Implementing Bounded Revision via Lexicographic Revision and C-revision

Meliha Sezgin, Gabriele Kern-Isberner

TU Dortmund University, Germany
meliha.sezgin@tu-dortmund.de, gabriele.kern-isberner@tu-dortmund.de

Abstract

New information in the context of real life settings usually is accompanied by some kind of supplementary information that indicates context, reliability, or expertise of the information's source. Bounded Revision (BR) displays an iterated belief revision mechanism that takes as input a new information accompanied by a reference sentence acting as supplementary information, which specifies the depth with which the new input shall be integrated in the posterior belief state. The reference sentence specifies which worlds in the prior belief state are affected by the change mechanism. We show that Bounded Revision can be characterized by three simple, yet elegant postulates and corresponds to a special case of a lexicographic revision, which inherits all relevant features of BR. Furthermore, we present methodological implementations of BR including conditional revision with c-revisions, making it directly usable for conditional revision tools.

Introduction

The vivid research area of iterated belief revision (Darwiche and Pearl 1997) investigates in which manner an agent incorporates new information that may be inconsistent with their current epistemic state represented via total preorders. During the revision inconsistencies are cleared out, yet meta-information accompanying the new input, e.g. reliability, are not taken into account. Extensive studies in psychology (see e.g. (Wolf, Rieger, and Knauff 2012; Sparks and Rapp 2011; Pornpitakpan 2004)) have shown that the rather naive acceptance of any kind of input information does not correspond to the way individuals revise their beliefs in real life settings. In fact, expertise, reliability, and other factors impact the way we incorporate new information. We strongly believe that realistic models of revision should provide the necessary means to represent this kind of meta-information.

We present a real life example illustrating what we have in mind. Anna knows that George Washington was the first president of the United States and she is also aware of the fact that all presidents of the US lived in the White House. In Washington, she learns from a tour guide that George Washington never lived in the White House and that he served as a president from 1789 to 1797. So, Anna needs to revise her beliefs but will do so only as far as she still accepts that George Washington was an American president and she trusts the tour guide's statement about Washington's term of office as president. The example illustrates how supplementary information affects the revision. In the course of the revision, Anna should not give up her prior belief that Washington was an US president. Also the input's context, e.g., the tour guide's trustworthiness, is crucial for the revision and should determine the depth with which Anna trusts his statement about Washington's term of office.

There have been several approaches to belief revision with input information that is accompanied by some kind of supplementary information, see e.g. (Hunter 2021; Ammar and Ismail 2021; Sezgin and Kern-Isberner 2022), indicating trust or reliance towards the input information. In this paper, we present a revision method in a qualitative and a semi-quantitative framework based on *Bounded Revision* (BR) which was firstly introduced by (Rott 2012). BR takes as input two pieces of information, an input sentence β and a reference sentence α, thus displaying a two-dimensional revision function, where the input sentence is accompanied by some kind of meta-information. The reference sentence α guides the revision process and the main goal of BR by β w.r.t. α can be expressed as follows: *Accept β in the posterior entrenchment relation as far as α and just a little further.* Despite its intuitive strength the formal implementation of BR remains unclear in (Rott 2012). Our investigations provide new insights in which way α influences the posterior belief state and how it acts as a bound for the acceptance of β.

Our main contributions are as follows:

- We clarify the strategy underlying BR by β w.r.t. α and summarize the worlds affected by the belief change by a single formula
- We present a representation theorem implementing semantic postulates that characterize the iterated belief change of BR by β w.r.t. to the reference sentence α
- We show that BR corresponds to a special lexicographic revision and thus implements a more fine-grained revision mechanism via supplementary information
- We define BR for the semi-quantitative framework of ranking functions (Spohn 1988) and implement it methodologically as a conditional revision making the ensuing application of BR more explicit and therefore di-

rectly usable for revision tools, like the one presented in (Haldimann, Beierle, and Kern-Isberner 2021)

The outline of this paper is as follows. In Section , we define basics of belief revision and fix our notation. Then, in Section , we briefly recall the basics of BR from (Rott 2012). In Section , we clarify the methodology of BR in the possible worlds reading and define a specific formula that subsumes the change implemented by BR and provides ground for a representation theorem. Employing this formula, we show that each BR corresponds to a lexicographic revision. Then, we transfer our results to ranking functions, first by a direct translation of the revision mechanism in Section and then via a conditional revision employing c-revisions in Section . We show that these implementations inherit all relevant features of BR and make the operation more easily accessible. In Section , we conclude. Most of the technical proofs of this paper are omitted due to lack of space and can be found in the technical appendix.

Formal and Semantic Basis of Belief Revision

In this section, after recalling some basics from propositional resp. conditional logic and fixing our notation, we present qualitative and semi-quantitative frameworks of representing beliefs for belief revision.

We denote by $\mathcal{L}$ a finitely generated, propositional language built over a signature Σ, equipped with the standard connectives *and* $\wedge$, *or* $\vee$ and *not* $\neg$. We omit the logical and-connector, writing $\alpha\beta$ instead of $\alpha \wedge \beta$, and overlining formulas indicates negation, i.e., $\overline{\alpha}$ means $\neg\alpha$. As usual, $\alpha \Rightarrow \beta$ is equivalent to $\overline{\alpha} \vee \beta$. We denote logical truths by $\top$ and contradictions by $\bot$. By Ω we denote the set of all propositional interpretations over Σ. We write $\omega \models \alpha$ when a world ω satisfies α, i.e., when ω is a model of α. By slight abuse of notation, we use ω both for the model and the conjunction of all corresponding literals. The set of all models of a formula α is denoted by $Mod(\alpha)$. For $\gamma, \delta \in \mathcal{L}$, it holds that $\gamma \models \delta$ iff $Mod(\gamma) \subseteq Mod(\delta)$. The deductively closed set $Th(W) = \{\alpha \in \mathcal{L} \,|\, \omega \models \alpha \text{ for all } \omega \in W\}$ which has exactly a subset $W \subseteq \Omega$ as models is called the *formal theory* of W. Let $(\mathcal{L}|\mathcal{L}) = \{(\beta|\alpha) \mid \alpha, \beta \in \mathcal{L}\}$ be the *conditional language* based on $\mathcal{L}$. The conditional $(\beta|\alpha)$ expresses *'If α, then (plausibly) β'*, and α is called the antecedent and β its consequent. Moreover, $(\beta|\alpha)$ is a three-valued logical entity (De Finetti 1975), with verification $\alpha\beta$, falsification $\alpha\overline{\beta}$ and neutrality $\overline{\alpha}$.

In 1997, Darwiche and Pearl introduced postulates called the DP-postulates for iterated belief revision. The semantic constraints following from the DP-postulates are defined for epistemic states Ψ equipped with total preorders on possible worlds, s.t. each epistemic state Ψ has an associated, deductively closed belief set $Bel(\Psi)$. There exists a wide variety of qualitative and semi-quantitative frameworks representing epistemic states each associated with different belief revision operators. Epistemic entrenchment relations represent an inner ordering of an agent's belief set, i.e., for two sentences $\alpha, \beta \in \mathcal{L}$, the notation $\alpha \leqslant_E \beta$ stands for *'β is at least as epistemically entrenched as α'*. Following (Nayak 1994), we define an epistemic entrenchment relation $\leqslant_E$ as a total preorder over $\mathcal{L}$, that satisfies (E1) If $\alpha \leqslant_E \beta$ and $\beta \leqslant_E \gamma$, then $\alpha \leqslant_E \gamma$, (E2) If $\alpha \models \beta$, then $\alpha \leqslant_E \beta$, (E3) $\alpha \leqslant_E \alpha \wedge \beta$ or $\beta \leqslant_E \alpha \wedge \beta$ and (E4) If $\alpha \leqslant_E \beta$ for all $\alpha \in \mathcal{L}$, then $\beta \equiv \top$. Then $\leqslant_E$ extracts a belief set $Bel(\leqslant_E) = \{\alpha \in \mathcal{L} \,|\, \bot <_E \alpha\}$ if $\bot <_E \alpha$ for some formula α, otherwise $Bel(\leqslant_E) = \mathcal{L}$. In contrast to epistemic entrenchment relations, total preorders (TPOs) over possible worlds rank worlds according to their closeness to a belief set, and therefore define an implausibility ordering on possible worlds, s.t. for two worlds $\omega, \omega' \in \Omega$, $\omega \preceq \omega'$ means that ω is at least as plausible as ω'. We denote for $W \subseteq \Omega$ by $\min(W, \preceq) = \{\omega \in \Omega \mid \omega \in W \text{ and } \omega \preceq \omega' \text{ for all } \omega' \in \Omega\}$ the set of minimal worlds in W w.r.t. $\preceq$. The belief set of $\preceq$ is defined via minimal worlds in Ω, s.t. $Bel(\preceq) = Th(\min(\Omega, \preceq))$ and a formula α is accepted by a TPO, $\preceq \models \alpha$, if α is part of its belief set, $\alpha \in Bel(\preceq)$. We call such TPOs over possible worlds *plausibilistic TPOs*. Each plausibilistic TPO on possible worlds $\preceq$ induces a plausibilistic relation on formulas via: $\alpha \preceq \beta$ iff $\min(Mod(\alpha), \preceq) \preceq \min(Mod(\beta), \preceq)$.

Plausibilistic TPOs and epistemic entrenchment relations are dual approaches to representing epistemic attitudes from which beliefs, but also an agent's preferences can be derived. Both formalisms are fundamental to belief revision (Katsuno and Mendelzon 1992; Gärdenfors and Makinson 1988). The following proposition summarizes the relationship between entrenchment relations and plausibility orderings.

Proposition 1 ((Peppas and Williams 1995)). *For each epistemic entrenchment $\leqslant_E$,*

$$\overline{\alpha} \preceq \overline{\beta} \text{ iff } \alpha \leqslant_E \beta \tag{1}$$

defines a plausibilistic TPO $\preceq$ s.t. $Bel(\leqslant_E) = Bel(\preceq)$ and vice versa.

Since each entrenchment relation $\leqslant_E$ is uniquely defined via the entrenchment classification of maximal disjunctions over Σ on which $\mathcal{L}$ is built, applying (1) on these entrenchment classifications leads immediately to a plausibilistic TPO $\preceq$ on possible worlds with the same belief set.

Now, we recall two commonly used iterated belief revision operators. Lexicographic revision displays an iterated belief revision operator that takes as input a sentence γ and is semantically defined on a plausibilistic TPOs as follows.

Definition 1 (Lexicographic Revision (Nayak, Pagnucco, and Peppas 2003)). *Let $\preceq$ be a plausibilistic TPO and $\gamma \in \mathcal{L}$. The lexicographic revision $\preceq *^{\ell} \gamma = \preceq^{\ell}_{\gamma}$ of $\preceq$ by γ satisfies*

(L1) *If $\omega, \omega' \models \gamma$, then $\omega \preceq^{\ell}_{\gamma} \omega'$ iff $\omega \preceq \omega'$*

(L2) *If $\omega, \omega' \not\models \gamma$, then $\omega \preceq^{\ell}_{\gamma} \omega'$ iff $\omega \preceq \omega'$*

(L3) *If $\omega \models \gamma$ and $\omega' \not\models \gamma$, then $\omega \prec^{\ell}_{\gamma} \omega'$*

C-revisions, introduced in 2001 by Kern-Isberner, display an iterated belief revision operator for conditional information. C-revisions employ the semi-quantitative framework of *ranking functions* resp. *ordinal conditional functions* (OCFs) (Spohn 1988) $\kappa : \Omega \rightarrow \mathbb{N}^{\infty}_0$, with $\kappa^{-1}(0) \neq \emptyset$, which substantiate plausibilistic TPOs via assigning to each world ω an implausibility rank $\kappa(\omega)$. The higher $\kappa(\omega)$ the less plausible ω is, and $\kappa^{-1}(0) \neq \emptyset$ ensures that worlds with maximal plausibility have rank zero. The belief set $Bel(\kappa) = Th(\kappa^{-1}(0))$ is defined via worlds with minimal ranks. It holds that $\kappa(\alpha) := \min\{\kappa(\omega) \mid \omega \models \alpha\}$ and $\kappa \models \alpha$ if $\kappa(\overline{\alpha}) > 0$ holds. A conditional $(\beta|\alpha)$ is accepted, $\kappa \models (\beta|\alpha)$, if $\kappa(\alpha\beta) < \kappa(\alpha\overline{\beta})$. It is easy to see that each ranking function corresponds to a plausibilistic TPO $\preceq_\kappa$ via the following equivalence

$$\omega \preceq_\kappa \omega' \text{ iff } \kappa(\omega) \leqslant \kappa(\omega') \text{ for all } \omega, \omega' \in \Omega. \tag{2}$$

C-revisions with a single conditional are defined as follows:

Definition 2 (C-revision (Kern-Isberner 2001))**.** *Let κ be an OCF and $(\delta|\gamma) \in (\mathcal{L}|\mathcal{L})$. The c-revision $\kappa^c = \kappa *^c (\delta|\gamma)$ is defined as follows:*

$$\kappa^c(\omega) = -\kappa(\overline{\delta} \vee \gamma) + \kappa(\omega) + \begin{cases} \nu^-, & \text{for } \omega \models \gamma\overline{\delta} \\ 0 & \text{else} \end{cases}, \tag{3}$$

s.t. the non-negative impact factor ν^- satisfies

$$\nu^- > \kappa(\gamma\delta) - \kappa(\gamma\overline{\delta}). \tag{4}$$

$-\kappa(\overline{\delta} \vee \gamma)$ ensures that κ^c is a well-defined OCF, s.t. $\kappa^c(0)^{-1} \neq \emptyset$ (Kern-Isberner 2004). The impact factor ν^- is defined via the inequality in (4) and it ensures that $\kappa^c \models (\delta|\gamma)$. Since we do not have interactions with other conditionals to be adopted, it is straightforward to single out a unique, non-negative impact factor

$$\nu^- = \max\{0, \kappa(\gamma\delta) - \kappa(\gamma\overline{\delta}) + 1\}. \tag{5}$$

A c-revision with ν^- as defined in (5) is called the *minimal c-revision with* $(\gamma|\delta)$. Since c-revisions rely on numerical ranks given by OCFs they are predestined for implementation. A high-level implementation using Abstract State Machines (Börger and Stärk 2003) and providing the core functionality of a series of belief change operations is presented in (Beierle, Kern-Isberner, and Koch 2008; Beierle and Kern-Isberner 2008). The Java-based implementation of belief change operations given in (Haldimann, Beierle, and Kern-Isberner 2021) also provides an implementation of c-revision because c-revision can be viewed as a special case of conditional descriptor revision (Hansson 2014).

Basics of Bounded Revision

In this section, we state the basic methodology and properties of BR in the framework of epistemic entrenchment relations as it was presented in (Rott 2012). BR displays an iterated belief revision mechanism that takes as input not only a classical input sentence β but also a reference sentence α. The reference sentence α acts as supplementary information for the revision and marks the depth with which the new belief is to be anchored in the posterior entrenchment relation.

Definition 3 ((Rott 2012))**.** *Let $\leqslant_E$ be an entrenchment relation and $\alpha, \beta \in \mathcal{L}$. The* Bounded Revision by β w.r.t. α of an entrenchment relation $\leqslant_E *_\alpha \beta = \leqslant_E^*$ *is defined as follows*

$\gamma \leqslant_E^* \delta$ *iff*

$$\begin{cases} (\beta \Rightarrow \gamma) \leqslant_E (\beta \Rightarrow \delta), & \text{if } \beta \Rightarrow (\gamma \wedge \delta) \leqslant_E (\beta \Rightarrow \alpha) \\ \gamma \leqslant_E \delta, & \text{otherwise} \end{cases} \tag{6}$$

for any arbitrary sentences $\gamma, \delta \in \mathcal{L}$.

The informal goal of BR can be partly formalized as the following success condition for BR:

(BR)$_E$ β is strictly more entrenched than α: $\alpha <_E^* \beta$

Note that, from Definition 3 and $(\text{BR})_E$ it remains unclear how much more plausible the input β shall be in the posterior ordering, i.e., what 'Accept β as far as α and *just a little further*' Furthermore, it holds that $\leqslant_E^*$ is a well-defined entrenchment relation and the BR operator $*_\alpha\beta$ satisfies the DP-postulates for iterated belief revision with β regardless of the choice of α (Rott 2012). Also, it holds that the resulting belief set $Bel(\leqslant_E^*)$ is insensitive to the choice of the reference sentence α Rott calls this property the *Same Beliefs Condition (SBC)$_E$*.

(SBC)$_E$ $Bel(\leqslant_E *_\alpha \beta) = Bel(\leqslant_E *_\gamma \beta)$ for any $\alpha, \gamma \in \mathcal{L}$

BR implements the idea that α does not define whether β will be accepted or not, but only how firmly it will be entrenched in the posterior state. Note that, since BR displays above all an iterated belief revision with β, the standard success condition $\beta \in Bel(\leqslant_E^*)$ holds.

Bounded Revision for Plausibilistic TPOs via Lexicographic Revision

The constraints in (6), defining BR for epistemic entrenchment relations, can be transferred to constraints for plausibilistic TPOs via (1). For γ and δ, we take maximal disjunctions, then their negations $\overline{\gamma}$ and $\overline{\delta}$ are maximal conjunctions, which correspond to possible worlds ω and ω' in the set Ω.

Definition 4. *Let $\preceq$ be a plausibilistic TPO and $\alpha, \beta \in \mathcal{L}$. The* Bounded Revision by β w.r.t. α of a plausibilistic TPO, $\preceq_{\alpha,\beta}^* = \preceq *_\alpha \beta$*, is defined as follows*

$$\omega \preceq_{\alpha,\beta}^* \omega' \text{ iff } \begin{cases} \beta\omega \preceq \beta\omega', & \text{if } \beta \wedge (\omega \vee \omega') \preceq \overline{\alpha}\beta \\ \omega \preceq \omega', & \text{otherwise} \end{cases} \tag{7}$$

Again, via applying (1), we transfer the success condition for BR $(\text{BR})_E$ and the Same Beliefs Condition $(\text{SBC})_E$ to the framework of plausibilistic TPOs.

(BR)$_\preceq$ $\overline{\alpha}$ is strictly more plausible than $\overline{\beta}$: $\overline{\alpha} \prec_{\alpha,\beta}^* \overline{\beta}$

(SBC)$_\preceq$ $Bel(\preceq *_\alpha \beta) = Bel(\preceq *_\gamma \beta)$ for any $\alpha, \gamma \in \mathcal{L}$

Since (7) is an equivalent reformulation of Definition 3, it is obvious that $(\text{BR})_\preceq$ and $(\text{SBC})_\preceq$ hold for each BR revised TPO $\preceq_{\alpha,\beta}^*$ defined by (7). Also, $\preceq_{\alpha,\beta}^* \models \beta$ holds, i.e., $\preceq_{\alpha,\beta}^*$ displays a revision with β.

The following proposition fully integrates BR by β w.r.t. α in the possible worlds reading.

Proposition 2. *For BR $\preceq_{\alpha,\beta}^*$ by β w.r.t. α, it holds that* (7) *is equivalent to the following constraints:*

$$\omega \preceq_{\alpha,\beta}^* \omega' \text{ iff } \begin{cases} \omega \preceq \omega', \text{if } \big(\omega, \omega' \models \beta \text{ and } \omega, \omega' \preceq \overline{\alpha}\beta\big) \\ \qquad \text{or } \big(\omega, \omega' \models \overline{\beta}\big) \\ \qquad \text{or } \big(\omega, \omega' \models \beta \text{ and } \overline{\alpha}\beta \prec \omega, \omega'\big) \\ \top, \qquad \text{if } \omega \models \beta, \omega \preceq \overline{\alpha}\beta \text{ and} \\ \qquad \big(\omega' \models \overline{\beta} \text{ or } \omega' \models \beta, \overline{\alpha}\beta \prec \omega'\big) \end{cases} \tag{8}$$

In (8), we employ exclusive cases which lead to more comprehensible constraints for BR. The first case in (8) implies that the relations among β-worlds, that are more plausible than $\bar{\alpha}\beta$ are kept and also the relations among worlds in the remaining set of Ω are kept. BR by β w.r.t. α affects the relations in between these exclusive sets and promotes β-worlds that are more plausible than $\bar{\alpha}\beta$. In a BR revised state these worlds are *strictly* more plausible than the rest of Ω which follows from the second case in (8). The strictness of the relation follows from the fact that we cannot swap ω and ω' in the inequality. Thus, we can derive from (8) that worlds $\omega \models \beta$, s.t. $\omega \preceq \bar{\alpha}\beta$ are at the center of the iterated belief change implemented by BR by β w.r.t. α. We summarize them by a unique formula.

Definition 5 (Core of Bounded Revision)**.** *For a plausibilistic TPO $\preceq$ and a BR operator $*_\alpha\beta$, we define the* core of Bounded Revision by β w.r.t. α *as the formula*

$$\varphi_{\alpha,\beta} = \beta \wedge (\bigvee_{\omega \preceq \bar{\alpha}\beta} \omega).$$

We call the set of possible worlds satisfying $\varphi_{\alpha,\beta}$ the core set of Bounded Revision *and notate* $\mathcal{C}_{\alpha,\beta} = Mod(\varphi_{\alpha,\beta}) = \{\omega \in \Omega \,|\, \omega \in Mod(\beta), \omega \preceq \bar{\alpha}\beta\}$.

The formula $\varphi_{\alpha,\beta}$ specifies via the reference sentence α, which β-worlds are *sufficiently* plausible to be promoted by BR by β w.r.t. α, i.e., to which extent the plausibility of β shall be increased relatively to the remaining worlds depending on the plausibility of $\bar{\alpha}$. Thus, the idea to "*accept the input sentence as far as the reference sentence and just a little further*", expressed in (Rott 2012), is substantiated by the formula $\varphi_{\alpha,\beta}$ which places it at the core of BR. We state the following representation theorem, which proves that $\varphi_{\alpha,\beta}$ displays the right choice to characterize the change mechanism of BR by β w.r.t. α.

Theorem 1 (Representation Theorem for BR)**.** *Let $*_\alpha\beta$ be a BR operator by β w.r.t. α. Let $\preceq$ be a plausibilistic TPO and $\preceq *_\alpha\beta = \preceq^*_{\alpha,\beta}$ be the corresponding BR revised plausibilistic TPO. Then $\preceq$ and $\preceq^*_{\alpha,\beta}$ satisfy* (8) *iff $\preceq$ and $\preceq^*_{\alpha,\beta}$ satisfy:*

(BR1) *If $\omega, \omega' \models \varphi_{\alpha,\beta}$, then $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$*

(BR2) *If $\omega, \omega' \not\models \varphi_{\alpha,\beta}$, then $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$*

(BR3) *If $\omega \models \varphi_{\alpha,\beta}$ and $\omega' \not\models \varphi_{\alpha,\beta}$ then $\omega \prec^*_{\alpha,\beta} \omega'$*

Proof. $\underline{\Rightarrow}$: Assume that $\preceq$ and $\preceq^*_{\alpha,\beta}$ satisfy (8).

(BR1): If $\omega, \omega' \models \varphi_{\alpha,\beta}$, then $\omega, \omega \models \beta$ and $\omega, \omega' \preceq \bar{\alpha}\beta$ and therefore the first alternative in the first case of (8) applies, i.e., $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$, and (BR1) holds.

(BR2): If $\omega, \omega' \not\models \varphi_{\alpha,\beta}$, then $\omega, \omega \not\models \beta$ or $\bar{\alpha}\beta \prec \omega, \omega'$. Note that the two alternatives are not disjunct, since some of the $\bar{\beta}$-worlds are less plausible than $\bar{\alpha}\beta$. We exclude these doubly named worlds and obtain the following exclusive formulation: $\omega, \omega' \models \bar{\beta}$ or ($\omega, \omega' \models \beta$ and $\bar{\alpha}\beta \prec \omega, \omega'$). This corresponds to the second or third alternative in the first case of (8). Thus, $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$, and (BR2) holds.

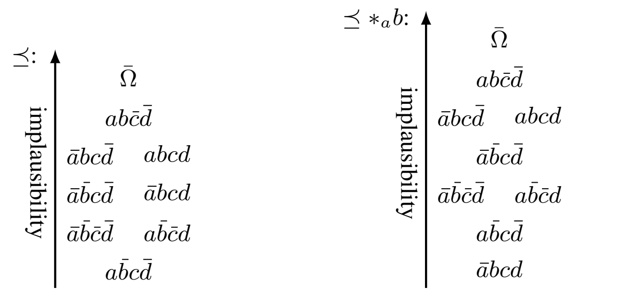

(a) Plausibilistic TPO $\preceq$ (b) BR revised TPO $\preceq *_a b$

Figure 1: BR by a w.r.t. b

(BR3): If $\omega \models \varphi_{\alpha,\beta}$ and $\omega' \not\models \varphi_{\alpha,\beta}$, then it holds that $\omega \models \beta$ and $\omega \preceq \bar{\alpha}\beta$, and for ω', it holds that either $\omega' \models \bar{\beta}$ or $\bar{\alpha}\beta \prec \omega'$. As for (BR2), we exclude the doubly named worlds in $\bar{\beta}$ and assume that $\omega' \models \bar{\beta}$ or ($\omega' \models \beta$ and $\bar{\alpha}\beta \prec \omega'$). Then the conditions of the second case of (8) are satisfied and, we can conclude that $\omega \preceq^*_{\alpha,\beta} \omega'$ but not $\omega' \preceq^*_{\alpha,\beta} \omega$, i.e. (BR3) holds.

$\underline{\Leftarrow}$: Assume that $\preceq$ and $\preceq^*_{\alpha,\beta}$ satisfy (BR1) – (BR3).

1. Assume that $\omega, \omega' \models \beta$ and $\omega, \omega' \preceq \bar{\alpha}\beta$, i.e. the first condition in the first case of (8) holds. Then $\omega, \omega' \models \varphi_{\alpha,\beta}$, i.e. we can conclude from (BR1), that $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$ holds.

2. Assume that $\omega, \omega' \models \bar{\beta}$, i.e. the second condition in the first case of (8) holds. Then $\omega, \omega' \not\models \varphi_{\alpha,\beta}$, i.e. we can conclude from (BR2), that $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$ holds.

3. Assume that $\omega, \omega' \models \beta$ and $\bar{\alpha}\beta \prec \omega, \omega'$, i.e. the third condition in the first case of (8) holds. Then $\omega, \omega' \not\models \varphi_{\alpha,\beta}$, i.e. we can conclude from (BR2), that $\omega \preceq \omega'$ iff $\omega \preceq^*_{\alpha,\beta} \omega'$ holds.

4. Assume that $\omega \models \beta, \omega \preceq \bar{\alpha}\beta$ and $\omega' \models \bar{\beta}$ or $\omega' \models \beta, \bar{\alpha}\beta \prec \omega'$, i.e. the second case of (8) applies. Then $\omega \models \varphi_{\alpha,\beta}$ and $\omega' \not\models \varphi_{\alpha,\beta}$ and we can conclude from (BR3) that $\omega \prec^*_{\alpha,\beta} \omega'$ holds, which implies that $\omega \preceq^*_{\alpha,\beta} \omega'$ holds.

□

The following proposition summarizes the characteristics of BR defined via (BR1) – (BR3).

Proposition 3. *For a plausibilistic TPO $\preceq$ and BR operator by β w.r.t. α, s.t. (BR1) – (BR3) hold for $\preceq *_\alpha\beta = \preceq^*_{\alpha,\beta}$, the following statements hold:*

1. $\beta \in Bel(\preceq^*_{\alpha,\beta}) = Th(\min(Mod(\beta), \preceq))$
2. $\bar{\alpha} \prec^*_{\alpha,\beta} \bar{\beta}$*, i.e., (BR)$_\preceq$ holds*
3. $Bel(\preceq *_\alpha\beta) = Bel(\preceq *_\gamma\beta)$ *for* $\gamma \in \mathcal{L}$*, so, (SBC)$_\preceq$ holds*

In the following example, we illustrate the versatile character of BR which depends on the interplay of input and reference sentence.

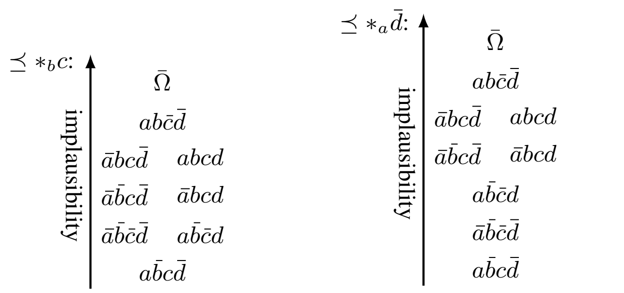

(a) BR revised TPO $\preceq *_b c$ (b) BR revised TPO $\preceq *_a \bar{d}$

Figure 2: BR by c resp. $\bar{d}$ w.r.t. b resp. a

Example 1. *In Figure 1a) a plausibilistic TPO $\preceq$ with signature $\Sigma = \{a, b, c, d\}$ is given, where $\bar{\Omega}$ denotes all remaining worlds on the same level of plausibility, which are not shown explicitly. We perform three different BR operations in Figures 1b), 2a) and 2b) to illustrate the strength and special features of BR. In Figure 1b) the BR $\preceq *_a b$ with input $\beta = b$ and reference sentence $\alpha = a$ is depicted. It holds that $\mathcal{C}_{a,b} = \{\bar{a}bcd\}$ and obviously $\preceq^*_{a,b} \models b$ and (BR)$_\preceq$ hold. Note that, BR yields the same belief set of the posterior TPO for a different reference sentence, like, e.g. $\alpha = \bar{a}$ due to (SBC)$_\preceq$. In Figure 2a), the outcome of $\preceq *_b c$ with reference sentence $\alpha = b$ and input $\beta = c$ is illustrated. BR does not change the prior ordering since (BR1) – (BR3) are already satisfied. This example shows that the change implemented by BR is vacuous under the condition that all worlds in $\mathcal{C}_{\alpha,\beta}$ are already more plausible than worlds outside of it in the prior ordering. Otherwise, (BR1) – (BR3) imply the strengthening of the β-belief as can be seen in Figure 2b), where BR is performed with input $\beta = \bar{d}$ and reference sentence $\alpha = a$. Here, the input sentence $\bar{d}$ gets promoted via BR, even though it is already believed in the prior ordering. Note that this promotion of the input belief depends on the choice of the reference sentence, because the reference sentence specifies how much more plausible the input shall be in the posterior ordering.*

In general, for BR by β w.r.t. α the reference sentence is used to specify the distance between β- and $\bar{\beta}$-worlds, since it is crucial for $\varphi_{\alpha,\beta}$ and thus defines how much more plausible the input β shall be in the posterior ordering $\preceq^*_{\alpha,\beta}$. In general, it holds that $\bar{\alpha}\beta \approx^*_{\alpha,\beta} \bar{\alpha} \prec^*_{\alpha,\beta} \bar{\beta}$ due to the definition of $\varphi_{\alpha,\beta}$ and (BR)$_\preceq$. Yet, BR does not guarantee that (BR)$_\preceq$ is satisfied in the slightest possible way, i.e., it is possible that there exists a world ω s.t. $\bar{\alpha} \prec^*_{\alpha,\beta} \omega \prec^*_{\alpha,\beta} \bar{\beta}$ as the following example shows.

Example 2. *Let $\bar{a}b \prec ab \prec a\bar{b} \prec \bar{a}\bar{b}$ be a plausibilistic TPO over $\Sigma = \{a, b\}$. For BR with reference $\alpha = a$ and input sentence $\beta = b$, it holds that $\mathcal{C}_{a,b} = \{\bar{a}b\}$ and thus, the constraints (BR1) – (BR3) hold and BR does not change the prior ordering. Note that, (BR)$_\preceq$ is satisfied, but for the posterior ordering it holds that $\bar{a} \prec^*_{\alpha,\beta} ab \prec^*_{\alpha,\beta} \bar{b}$*

In Theorem 1 constraints (BR1) – (BR3) are given which characterize BR by β w.r.t. α. If we compare these constraints to (L1) – (L3) from Definition 1, i.e., the semantic constraints defining a lexicographic revision with a statement γ, we notice that the revision mechanisms are related to each other. The following theorem proves that BR by β w.r.t. α displays a lexicographic revision with the corresponding core of BR.

Theorem 2. *Let $\preceq$ be a plausibilistic TPO, $*_\alpha\beta$ be a BR operator by β w.r.t. α and $*^\ell$ be a lexicographic revision operator. For $\preceq *_\alpha\beta$ and $\preceq *^\ell\varphi_{\alpha,\beta}$ with $\varphi_{\alpha,\beta}$ being the core of BR by β w.r.t. α, it holds that*

$$\omega\ (\preceq *_\alpha\beta)\ \omega' \text{ iff } \omega\ (\preceq *^\ell\varphi_{\alpha,\beta})\ \omega'.$$

Theorem 2 states that BR by β w.r.t. α corresponds to a lexicographic revision by $\varphi_{\alpha,\beta}$. Thus, we encode the supplementary information given in the reference sentence α from the meta-level to the more clearly defined and directly usable object level via $\varphi_{\alpha,\beta}$. Hence, lexicographic revision with $\varphi_{\alpha,\beta}$ can be seen as a reduction of BR by β w.r.t. α to the framework of lexicographic revisions, which makes it directly usable for lexicographic revision solvers (see for e.g. (Amor et al. 2018)). Moreover, in contrast to a standard lexicographic revision with β, where (L3) forces quite rough changes on the prior ordering by making all β-worlds more plausible than $\bar{\beta}$-worlds, BR implements a more fine-grained revision with input β. The incorporation of α leads to the corresponding revision with $\varphi_{\alpha,\beta}$, where α marks to which plausibility level worlds in $Mod(\beta)$ are promoted in the posterior TPO. This supports the idea that α serves as an indicator for the reliability of β.

Realizing Bounded Revision for OCFs

In this section, we present a realization of BR in the framework of ranking functions via a straightforward implementation of BR by β w.r.t. α for ranking functions $\kappa *_\alpha \beta$. And, we show that the corresponding plausibilistic TPO satisfies the representation theorem stated in the previous section.

The following formulation of BR for ranking functions makes the mechanism of BR more explicit.

Definition 6. *Let κ be a ranking function. We define* Bounded Revision by β w.r.t. α for ranking functions $\kappa *_\alpha \beta = \kappa^*_{\alpha,\beta}$ *as follows:*

$$\kappa^*_{\alpha,\beta}(\omega) = \kappa_0 + \begin{cases} \kappa(\omega) + \kappa(\bar{\alpha}\beta) + 1, & \omega \not\models \varphi_{\alpha,\beta} \\ \kappa(\omega), & \textit{otherwise} \end{cases} \tag{9}$$

with $\kappa_0 = -\kappa(\beta)$ as a normalization constant.

In general, for the normalization constant κ_0 in (9) it holds that $\kappa_0 = -\min\{\min_{\omega\models\varphi_{\alpha,\beta}}\{\kappa(\omega)\}, \min_{\omega'\not\models\varphi_{\alpha,\beta}}\{\kappa(\omega') + \kappa(\bar{\alpha}\beta) + 1\}\}$. Since for $\omega \models \varphi_{\alpha,\beta}$ it holds that $\kappa(\omega) \leqslant \kappa(\bar{\alpha}\beta)$ and $\omega \models \beta$. We get that $\kappa_0 = -\kappa(\varphi_{\alpha,\beta}) = -\kappa(\beta)$, since all minimal worlds satisfying $\varphi_{\alpha,\beta}$ are also minimal worlds satisfying β and thus, $\kappa^*_{\alpha,\beta}$ is well-defined. The following theorem shows that $\kappa^*_{\alpha,\beta}$ satisfies (BR1) – (BR3) and therefore displays a variant of a simple, yet elegant implementation of BR for ranking functions.

Theorem 3. *Let κ be a ranking function and $\kappa *_\alpha \beta = \kappa^*_{\alpha,\beta}$ as defined in (9). Then (BR1) – (BR3) hold for the corresponding plausibilistic TPOs $\preceq_\kappa$ and $\preceq_{\kappa^*_{\alpha,\beta}}$.*

Note that, BR by β w.r.t. α for ranking functions as defined in (9) is a variant of a realization of BR for OCFs. Due to the use of numerical ranks, other realizations are possible. Like, e.g., variants that implement a greater stretch between worlds $\omega \models \varphi_{\alpha,\beta}$ and $\omega' \not\models \varphi_{\alpha,\beta}$ by adding a constant number higher than 1 to worlds $\omega \not\models \varphi_{\alpha,\beta}$ in (9). Yet, all variants correspond to the same TPO, since the addition of a constant number greater than 1 solely adds empty layers to $\kappa^*_{\alpha,\beta}$.

From Proposition 3 and Theorem 3 it follows that, $\kappa^*_{\alpha,\beta}(\overline{\alpha}) < \kappa^*_{\alpha,\beta}(\overline{\beta})$ and $Bel(\kappa *_\alpha \beta) = Bel(\kappa *_\gamma \beta)$ for $\alpha, \gamma \in \mathcal{L}$ hold, i.e., (BR)$_\preceq$ and (SBC)$_\preceq$ reformulated via (2) for ranking functions are satisfied.

Bounded Revision via Conditional C-revision

In this section, we realize BR for OCFs from (9) via a conditional revision of ranking functions with a single conditional. This makes BR directly usable for existing frameworks of belief revision which are capable to revise with conditional information, such as the one presented in (Haldimann, Beierle, and Kern-Isberner 2021).

We have seen so far, that BR yields a posterior ordering by shifting worlds outside the core set of BR to plausibility levels strictly higher than the ones where worlds inside the core set reside. Now, we will show that this shift can be expressed by the following conditional which uses the core formula $\varphi_{\alpha,\beta}$ corresponding to a BR by β w.r.t. α:

$$\Phi_{\alpha,\beta} = (\varphi_{\alpha,\beta} | \overline{\varphi_{\alpha,\beta}} \vee (\overline{\alpha}\beta))$$

We call $\Phi_{\alpha,\beta}$ the *core conditional of BR by β w.r.t. α* and yield the corresponding verification resp. falsification formulas as follows

$$(\overline{\varphi_{\alpha,\beta}} \vee (\overline{\alpha}\beta)) \wedge (\varphi_{\alpha,\beta}) \equiv \varphi_{\alpha,\beta} \wedge \overline{\alpha}\beta \equiv \overline{\alpha}\beta \wedge (\bigvee_{\omega \preceq \overline{\alpha}\beta} \omega) \tag{10}$$

$$(\overline{\varphi_{\alpha,\beta}} \vee (\overline{\alpha}\beta)) \wedge (\overline{\varphi_{\alpha,\beta}}) \equiv \overline{\varphi_{\alpha,\beta}} \tag{11}$$

Equation (10) corresponds to the verification of $\Phi_{\alpha,\beta}$ and (11) to its falsification.

The minimal c-revision $\kappa *^c \Phi_{\alpha,\beta} = \kappa^c_{\alpha,\beta}$ according to Definition 2 is obtained by employing (10) and (11) in the definition of the minimal, non-negative impact factors in (5). Note that, since the plausibility of $\overline{\alpha}\beta$ is defined via minimal models, it holds that only these minimal models satisfy $\overline{\alpha}\beta$ and $\omega \preceq \overline{\alpha}\beta$ in (10) at the same time. And therefore the OCF ranking of (10) equals $\kappa(\overline{\alpha}\beta)$. We get the following compact definition of the impact factor for $\Phi_{\alpha,\beta}$:

$$\nu^-_\Phi = \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi_{\alpha,\beta}}) + 1\} \tag{12}$$

And we get the following c-revision

$$\begin{aligned}\kappa^c_{\alpha,\beta} &= -\kappa((\varphi_{\alpha,\beta} \wedge (\alpha \Rightarrow \beta)) \vee \varphi_{\alpha,\beta}) + \kappa(\omega) \\ &\quad + \begin{cases} \nu^-_\Phi, & \text{for } \omega \models \overline{\varphi_{\alpha,\beta}} \\ 0, & \text{else} \end{cases} \\ &= -\kappa(\beta) + \kappa(\omega) + \begin{cases} \nu^-_\Phi, & \text{for } \omega \not\models \varphi_{\alpha,\beta} \\ 0, & \text{else} \end{cases}\end{aligned} \tag{13}$$

For the normalization constant it holds that $(\varphi_{\alpha,\beta} \wedge (\alpha \Rightarrow \beta)) \vee \varphi_{\alpha,\beta} \equiv \varphi_{\alpha,\beta}$. And, as before for (9), it holds for all minimal worlds satisfying $\varphi_{\alpha,\beta}$ that they are also minimal worlds satisfying β. Thus, we get that $-\kappa(\varphi_{\alpha,\beta}) = -\kappa(\beta)$ for the normalization constant in (13) and $\kappa *^c \Phi_{\alpha,\beta}$ is well-defined. The non-negative impact factor ν^-_Φ specifies the plausibility stretch between the input and the reference sentence in the revised state. This is crucial to reduce the meta-level revision with a supplementary information in (9) to the object level encoded as the c-revision with a single conditional, which captures all relevant features of BR as the following theorem shows.

Theorem 4. *Let κ be a ranking function. For the minimal c-revision $\kappa *^c \Phi_{\alpha,\beta} = \kappa^c_{\alpha,\beta}$ as defined in* (13) *and $\kappa *_\alpha \beta = \kappa^*_{\alpha,\beta}$ from* (9)*, it holds that the corresponding plausibilistic TPOs $\preceq_{\kappa^c_{\alpha,\beta}}$ and $\preceq_{\kappa^*_{\alpha,\beta}}$ are the same, i.e.,*

$$\omega \preceq_{\kappa^c_{\alpha,\beta}} \omega' \text{ iff } \omega \preceq_{\kappa^*_{\alpha,\beta}} \omega'.$$

Proof. We show that TPO $\preceq_{\kappa^c_{\alpha,\beta}}$ corresponding to the c-revision $\kappa^c_{\alpha,\beta}$ satisfies (BR1) – (BR3) from Theorem 1. Since (BR1) – (BR3) define a unique TPO and $\preceq_{\kappa^*_{\alpha,\beta}}$, i.e., the TPO corresponding to $\kappa^*_{\alpha,\beta}$, also satisfies (BR1) – (BR3) as we have shown in Theorem 3, we can immediately conclude that $\preceq_{\kappa^c_{\alpha,\beta}} = \preceq_{\kappa^*_{\alpha,\beta}}$ holds.

- **(BR1):** Let $\omega, \omega' \models \varphi_{\alpha,\beta}$. Then both worlds do not falsify $\Phi_{\alpha,\beta}$, thus their ranks do not change during the c-revision with $\Phi_{\alpha,\beta}$. We get that $\kappa^c_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) \leqslant \kappa_0 + \kappa(\omega') = \kappa^c_{\alpha,\beta}(\omega')$ iff $\kappa(\omega) \leqslant \kappa(\omega')$ and therefore, (BR1) follows for $\preceq_\kappa$ resp. $\preceq_{\kappa^*_{\alpha,\beta}}$ via (2).
- **(BR2):** Let $\omega, \omega' \not\models \varphi_{\alpha,\beta}$. Then $\omega, \omega' \models \overline{\varphi_{\alpha,\beta}}$ holds and both worlds falsify $\Phi_{\alpha,\beta}$. For the c-revision with $\Phi_{\alpha,\beta}$ and the corresponding impact factor ν^-_Φ as defined in (12), we get that $\kappa^c_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) + \nu^-_\Phi \leqslant \kappa_0 + \kappa(\omega') + \nu^-_\Phi = \kappa^c_{\alpha,\beta}(\omega')$ iff $\kappa(\omega) \leqslant \kappa(\omega')$, since ν^-_Φ is a constant. Thus, (BR2) follows for $\preceq_\kappa$ resp. $\preceq_{\kappa^*_{\alpha,\beta}}$ via (2).
- **(BR3):** Let $\omega \models \varphi_{\alpha,\beta}$ and $\omega' \not\models \varphi_{\alpha,\beta}$, then it holds that ω does not falsify $\Phi_{\alpha,\beta}$, while ω' does and therefore, we get that $\kappa^c_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega)$ and $\kappa^c_{\alpha,\beta}(\omega') = \kappa_0 + \kappa(\omega') + \nu^-_\Phi$ holds, with $\nu^-_\Phi = \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi_{\alpha,\beta}}) + 1\}$ as in (12). The following statements hold for each ranking function κ and will be useful in the course of this proof.

$$\kappa(\overline{\varphi_{\alpha,\beta}}) = \kappa(\bigvee_{\tilde{\omega} \not\models \varphi_{\alpha,\beta}} \tilde{\omega}) = \min_{\tilde{\omega} \not\models \varphi_{\alpha,\beta}} \{\kappa(\tilde{\omega})\}. \tag{14}$$

$$\text{For } \omega \models \varphi_{\alpha,\beta}, \text{ it holds that } \kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta). \tag{15}$$

The first statement follows from $\overline{\varphi_{\alpha,\beta}} \equiv \bigvee_{\tilde{\omega} \not\models \varphi_{\alpha,\beta}} \tilde{\omega}$ and the properties of ranking functions. The second one holds since $\omega \in \mathcal{C}_{\alpha,\beta}$.

Now, we show (BR3) for the different forms of the impact factor ν^-_Φ.

1. Assume that $\nu^-_\Phi = 0$, then it holds that $\kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi_{\alpha,\beta}}) + 1 \leqslant 0$, s.t., $\kappa(\overline{\alpha}\beta) + 1 \leqslant \kappa(\overline{\varphi_{\alpha,\beta}})$ and therefore $(*)$ $\kappa(\overline{\alpha}\beta) < \kappa(\overline{\varphi_{\alpha,\beta}})$ holds. Together with (15), $(*)$

$\omega \in \Omega$	κ	$\kappa^*_{a,b}$	$\kappa^c_{a,b}$	$\kappa^*_{b,c}$	$\kappa^c_{b,c}$	$\kappa^*_{a,\bar{d}}$	$\kappa^*_{a,\bar{d}}$
$a\bar{b}c\bar{d}$	0	1	1	0	0	0	0
$a\bar{b}\bar{c}d$	1	2	2	2	1	3	2
$\bar{a}\bar{b}\bar{c}\bar{d}$	1	2	2	2	1	1	1
$\bar{a}bcd$	2	0	0	3	2	4	3
$\bar{a}\bar{b}c\bar{d}$	2	3	3	3	2	4	3
$abcd$	3	4	4	4	3	5	4
$\bar{a}bc\bar{d}$	3	4	4	4	3	5	4
$ab\bar{c}\bar{d}$	4	5	5	5	4	6	5
$\bar{\Omega}$	5	6	6	6	5	7	6

Table 1: Prior κ and the BR revised κ^* resp. c-revised $\kappa^c_{\alpha,\beta}$.

and (14), we get that $\kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta) < \kappa(\overline{\varphi_{\alpha,\beta}}) \leqslant \kappa(\omega')$ holds. Thus, $\kappa^c_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) < \kappa_0 + \kappa(\omega') = \kappa^c_{\alpha,\beta}(\omega')$ holds.
2. Assume that $\nu^-_\Phi = \kappa(\bar{\alpha}\beta) - \kappa(\overline{\varphi_{\alpha,\beta}}) + 1$. Then, it follows from (14) that $(**)\ 0 \leqslant \kappa(\omega') - \kappa(\overline{\varphi_{\alpha,\beta}}) < \kappa(\omega') - \kappa(\overline{\varphi_{\alpha,\beta}}) + 1$ holds. Together with (15) and $(**)$, it holds that $\kappa^c_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) \leqslant \kappa_0 + \kappa(\overline{\alpha}\beta) < \kappa_0 + \kappa(\overline{\alpha}\beta) + \kappa(\omega') - \kappa(\overline{\varphi_{\alpha,\beta}}) + 1 = \kappa^c_{\alpha,\beta}(\omega')$. All in all, (BR3) follows for $\preceq_\kappa$ resp. $\preceq_{\kappa^*_{\alpha,\beta}}$ via (2).

□

The theorem shows that the core conditional encodes the specific decrease of plausibility for all worlds $\omega \not\models \varphi_{\alpha,\beta}$ in a single, easily accessible logical entity, rather than introducing a meta-level to the revision operator itself as it is the case for standard BR operators. Thus, we have shown that the supplementary information provided by the reference sentence in BR can be incorporated naturally to the object level of an existing revision operator. Since $\preceq_{\kappa^*_{\alpha,\beta}}$ satisfies (BR1) – (BR3), it follows immediately from Theorem 4 that $\preceq_{\kappa^c_{\alpha,\beta}}$ also satisfies (BR1) – (BR3).

Yet, the ranking functions $\kappa^*_{\alpha,\beta}$ from (9) and $\kappa^c_{\alpha,\beta}$ in (13) are not identical. This is due to special features of $\kappa^*_{\alpha,\beta}$ vs. $\kappa^c_{\alpha,\beta}$, which we illustrate via the following example. But first, we transfer the prior TPO $\preceq$ from Example 1 in Figure 1a) to a ranking function by assigning each layer of $\preceq$ a plausibility rank starting with zero for the lowermost layer.

Example 3. *In Table 1 the prior ranking function κ which corresponds to $\preceq$ from Example 1 is depicted. Alongside, with the two ranking functions $\kappa^*_{\alpha,\beta} = \kappa *_\alpha \beta$ resp. $\kappa^c_{\alpha,\beta} = \kappa * \Phi_{\alpha,\beta}$ for each corresponding BR.*

For the first BR by $\beta = b$ w.r.t. $\alpha = a$, it holds that $\kappa^*_{a,b}(\omega) = \kappa^c_{a,b}(\omega)$, i.e., BR for OCFs as defined in Definition 6 and the c-revision with the corresponding conditional yield the same result. For the next BR with reference $\alpha = b$ and input sentence $\beta = c$, it holds that $c \in Bel(\kappa)$ and $\mathcal{C}_{b,c} = \{a\bar{b}c\bar{d}\}$. So, the input is already accepted in the prior ranking function and (BR3) is satisfied. Here, BR for OCFs $\kappa^*_{b,c}$ strengthens the input c via introducing an empty layer, s.t., $\kappa^*_{b,c}(c) = 0 < 2 = \kappa^*_{b,c}(\bar{c})$, instead of $\kappa(c) = 0 < 1 = \kappa(\bar{c})$ as before. In contrast, c-revisions as in $\kappa^c_{b,c}$ employ minimal change since (BR3) is already satisfied, and therefore the prior ranking function is kept. For the next BR with $\alpha = a$ and $\beta = \bar{d}$, κ accepts the input $\bar{d}$, but (BR3) is not satisfied, since $\mathcal{C}_{a,\bar{d}} = \{a\bar{b}c\bar{d}, \bar{a}\bar{b}\bar{c}\bar{d}\}$ and $\kappa(\bar{a}\bar{b}\bar{c}\bar{d}) = \kappa(a\bar{b}\bar{c}d)$ holds. Again, the change implemented by $\kappa^c_{a,\bar{d}}$ is minimal, in the sense that the ranks of all worlds outside the core set are increased by exactly the stretch needed to satisfy (BR3), namely $\kappa(\bar{a}\bar{d}) - \kappa(\overline{\varphi}_{a,\bar{d}}) + 1 = 1$. BR for OCFs $\kappa^*_{a,\bar{d}}$ leads to a strengthening of the input $\bar{d}$, while satisfying (BR1) – (BR3). Note that, it holds for all BR revised and c-revised ranking functions, that their corresponding plausibilistic TPOs coincide with the BR revised TPOs from Figures 1b), 2a) and 2b).

The example shows that c-revisions do not change the prior ordering when it is not necessary. This corresponds to a minimal implementation of (BR1) and (BR2), in the sense that the specific ranks of worlds in and outside the core set of BR are kept as long as $\Phi_{\alpha,\beta}$ holds. On the other hand, BR for OCFs can also be used to strengthen the input beliefs by decreasing the plausibility level of worlds outside the core, even though the input is already accepted. Note that BR in general increases the number of layers in the posterior ordering, making the belief state more fine-grained.

Conclusion

BR displays a versatile iterated belief revision operator with dual input information β and a reference sentence α. A unique feature of BR is that α determines to which extent β is accepted by the agent and thus guides the revision process. In this paper, we have elaborated the intuitive strengths of BR and provided a representation theorem in Section , which clarifies the methodology of BR in an elegant way and clearly indicates its connection to lexicographic revision. We have shown that BR by β w.r.t. α can be implemented as a lexicographic revision by the corresponding core formula $\varphi_{\alpha,\beta}$, which reduces the supplementary information in α from the meta-level to the object level. The translations into the framework of ranking functions paves the way to the implementation of interesting operators capable of dealing with two-dimensional belief change operators and it integrates BR into an existing framework of conditional revision operators. Now, BR can be implemented as a conditional revision by the core conditional $\Phi_{\alpha,\beta}$ for OCFs.

Summarizing, we have made BR fully usable and implementable as an iterated revision operator for TPOs and OCFs, elaborating in these frameworks clearly its underlying methodology. This provides immediately revision operators that allow for taking supplementary information in the form of a reference sentence also for TPOs and OCFs. As part of our ongoing and future work, we will use the insights into BR gained from this paper for elaborating its relations to other iterated revision operators such as natural revision (Boutilier 1996) more deeply.

References

Ammar, Y.; and Ismail, H. O. 2021. Trust Is All You Need: From Belief Revision to Information Revision. In Faber, W.; Friedrich, G.; Gebser, M.; and Morak, M., eds., *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021, Virtual Event, May 17-20, 2021, Proceedings*, volume 12678 of *Lecture Notes in Computer Science*, 50–65. Springer.

Amor, N. B.; Khalfi, Z. E.; Fargier, H.; and Sabbadin, R. 2018. Lexicographic refinements in stationary possibilistic Markov Decision Processes. *Int. J. Approx. Reason.*, 103: 343–363.

Beierle, C.; and Kern-Isberner, G. 2008. A Verified AsmL Implementation of Belief Revision. In Börger, E.; Butler, M.; Bowen, J. P.; and Boca, P., eds., *Abstract State Machines, B and Z, First International Conference, ABZ 2008, London, UK, September 16-18, 2008. Proceedings*, volume 5238 of *LNCS*, 98–111. Springer. ISBN 978-3-540-87602-1.

Beierle, C.; Kern-Isberner, G.; and Koch, N. 2008. A High-Level Implementation of a System for Automated Reasoning with Default Rules (System Description). In Armando, A.; Baumgartner, P.; and Dowek, G., eds., *Proc. of the 4th International Joint Conference on Automated Reasoning (IJCAR-2008)*, volume 5195 of *LNCS*, 147–153. Springer.

Börger, E.; and Stärk, R. F. 2003. *Abstract State Machines. A Method for High-Level System Design and Analysis*. Springer. ISBN 3540007024.

Boutilier, C. 1996. Iterated revision and minimal change of conditional beliefs. *J. Philos. Log.*, 25(3): 263–305.

Darwiche, A.; and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence*, 89: 1–29.

De Finetti, B. 1975. *Theory of Probability: A critical introductory treatment*. Wiley Series in Probability and Statistics. Wiley. ISBN 9781119286370.

Gärdenfors, P.; and Makinson, D. 1988. Revisions of Knowledge Systems Using Epistemic Entrenchment. In Vardi, M. Y., ed., *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, USA, March 1988*, 83–95. Morgan Kaufmann.

Haldimann, J.; Beierle, C.; and Kern-Isberner, G. 2021. Syntax Splitting for Iterated Contractions, Ignorations, and Revisions on Ranking Functions Using Selection Strategies. In Faber, W.; Friedrich, G.; Gebser, M.; and Morak, M., eds., *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021, Virtual Event, May 17-20, 2021, Proceedings*, volume 12678 of *LNCS*, 85–100. Springer.

Hansson, S. O. 2014. Descriptor Revision. *Stud Logica*, 102(5): 955–980.

Hunter, A. 2021. Building Trust for Belief Revision. In Pham, D. N.; Theeramunkong, T.; Governatori, G.; and Liu, F., eds., *PRICAI 2021: Trends in Artificial Intelligence - 18th Pacific Rim International Conference on Artificial Intelligence, PRICAI 2021, Hanoi, Vietnam, November 8-12, 2021, Proceedings, Part I*, volume 13031 of *Lecture Notes in Computer Science*, 543–555. Springer.

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

Kern-Isberner, G. 2001. *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*, volume 2087 of *Lecture Notes in Computer Science*. Berlin: Springer. ISBN 3-540-42367-2.

Kern-Isberner, G. 2004. A thorough axiomatization of a principle of conditional preservation in belief revision. *Annals of Mathematics and Artificial Intelligence*, 40(1-2): 127–164.

Nayak, A. C. 1994. Iterated Belief Change Based on Epistemic Entrenchment. *Erkenntnis*, 41(3): 353–390.

Nayak, A. C.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *Artif. Intell.*, 146(2): 193–228.

Peppas, P.; and Williams, M.-A. 1995. Constructive Modelings for Theory Change. *Notre Dame Journal of Formal Logic*, 36(1): 120 – 133.

Pornpitakpan, C. 2004. The Persuasiveness of Source Credibility: A Critical Review of Five Decades' Evidence. *Journal of Applied Social Psychology*, 34: 243 – 281.

Rott, H. 2012. Bounded Revision: Two-Dimensional Belief Change Between Conservative and Moderate Revision. *J. Philos. Log.*, 41(1): 173–200.

Sezgin, M.; and Kern-Isberner, G. 2022. Revision by Comparison for Ranking Functions. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, 2734–2740. ijcai.org.

Sparks, J. R.; and Rapp, D. N. 2011. Unreliable and Anomalous: How the Credibility of Data Affects Belief Revision. In Carlson, L. A.; Hölscher, C.; and Shipley, T. F., eds., *Proceedings of the 33th Annual Meeting of the Cognitive Science Society, CogSci 2011, Boston, Massachusetts, USA, July 20-23, 2011*. cognitivesciencesociety.org.

Spohn, W. 1988. Ordinal Conditional Functions: A Dynamic Theory of Epistemic States. In Harper, W. L.; and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, 105–134. Dordrecht: Springer. ISBN 978-94-009-2865-7.

Wolf, A. G.; Rieger, S.; and Knauff, M. 2012. The effects of source trustworthiness and inference type on human belief revision. *Thinking & Reasoning*, 18(4): 417–440.