq \min d(w_2, \kappa), \\
w_1 \leq_\kappa^{d,\, \mathrm{lmin}} w_2 & \text{iff} & \overrightarrow{d(w_1, \kappa)} \leq_{\mathsf{lex}} \overrightarrow{d(w_2, \kappa)}, \\
w_1 \leq_\kappa^{d,\, \mathrm{agr}} w_2 & \text{iff} & \mathrm{agr}(w_1, \kappa) \leq \mathrm{agr}(w_2, \kappa), \\
w_1 \leq_\kappa^{d,\, \max} w_2 & \text{iff} & \max d(w_1, \kappa) \leq \max d(w_2, \kappa), \\
w_1 \leq_\kappa^{d,\, \mathrm{lmax}} w_2 & \text{iff} & \overleftarrow{d(w_1, \kappa)} \leq_{\mathsf{lex}} \overleftarrow{d(w_2, \kappa)}, \\
w_1 \leq_\kappa^{d,\, \mathrm{dagr}} w_2 & \text{iff} & \mathrm{dagr}(w_1, \kappa) \leq \mathrm{dagr}(w_2, \kappa), \\
w_1 \leq_\kappa^{d,\, \mathrm{sum}} w_2 & \text{iff} & \sum d(w_1, \kappa) \leq \sum d(w_2, \kappa),
\end{array}
$$

Table 1: Types of rankings, defined for a distance d

We illustrate these operators with the following examples.

Example 10. Take $\kappa = \{\neg a \vee \neg b \vee (a \wedge b \wedge c)\}$, for which we get that $[\kappa] = \{\emptyset, a, b, abc\}$. For the interpretation $w = \emptyset$, we get that $\overrightarrow{d_{\mathsf{H}}(w, \kappa)} = (0113)$, $\overleftarrow{d_{\mathsf{H}}(w, \kappa)} = (3110)$, $\min d_{\mathsf{H}}(w, \kappa) = 0$, $\max d_{\mathsf{H}}(w, \kappa) = 3$ and $\sum d_{\mathsf{H}}(w, \kappa) = 5$. The distances and computed aggregation functions for each interpretation are depicted in Table 2. Notice how the models of κ are distributed when the interpretations are ranked according to the different comparison functions used: we have $\emptyset \approx_\kappa^{\mathsf{H},\, \min} a$, since $\min d_{\mathsf{H}}(a, \kappa) = \min d_{\mathsf{H}}(a, \kappa) = 0$, but $\emptyset <_\kappa^{\mathsf{H},\, \mathrm{lmin}} a$, since $(0113) \leq_{\mathsf{lex}} (0122)$. Also, we have that $c <_\kappa^{\mathsf{H},\, \max} abc$, $c <_\kappa^{\mathsf{H},\, \mathrm{lmax}} abc$ and $ab <_\kappa^{\mathsf{H},\, \mathrm{sum}} abc$, i.e., models of κ are not minimal in $\leq_\kappa^{\mathsf{H},\, \max}$, $\leq_\kappa^{\mathsf{H},\, \mathrm{lmax}}$ and $\leq_\kappa^{\mathsf{H},\, \mathrm{sum}}$. In particular, $\leq_\kappa^{\mathsf{H},\, \max}$ makes the models of $\overline{\kappa}$ (i.e., abc, bc, ac and $\emptyset$) the least plausible interpretations.

	$\emptyset$	a	b	abc	$\overrightarrow{d_{\mathsf{H}}(w, \kappa)}$	$\overleftarrow{d_{\mathsf{H}}(w, \kappa)}$	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 2: Table of Hamming distances for κ from Example 10

The agreement and disagreement operators ($\circ^{d,\, \mathrm{agr}}$ and $\circ^{d,\, \mathrm{dagr}}$) are simpler than they appear: $\circ^{d,\, \mathrm{agr}}$ works as a scoring rule allowing interpretations other than the models of κ as the minimal elements of the preorder $\leq_\kappa$. Notice that the score of an interpretation in $\leq_\kappa^{d,\, \mathrm{agr}}$ 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 $\circ^{d,\, \mathrm{dagr}}$ does something similar, by making models of $\overline{\kappa}$ and interpretations minimally equidistant to them the least plausible interpretations in $\leq_\kappa^{d,\, \mathrm{dagr}}$.

Example 11. Take κ such that $[\kappa] = \{a, b, c\}$, and notice that $d_{\mathsf{H}}(\emptyset, \kappa) = (111)$ and $d_{\mathsf{H}}(abc, \kappa) = (222)$, 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 $\mathrm{agr}(\emptyset, \kappa) = 0$. Thus, what $\leq_\kappa^{\mathsf{H},\, \mathrm{agr}}$ does is to give a minimal score to models of κ and to the minimally equidistant interpretation $\emptyset$. By contrast, $\leq_\kappa^{\mathsf{H},\, \mathrm{dagr}}$ gives a maximal score to the models of $\overline{\kappa}$ and to the maximally equidistant interpretation abc.

All operators proposed here generate a total preorder $\leq_\kappa$ over interpretations, but differ in how they arrange the models of κ in $\leq_\kappa$, corresponding to different attitudes an agent can have towards its beliefs κ prior to any revision. The operator $\circ^{\mathsf{H},\, \min}$, known as Dalal's operator (Dalal 1988), considers all models of κ as the most plausible elements in $\leq_\kappa$ and is the only operator for which $\kappa \circ \mu$ is equivalent to $\kappa \wedge \mu$ when $\kappa \wedge \mu$ is consistent. Similarly, $\circ^{\mathsf{H},\, \mathrm{lmin}}$ also ranks models of κ as more plausible than any other interpretation, but discriminates among models of κ. The operators $\leq_\kappa^{\mathsf{H},\, \max}$ and $\leq_\kappa^{\mathsf{H},\, \mathrm{lmax}}$ work by pushing away models of $\overline{\kappa}$, under the assumption that they are the most implausible possible worlds. They differ as to how they arrange the models of $\overline{\kappa}$: $\leq_\kappa^{\mathsf{H},\, \max}$ considers them equally implausible, whereas $\leq_\kappa^{\mathsf{H},\, \mathrm{lmax}}$ uses the more fine-grained lexicographic approach. The operator $\circ^{\mathsf{H},\, \mathrm{agr}}$ makes models of κ the most plausible elements in $\leq_\kappa$ but does not stop here and allows

$$\circ^{d,\,\mathrm{lmin}} \rightarrow \circ^{d,\,\mathrm{min}} \longrightarrow \circ^{d,\,\mathrm{agr}} \qquad \circ^{d,\,\mathrm{lmax}} \rightarrow \circ^{d,\,\mathrm{max}} \qquad \circ^{d,\,\mathrm{sum}}$$

Figure 5: An arrow from x to y indicates that the output of operator x implies output of operator y for the same input

other interpretations on that position, in particular certain interpretations that are equidistant to κ as per Example 11. The intuition here is that an interpretation equally distanced from models of κ is something like a compromise point of view, with good chances of being correct if it is close to κ. The operator $\circ^{\mathsf{H},\,\mathrm{dagr}}$ is the dual of $\circ^{\mathsf{H},\,\mathrm{agr}}$ and, finally, operator $\circ^{\mathsf{H},\,\mathrm{sum}}$ evokes utilitarian approaches by choosing interpretations that minimize the sum of the distances to each model of κ, i.e., are close to κ on an aggregate level.

Plugging in the drastic and Hamming distances in the definition in Table 1 would seem to give us a considerable number of operators, but 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* $\circ^{\mathsf{dr}}$ defined, for $\kappa \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$ and $\mu \in \mathsf{Prop}$, as $\kappa \circ^{\mathsf{dr}} \mu = \kappa \wedge \mu$, if $\kappa \wedge \mu$ is consistent, and μ otherwise, and the *forgetful revision operator* $\circ^{\mathsf{fg}}$ defined as $\kappa \circ^{\mathsf{fg}} \mu = \mu$.

Proposition 6. For any knowledge base κ and formula μ, it holds that $\kappa \circ^{\mathsf{D},\,\mathrm{min}} \mu \equiv \kappa \circ^{\mathsf{D},\,\mathrm{lmin}} \mu \equiv \kappa \circ^{\mathsf{D},\,\mathrm{lmax}} \mu \equiv \kappa \circ^{\mathsf{D},\,\mathrm{sum}} \mu \equiv \kappa \circ^{\mathsf{dr}} \mu$. Moreover, $\kappa \circ^{\mathsf{D},\,\mathrm{agr}} \mu \equiv \kappa \circ^{\mathsf{fg}} \mu$ and:

$$\kappa \circ^{\mathsf{D},\,\mathrm{max}} \mu \equiv \kappa \circ^{\mathsf{D},\,\mathrm{dagr}} \mu \equiv \begin{cases} \kappa \circ^{\mathsf{dr}} \mu, & \text{if } \kappa \text{ is complete,} \\ \kappa \circ^{\mathsf{fg}} \mu, & \text{otherwise.} \end{cases}$$

Proof. If κ is a complete knowledge base and $[\kappa] = \{w_0\}$, then $d_{\mathsf{D}}(w,\kappa) = d_{\mathsf{D}}(w,w_0)$, for any interpretation w. In other words, $\overrightarrow{d_{\mathsf{D}}(w,\kappa)}$ yields (0) if $w = w_0$ and (1) otherwise. Thus, $\max d_{\mathsf{D}}(w,\kappa) = 0$ if $w = w_0$, and $\max d_{\mathsf{D}}(w,\kappa) = 1$, otherwise. It is then straightforward to see that if $w_0 \in [\mu]$, then $[\kappa \circ^{\mathsf{D},\,\mathrm{max}} \mu] = \{w_0\} = [\kappa \wedge \mu]$, and if $w_0 \notin [\mu]$, then $[\kappa \circ^{\mathsf{D},\,\mathrm{max}} \mu] = [\mu]$. If κ is not complete, then $\overrightarrow{d_{\mathsf{D}}(w,\kappa)}$ yields $(01\ldots1)$ if $w \in [\kappa]$, and $(1\ldots1)$ otherwise. It is now straightforward to see that the remaining statements of Proposition 6 hold. ☐

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.

Proposition 7. For any $\kappa \in 2^{\mathsf{Prop}}_{\mathsf{cons}}$ and $\mu \in \mathsf{Prop}$, we have that $\kappa \circ^{\mathsf{H},\,\mathrm{lmin}} \mu \models \kappa \circ^{\mathsf{H},\,\mathrm{min}} \mu \models \kappa \circ^{\mathsf{H},\,\mathrm{agr}} \mu$ and $\kappa \circ^{\mathsf{H},\,\mathrm{lmax}} \mu \models \kappa \circ^{\mathsf{H},\,\mathrm{max}} \mu \models \kappa \circ^{\mathsf{H},\,\mathrm{dagr}} \mu$.

The relationship between the different operators is depicted in Figure 5. One can see that lexicographic operators are the most discriminating, in the sense that they pick formulas with fewer models (i.e., more specific formulas).

Since all operators generate total preorders over interpretations, by Theorem 1 they all satisfy axioms R_{1-5}. Satisfaction with respect to the newly introduced postulates is clarified by the following result.

Proposition 8. For a distance $d \in \{\mathsf{D}, \mathsf{H}\}$ and a comparison function $f \in \{\min, \mathrm{lmin}, \max, \mathrm{lmax}, \mathrm{agr}, \mathrm{dagr}, \mathrm{sum}\}$, the operators $\circ^{d,\,f}$ satisfy postulates R_{6-10} as shown in Table 3.

	R_6	R_7	R_8	R_9	R_{10}	R_{N}	$\mathsf{R}_{\mathsf{IDF}}$	$\mathsf{R}_{\mathsf{STB}}$
$\circ^{\mathsf{H},\,\mathrm{min}}$	✓	✓	×	×	✓	✓	✓	✓
$\circ^{\mathsf{H},\,\mathrm{lmin}}$	×	✓	×	×	✓	✓	×	✓
$\circ^{\mathsf{H},\,\mathrm{agr}}$	✓	×	×	×	✓	✓	✓	✓
$\circ^{\mathsf{H},\,\mathrm{max}}$	×	×	✓	✓	×	✓	×	✓
$\circ^{\mathsf{H},\,\mathrm{lmax}}$	×	×	×	✓	×	✓	×	✓
$\circ^{\mathsf{H},\,\mathrm{dagr}}$	×	×	✓	×	×	✓	✓	✓
$\circ^{\mathsf{H},\,\mathrm{sum}}$	×	×	×	×	✓	✓	×	✓
$\circ^{\mathsf{dr}}$	✓	✓	×	×	✓	✓	✓	✓
$\circ^{\mathsf{fg}}$	✓	×	✓	×	✓	✓	✓	✓

Table 3: Satisfaction of axioms

Proof. It is already known that $\leq^{d,\,\mathrm{min}}_{\kappa}$ (known as Dalal's operator (Dalal 1988; Katsuno and Mendelzon 1991)) satisfies axioms R_{5-6}. To see why the operator $\circ^{\mathsf{H},\,\mathrm{lmin}}$ satisfies R_7, notice that if $w_1 \in [\kappa]$ and $w_2 \notin [\kappa]$, then the first element in $d(w_1,\kappa)$ is 0, while the first element in $d(w_2,\kappa)$ is strictly greater than 0. This implies that $\overrightarrow{d(w_1,\kappa)} <_{\mathsf{lex}} \overrightarrow{d(w_2,\kappa)}$, which in turn implies that $w_1 <^{\mathsf{H},\,\mathrm{lmin}}_{\kappa} w_2$. Hence property o_7 is satisfied, which implies that axiom R_7 is satisfied. Additionally, it cannot be the case that $w_3 <^{\mathsf{H},\,\mathrm{lmin}}_{\kappa} w_1$, for any $w_3 \notin [\kappa]$, which shows that property o_6 (and hence axiom R_6) is satisfied. For $\circ^{\mathsf{H},\,\mathrm{lmin}}$ and axiom R_6, take $[\kappa] = \{a, b, ab\}$ and $[\mu] = \{a, b, ab\}$. We get that $[\kappa \circ^{\mathsf{H},\,\mathrm{lmin}} \mu] = \{a, b\}$. The operator $\circ^{\mathsf{H},\,\mathrm{agr}}$ satisfies axiom R_6 because it makes all models of κ, and potentially other interpretations as well (which is the reason why it does not satisfy axiom R_7), as the equally most plausible interpretations in $\leq^{\mathsf{H},\,\mathrm{agr}}_{\kappa}$. Since all these operator place the models of κ on the lowest levels of $\leq_{\kappa}$, they all satisfy axiom R_{10}.

To see why postulates R_{8-9} are not satisfied by $\circ^{d,\,\mathrm{min}}$, $\circ^{d,\,\mathrm{lmin}}$ or $\circ^{d,\,\mathrm{dagr}}$, notice that these operators do not make models of $\overline{\kappa}$ as the least plausible interpretations in $\leq_{\kappa}$. Thus, if $\kappa = a \vee b$, then $\overline{\kappa}$ shares some models with κ, yet these models (along with all other models of $a \vee b$) will be among the most plausible interpretations in $\leq^{d,\,\mathrm{min}}_{\kappa}$, $\leq^{d,\,\mathrm{lmin}}_{\kappa}$ and $\leq^{d,\,\mathrm{agr}}_{\kappa}$. The one exception is the forgetful operator $\circ^{\mathsf{fg}}$, which satisfies R_8 trivially.

The case for $\circ^{\mathsf{H},\,\mathrm{max}}$, $\circ^{\mathsf{H},\,\mathrm{lmax}}$ and $\circ^{\mathsf{H},\,\mathrm{dagr}}$ is analogous to the one for $\circ^{\mathsf{H},\,\mathrm{min}}$, $\circ^{\mathsf{H},\,\mathrm{lmin}}$ and $\circ^{\mathsf{H},\,\mathrm{agr}}$, as they can be seen as duals of each other. For the operator $\circ^{\mathsf{H},\,\mathrm{sum}}$, take $[\kappa] = \{a, b, c\}$ and $[\mu] = \{\emptyset, a, b, c\}$. We get that $[\kappa \circ^{\mathsf{H},\,\mathrm{sum}} \mu] = \{\emptyset\}$, 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 $\circ^{\mathsf{H},\,\mathrm{sum}}$ and R_{10}, 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 $\leq^{\mathsf{H},\,\mathrm{sum}}_{\kappa \vee \kappa'}$ does not increase with respect to $\leq^{\mathsf{H},\,\mathrm{sum}}_{\kappa}$. Satisfaction of $\mathsf{R}_{\mathsf{IDF}}$ and $\mathsf{R}_{\mathsf{STB}}$ is straightforward, keeping in mind how the various operators arrange the models of κ in the generated preorders. The neutrality axiom R_{N} is discussed separately. ☐

With respect to neutrality, Table 3 shows that all operators introduced so far satisfy postulate R_{N}. This is guaranteed by a property of the distances which we will, by an overload of notation, call by the same name. Thus, a distance d is *neutral* if, for any renaming ρ and interpretations w_1 and w_2, it holds that $d(w_1, w_2) = d(\rho(w_1), \rho(w_2))$. It is straightforward to see that the drastic and Hamming distances are neutral. Furthermore, if d is neutral, then $d(w, \kappa) = d(\rho(w), \rho(\kappa))$, for any $w \in \mathcal{U}$, and $w_1 \leq_\kappa^{d,f} w_2$ iff $\rho(w_1) \leq_{\rho(\kappa)}^{d,f} \rho(w_2)$, for all selection functions introduced so far. Thus, the preorders $\leq_\kappa^{d,f}$ satisfy property $\mathsf{o_N}$ and, by Theorem 2, the operators represented by them satisfy postulate $\mathsf{R_N}$. It should be kept in mind that neutrality is not guaranteed by the other postulates, but the way in which concrete operators are usually defined (i.e., by appeal to neutral distances) indicates that neutrality is part of our basic understanding of how a revision operator should behave. And, in general, there seems to be no *a priori* reason for looking at non-neutral operators. However, it turns out that such operators cannot be avoided when we move to a fragment of propositional logic.

Revision of Horn theories

Discussions of stability and varying attitudes to initial beliefs notwithstanding, one might still question the rationale behind giving up the KM axiom R_2 (axioms R_{6-7} in the present work): indeed, why fix something if it is not broken? In response, we will use this section to argue that there are situations where revision is warranted, but in which axioms R_{6-7} *cannot* occur together.

Recall that a clause is called *Horn* if at most one of its literals is positive, a *Horn formula* is a conjunction of Horn clauses and a *Horn knowledge base* is a finite set of Horn formulas. The set of all Horn formulas is $\mathcal{L}_{\mathsf{Horn}}$. Horn formulas are characterized semantically by closure under intersection, i.e., if W is a set of interpretations then there exists a Horn formula φ such that $[\varphi] = W$ iff $\mathrm{Cl}_\cap(W) = W$.

As mentioned in the introduction, there is good reason to want to do revision on Horn knowledge bases, and to ask that the result is still a Horn knowledge base. We thus consider H-revision operators $\bullet\colon 2^{\mathcal{L}_{\mathsf{Horn}}}_{\mathsf{cons}} \times \mathsf{Prop} \to \mathcal{L}_{\mathsf{Horn}}$, mapping a Horn knowledge base κ and a propositional formula μ to a Horn knowledge base $\kappa \bullet \mu$. The natural next step now would be to adapt the regular propositional postulates to the new setting of H-revision. Postulates R_{1-5} can be adapted seamlessly, and we may denote the adapted postulates $\mathsf{R}^{\mathsf{H}}_{1-5}$, but consider what happens if we try to introduce a postulate that is an adapted version of the standard KM postulate R_2, saying that $\kappa \bullet \mu \equiv \kappa \wedge \mu$, if $\kappa \wedge \mu$ is consistent. For ease of reference, we split this postulate into two weaker postulates:

($\mathsf{R}^{\mathsf{H}}_6$) If $\kappa \wedge \mu$ is consistent, then $\kappa \wedge \mu \models \kappa \bullet \mu$.

($\mathsf{R}^{\mathsf{H}}_7$) If $\kappa \wedge \mu$ is consistent, then $\kappa \bullet \mu \models \kappa \wedge \mu$.

It turns out that postulate $\mathsf{R}^{\mathsf{H}}_6$ cannot be used in this context.

Example 12. Take a Horn knowledge base $\kappa = \{\neg a \vee \neg b\}$ and a $\mu = a \leftrightarrow \neg b$. Clearly, $\kappa \wedge \mu$ is satisfiable and, moreover, $\kappa \wedge \mu \equiv \mu$. However, $[\kappa \wedge \mu] = \{a, b\}$, which is not equal to $\mathrm{Cl}_{\mathsf{Horn}}([\kappa \wedge \mu])$ and thus does not represent any Horn formula. If $\mathsf{R}^{\mathsf{H}}_6$ were true, then with $\mathsf{R}^{\mathsf{H}}_1$ we would have to conclude that $[\kappa \bullet \mu] = [\kappa \wedge \mu] = [\mu]$, a contradiction.

Thus, it seems that H-revision operators cannot be axiomatized in a way that is analogous to propositional operators satisfying axioms R_1 and R_6, or, equivalently: we cannot model an agent who, when revising a Horn knowledge base κ, always makes the models of κ equally plausible. This can be stated as a corollary, following directly from Example 12.

Corollary 9. If an H-revision operator satisfies axiom $\mathsf{R}^{\mathsf{H}}_1$, then it does not satisfy axiom $\mathsf{R}^{\mathsf{H}}_6$.

Since we are not prepared to sacrifice axiom R_1, we are left with satisfying axiom R_7. What about neutrality? It turns out that this is also problematic for H-revision operators: since H-revision by $\mu = a \leftrightarrow \neg b$ must return, by $\mathsf{R}^{\mathsf{H}}_{1,3}$, a consistent result which implies μ and is a Horn formula, such an operator must effectively choose exactly one of the interpretations a and b. This leads to a clash with the adapted neutrality postulate, which we may call $\mathsf{R}^{\mathsf{H}}_{\mathsf{N}}$.

Proposition 10. If an H-revision operator satisfies axioms $\mathsf{R}^{\mathsf{H}}_{1-3}$, then it does not satisfy axiom $\mathsf{R}^{\mathsf{H}}_{\mathsf{N}}$.

Proof. Take $\kappa = a \wedge b$ and $\mu = a \leftrightarrow \neg b$. Suppose $\kappa \bullet \mu \equiv a \wedge \neg b$, and take a renaming ρ such that $\rho(a) = b$ and $\rho(b) = a$ to get a contradiction. $\square$

The move to be explicit about neutrality and to split the standard KM postulate R_2 into two distinct properties (postulates $\mathsf{R}_{6,7}$), either of which can be turned off, finds additional justification here: we can see now that properties taken for granted in the propositional case break down when restricting the language, and a thorough analysis of what are rational, or desirable, properties for revision must take this into account.

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 (KM axiom R_2), By varying this axiom and calling attention to a commonly overlooked neutrality property, we were able to put forward and characterize a wide range of revision operators, and refine previously entangled intuitions in the process. We also showed that this level of analysis is needed when working in restricted fragments of propositional logic, where the KM axiom R_2 cannot be satisfied and must therefore be broken down into two separate components (axioms R_{6-7} in the current work).

Analysis of the new operators uncovered the principles of indifference to already held beliefs ($\mathsf{R}_{\mathsf{IDF}}$) and stability ($\mathsf{R}_{\mathsf{STB}}$). Further work is needed to link these notions to the other axioms, to map out the interplay between them, and to provide them with semantic characterizations. Following the line of reasoning initiated in the previous section, a natural follow-up would be to consider the proposed postulates in fragments of propositional logic and to look for characterizations in terms of preorders on possible worlds. 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 put to use in an area such as learning (Kelly 1998; Baltag *et al.* 2011).

Acknowledgments

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

References

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.

Kenneth J. Arrow. Rational choice functions and orderings. *Economica*, 26(102):121–127, 1959.

Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. Belief revision as a truth-tracking process. In *Proc. of TARK 2011*, pages 187–190, 2011.

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.

Giacomo Bonanno. Rational choice and AGM belief revision. *Artif. Intell.*, 173(12-13):1194–1203, 2009.

Nadia Creignou, Odile Papini, Reinhard Pichler, and Stefan Woltran. Belief revision within fragments of propositional logic. *J. Comput. Syst. Sci.*, 80(2):427–449, 2014.

Mukesh Dalal. Investigations into a Theory of Knowledge Base Revision. In *Proc. of IJCAI 1988*, pages 475–479, 1988.

Adnan Darwiche and Judea Pearl. On the Logic of Iterated Belief Revision. *Artif. Intell.*, 89(1-2):1–29, 1997.

James P. Delgrande and Pavlos Peppas. Belief revision in Horn theories. *Artif. Intell.*, 218:1–22, 2015.

James P. Delgrande, Pavlos Peppas, and Stefan Woltran. General Belief Revision. *J. ACM*, 65(5):29:1–29:34, 2018.

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.

Eduardo Fermé and Sven Ove Hansson. *Belief Change: Introduction and Overview*. Springer, 2018.

Richard Foley. *Working Without a Net: A Study of Egocentric Epistemology*. Oxford University Press, 1993.

Simon French. *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood Limited, 1988.

Peter Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.

Sven Ove Hansson. A Survey of non-Prioritized Belief Revision. *Erkenntnis*, 50(2):413–427, 1999.

Adrian Haret, Andreas Pfandler, and Stefan Woltran. Beyond IC Postulates: Classification Criteria for Merging Operators. In *Proc. of ECAI 2016*, pages 372–380, 2016.

Andreas Herzig and Omar Rifi. Propositional Belief Base Update and Minimal Change. *Artif. Intell.*, 115(1):107–138, 1999.

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.

Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gärdenfors, editor, *Belief Revision*, pages 183–203. Cambridge University Press, 1992.

Hirofumi Katsuno and Alberto O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.*, 52(3):263–294, 1992.

Kevin T. Kelly. The Learning Power of Belief Revision. In *Proc. of TARK 1998*, pages 111–124, 1998.

Sébastien Konieczny and Ramón Pino Pérez. Logic Based Merging. *J. Philosophical Logic*, 40(2):239–270, 2011.

Sébastien Konieczny, Jérôme Lang, and Pierre Marquis. DA^2 merging operators. *Artif. Intell.*, 157(1-2):49–79, 2004.

Hector J Levesque. Making believers out of computers. *Artif. Intell.*, 30(1):81–108, 1986.

Pierre Marquis and Nicolas Schwind. Lost in translation: Language independence in propositional logic - application to belief change. *Artif. Intell.*, 206:1–24, 2014.

Hervé Moulin. *Axioms of cooperative decision making*. Number 15. Cambridge university press, 1991.

Pavlos Peppas and Mary-Anne Williams. Kinetic Consistency and Relevance in Belief Revision. In *Proc. of JELIA 2016*, pages 401–414, 2016.

Jörg Rothe, editor. *Economics and Computation*. Springer, 2015.

Hans Rott. *Change, choice and inference: A study of belief revision and nonmonotonic reasoning*. Number 42. Oxford University Press, 2001.

Mark Ryan. Belief Revision and Ordered Theory Presentations. In *Logic, Action, and Information*, pages 129–151, 1996.

Amartya K. Sen. *Collective Choice and Social Welfare*. NY: North-Holland, 2nd edition, 1984.

Fernando R. Velázquez-Quesada. On Subtler Belief Revision Policies. In *Proc. of LORI 2017*, pages 314–329, 2017.

Zhiqiang Zhuang, Maurice Pagnucco, and Yan Zhang. Inter-Definability of Horn Contraction and Horn Revision. *J. Philosophical Logic*, 46(3):299–332, 2017.