Revision by Comparison for Ranking Functions

Meliha Sezgin* and **Gabriele Kern-Isberner**
TU Dortmund University, Germany
{meliha.sezgin, gabriele.kern-isberner}@tu-dortmund.de

Abstract

Revision by Comparison (RbC) is a non-prioritized belief revision mechanism on epistemic states that specifies constraints on the plausibility of an input sentence via a designated reference sentence, allowing for kind of relative belief revision. In this paper, we make the strategy underlying RbC more explicit and transfer the mechanism together with its intuitive strengths to a semi-quantitative framework based on ordinal conditional functions where a more elegant implementation of RbC is possible. We furthermore show that RbC can be realized as an iterated revision by so-called weak conditionals. Finally, we point out relations of RbC to credibility-limited belief revision, illustrating the versatility of RbC for advanced belief revision operations.

1 Introduction

The aim of belief change operators [Alchourrón *et al.*, 1985; Katsuno and Mendelzon, 1992] is to incorporate new pieces of information into an agent's belief set, which means they are intrinsically linked to the principle of primacy of update. Non-prioritized belief revision operators break with this goal and allow that new information can be rejected for several reasons (see [Hansson, 2008] for a survey). *Revision by Comparison* (RbC) was firstly introduced by Fermé and Rott in [2004] and represents a belief change mechanism that allows to accept a new *input sentence* β only to a certain degree, via specifying the strength of the new belief in the posterior belief state via a so-called *reference sentence* α. This flexible approach to belief revision results in a hybrid belief change operator between revision with an input information and the contraction of a reference sentence. RbC adapts the kind of belief change to the prior belief state of an agent and thus links the reliability resp. priority of the new information to a reference sentence which is either specified via the input or can be selected freely by an agent. Except for a weaker form of belief change operators recently proposed in [Schwind *et al.*, 2018], this dynamic form of belief change is unique to the methodology of RbC and it leads to many interesting connections to other forms of non-prioritized belief change. Among them are *credibility-limited revision operators* [Hansson *et al.*, 2001; Booth *et al.*, 2012], which perform a revision of the prior belief state solely if the new piece of information is credible enough, whereby credibility is defined via a fixed set of propositions whose determination remains unclear. Here, RbC offers the chance to a more flexible approach to credibility using a reference sentence for each new piece of information. We give a motivating example:

Example 1. *On a sunny afternoon, an agent watches a video on BBC news that states that scientists discovered a rare penguin type that has the ability to fly! The agent is astonished! Even though she knows that penguins are birds and birds normally fly, she is also aware of the fact that penguins are an exception to this rule and do not fly. So the new information is highly unplausible, but on the other hand the BBC is a trustworthy source of information which acts as reference.*

For RbC the reference sentence, e.g. the trustworthiness of the BBC news, can act as a a marker of reliability and allows us to revise with a seemingly unplausible new information only to a certain degree or even the devaluation of the reference. Despite the intuitive strengths of RbC the implementation remains unclear in [Fermé and Rott, 2004]. We show that this is not due to the intricateness of the revision mechanism per se but rather because the determination of the affected worlds is highly dependent on the relative positioning of input and reference beliefs. By changing the framework of belief representation and adapting the underlying change mechanism, we make the methodology and the ensuing application of RbC more explicit and concise and therefore directly usable for revisions tools. Our main contributions are the following:

- We clarify the strategy underlying RbC and present a more elegant implementation for total preorders in a semi-quantitative framework of belief representation.
- Using negative information in the form of so-called weak conditionals, we characterize RbC as a conditional revision.
- We discuss the hybrid belief change character of RbC and compare our results in the context of credibility-limited revision operators.

We start in Section 2 by stating some formal preliminaries,

*Contact Author

whereby we discuss the relation between two qualitative approaches to belief representation in more detail. After recapping the belief revision mechanism of RbC in Section 3.1, we clarify its methodology by expressing it via possible worlds. This provides the grounds for presenting a more elegant implementation of RbC via ranking functions [Spohn, 1988] in Section 3.3. And in Section 3.4, we characterize RbC as a revision with a set of so-called weak conditionals using the highly adaptive framework of c-revisions [Kern-Isberner, 2001] and show that these implementations inherit all relevant features of RbC. Before we conclude in Section 5, we compare RbC operators to credibility-limited revision operators in Section 4. The proofs in this paper are straightforward, but technical, and therefore omitted due to lack of space.

2 Formal Preliminaries

In this section, we define some basics from propositional resp. conditional logic and fix our notation. Then we deal with two qualitative belief representation frameworks in more detail.

2.1 Propositional and Conditional Logic

We denote by $\mathcal{L}$ the set of formulas of a propositional language built over a finite signature Σ using logical connectives *and* $\wedge$, *or* $\vee$ and *not* $\neg$. For conciseness of notation, we will omit the logical and-connector, writing $\alpha\beta$ instead of $\alpha \wedge \beta$, and overlining formulas will indicate negation, i.e., $\bar{\alpha}$ means $\neg\alpha$. The material implication $\alpha \Rightarrow \beta$ '*From α it (always) follows that β*' is, as usual, equivalent to $\bar{\alpha} \vee \beta$. We denote by $\top$ logical truths or tautologies and by $\perp$ logical falsehoods or contradictions. The set of all propositional interpretations over Σ is denoted by Ω_Σ. As the signature will be fixed throughout the paper, we will usually omit the subscript and simply write Ω. As usual, we write $\omega \models \alpha$ when a world satisfies α, i.e. when ω is a model of α. By slight abuse of notation, we will use ω both for the model and the corresponding conjunction of all positive or negated atoms. The set of all models of a formula α is denoted by $Mod(\alpha)$. The set of classical consequences of a set of formulas $A \subseteq \mathcal{L}$ is $Cn(A) = \{\alpha \in \mathcal{L} \mid A \models \alpha\}$. The deductively closed set of formulas which has exactly a subset $W \subseteq \Omega$ as models is called the *formal theory* of W and defined as $Th(W) = \{\alpha \in \mathcal{L} \mid \omega \models \alpha \text{ for all } \omega \in W\}$. Let $(\mathcal{L}|\mathcal{L}) = \{(\beta|\alpha) \mid \alpha, \beta \in \mathcal{L}\}$ be a flat *conditional language* where α is called the antecedent of $(\beta|\alpha)$, and β is its consequent. $(\beta|\alpha)$ expresses *'If α, then (plausibly) β'*. We extend the standard conditional language to a *weak conditional language* $(|\mathcal{L}|\mathcal{L}|)$. Weak conditionals represent negated conditional information, as used in *Rational monotony* [Lehmann and Magidor, 1992]. For a weak conditional $(|\beta|\alpha|)$, we call α the antecedent and β the consequent and $(|\beta|\alpha|)$ expresses *'If α, then β might be the case but $\bar{\beta}$ is not plausible'*. Weak conditionals express an agent's insecure attitude towards the consequent β if α is true, the negation of β is not plausible, but on the other hand β (only) might be true. It holds that $(|\beta|\alpha|)$ implements the negation of the corresponding standard conditional $(\bar{\beta}|\alpha)$, the former is accepted iff the latter is not [Lewis, 1973]. The evaluation of a weak conditional corresponds to the evaluation of the standard conditional [De Finetti, 1975], with verification $\alpha\beta$, falsification $\alpha\bar{\beta}$ and neutrality $\bar{\alpha}$.

2.2 Belief Representation

In most of the formal developments for methods of non-prioritized belief revision *epistemic entrenchment relations*, noted as $\leqslant_E$, are used. Epistemic entrenchment relations mirror the attitude of an agent towards her current beliefs, i.e. represent an inner ordering of the belief set. Some sentences in the belief set have a higher degree of epistemic entrenchment than others, i.e., are more easily abandoned when a contraction is carried out. These degrees of entrenchments are measured only qualitatively, 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" and $\alpha <_E \beta$ is defined as $\alpha \leqslant_E \beta$ but not $\beta \leqslant_E \alpha$. Following [Nayak, 1994], we define an epistemic entrenchment relation as follows:

Definition 1 ([Nayak, 1994])**.** *A total preorder $\leqslant$ over $\mathcal{L}$ is called a relation of* epistemic entrenchment*, if it satisfies the following conditions:*

- *(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 $\beta \leqslant_E \alpha \wedge \beta$*
- *(E4)* *If $\alpha \leqslant_E \beta$ for all $\alpha \in \mathcal{L}$, then $\beta \equiv \top$*

The original definition of entrenchment relations, such as in [Gärdenfors and Makinson, 1988], is given relative to a belief set. However, belief sets can be extracted from each epistemic entrenchment as they contain enough information by themselves (as noted in [Nayak, 1994]) via $Bel(\leqslant_E) = \{\alpha \in \mathcal{L} \mid \perp <_E \alpha\}$ if $\perp <_E \alpha$ for some α, otherwise $Bel(\leqslant_E) = \mathcal{L}$. Note that each epistemic entrenchment relation $\leqslant_E$ is uniquely defined via the entrenchment classification of maximal disjunctions over the signature Σ that constitutes $\mathcal{L}$.

In contrast to epistemic entrenchment relations, *total preorders (TPOs) on possible worlds* rank worlds according to their closeness to a belief set, i.e. they take on the perspective of belief revision and define an implausibility ordering on possible worlds $\omega \in \Omega$, s.t. for two worlds $\omega, \omega' \in \Omega$, $\omega \preceq \omega'$ means that ω is at least as plausible as ω', and $\omega \prec \omega'$ holds if $\omega \preceq \omega'$ but not $\omega' \preceq \omega$. The belief set of $\preceq$ is defined via minimal worlds, s.t. $Bel(\preceq) = Th(\{\omega \mid \omega \preceq \omega' \text{ for all } \omega \in \Omega\})$ and $\preceq \models \alpha$ if $\alpha \in Bel(\preceq)$. 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 $\min(Mod(\alpha), \preceq) \preceq \min(Mod(\beta), \preceq)$. To ease the notation, we define the following subsets of Ω for each plausibilistic TPO $\preceq$ and two arbitrary sentences γ, δ:

$$Mod_\gamma^{\preceq}(\delta) = \{\omega \in \Omega \mid \omega \models \delta \text{ and } \omega \prec \gamma\}.$$

Plausibilistic TPOs and epistemic entrenchment relations are dual approaches to belief representation and both formalisms are fundamental to AGM belief revision, see [Katsuno and Mendelzon, 1992] for TPOs and [Gärdenfors and Makinson, 1988] for entrenchment relations. More precisely, the following proposition summarizes the relationship between entrenchment relations and plausibility orderings which was already discussed in [Peppas and Williams, 1995; Grove, 1988] and will prove to be useful in the context of this paper:

Proposition 1. *For each epistemic entrenchment* $\leqslant_E$*,*

$$\bar{\alpha} \preceq \bar{\beta} \textit{ iff } \alpha \leqslant_E \beta \tag{1}$$

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

Applying (1) on the entrenchment classification of maximal disjunctions in $\leqslant_E$ leads us to possible worlds ranked by plausibility according to $\leqslant_E$. Plausibilistic TPOs are substantiated using *ordinal conditional functions* (OCFs, also called *ranking functions*) $\kappa : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ [Spohn, 1988], with $\kappa^{-1}(0) \neq \emptyset$, assigning to each world ω an implausibility rank $\kappa(\omega) > 0$ and $\kappa(\perp) = \infty$, i.e. the higher $\kappa(\omega)$ the less plausible ω is. The normalization constraint calls for worlds having maximal plausibility and the belief set $Bel(\kappa) = Th(\kappa^{-1}\{0\})$ is defined via worlds with minimal ranks. It is easy to see that each ranking function defines a TPO $\preceq_\kappa$ via

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

It holds that $\kappa(\alpha) := \min\{\kappa(\omega) \mid \omega \models \alpha\}$ and it holds that $\kappa \models \alpha$ if $\kappa(\overline{\alpha}) > 0$. A standard conditional $(\beta|\alpha)$ is accepted, $\kappa \models (\beta|\alpha)$, if $\kappa(\alpha\beta) < \kappa(\alpha\overline{\beta})$. For weak conditionals $(|\beta|\alpha|)$ this condition is weakened s.t. $\kappa \models (|\beta|\alpha|)$, if $\kappa(\alpha\beta) \leqslant \kappa(\alpha\overline{\beta})$.

3 Revision by Comparison

In this section, we clarify the methodology of Revision by Comparison, i.e., a belief revision mechanism that specifies the strengths of new beliefs according to the credibility of a reference sentence its coming from, and discuss its special hybrid belief change character. Furthermore, we transfer RbC to the framework of ranking functions where its intuitive strengths become more apparent and implement it as a conditional revision.

3.1 Basics of Revision by Comparison

Revision by Comparison $*_\alpha\beta$ revises a prior belief state over a *reference sentence* α and an *input sentence* β. The main idea can be expressed as follows: Accept β in the posterior belief state with a degree of entrenchment that at least equals that of α.

We allow for arbitrary input sentences $\beta \in \mathcal{L}$ but exclude reference sentences $\alpha \equiv \top$ to avoid cases where we have to accept input sentences β as far as logical truths and thus would have a conflict with (E4). In [Fermé and Rott, 2004], RbC is defined via a function $*_\alpha\beta$ that can be applied to a belief set $Bel(\leqslant_E)$ extracted from an entrenchment relation $\leqslant_E$, yielding a new belief set $Bel(\leqslant_E) *_\alpha \beta$. Using belief states makes RbC suitable for iterated belief revision, thus in the following, we will neglect the belief sets and recall RbC as an operation on belief states represented by epistemic entrenchment relations.

Definition 2 (Revision by Comparison, $\leqslant_E^{*_{\alpha,\beta}}$)**.** *Let* $\leqslant_E$ *be an epistemic entrenchment relation and* $\leqslant_E^{*_{\alpha,\beta}}$ *be the posterior entrenchment relation after RbC with* α *as reference and* β *as input sentence. Then* $\leqslant_E^{*_{\alpha,\beta}}$ *is defined by*

$$\gamma \leqslant_E^{*_{\alpha,\beta}} \delta \textit{ iff } \begin{cases} \alpha \wedge (\beta \Rightarrow \gamma) \leqslant_E (\beta \Rightarrow \delta), & \textit{if } \gamma \leqslant_E \alpha \\ \gamma \leqslant_E \delta, & \textit{otherwise} \end{cases} \tag{3}$$

for any arbitrary sentences γ, δ*.*

Exceeding the basic success condition formulated above, Fermé and Rott formulated desirable properties for RbC in [Fermé and Rott, 2004] as follows: Let $\leqslant_E^{*_{\alpha,\beta}}$ be the posterior entrenchment relation after an RbC $*_\alpha\beta$. Then the following properties hold for $\leqslant_E^{*_{\alpha,\beta}}$:

(Success) β is at least as entrenched as α: $\alpha \leqslant_E^{*_{\alpha,\beta}} \beta$

(Lifting) It does no extra lifting: $\alpha <_E^{*_{\alpha,\beta}} \beta$ iff $\alpha <_E \beta$

(α-level) It does no extra lifting: $\alpha <_E^{*_{\alpha,\beta}} \beta$ iff $\alpha <_E \beta$

(Min) The entrenchment relation between two sentences that are at least as entrenched as α in the prior belief state will not be affected by the operation: Let $\gamma, \delta \in \mathcal{L}$, if $\alpha \leqslant_E \gamma$ and $\alpha \leqslant_E \delta$ then: $\gamma \leqslant_E^{*_{\alpha,\beta}} \delta$ iff $\gamma \leqslant_E \delta$

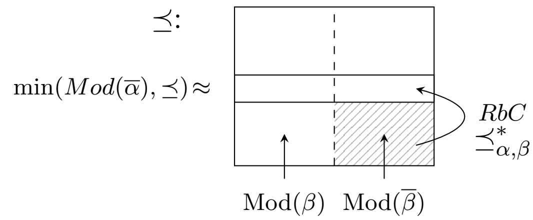

Figure 1: Schematic Revision by Comparison for plausibilistic preorders. The hatched area displays $Mod^{\preceq}_{\overline{\alpha}}(\beta)$

3.2 Revision by Comparison for Plausibilistic TPOs

We transfer the methodology underlying RbC in the framework of entrenchment relations to qualitative constraints for plausibilistic TPOs $\preceq$ and $\preceq^*_{\alpha,\beta}$ corresponding to the entrenchment relations $\leqslant_E$ and $\leqslant_E^{*_{\alpha,\beta}}$:

$$\omega \preceq^*_{\alpha,\beta} \omega' \text{ iff } \begin{cases} \alpha \Rightarrow (\beta\omega) \preceq \beta\omega', & \text{if } \omega \preceq \bar{\alpha} \\ \omega \preceq \omega', & \text{otherwise} \end{cases} \tag{4}$$

These constraints follow immediately from (3) via (1) for γ, δ as maximal disjunctions; their negations are then maximal conjunctions and correspond to possible worlds, yielding (4) for possible worlds ω resp. ω'.

It is not easy to understand the strategy of RbC and therefore also its special features are not clear. In order to discuss these in more detail, we clarify the strategy of RbC expressed in (3). In Figure 1 a schematic RbC for plausibilistic TPOs is depicted: The relations among worlds in $Mod(\beta)$ and among worlds in $Mod(\overline{\beta})$ that are less plausible than $\bar{\alpha}$ are kept when performing RbC. Because, for worlds $\omega_1 \in Mod(\beta)$ that are less or equally plausible than $\bar{\alpha}$, it holds that $\alpha \Rightarrow (\beta\omega_1) \equiv \bar{\alpha} \vee \omega_1 \equiv \omega_1$, i.e. the first case in (4) is equivalent to $\omega_1 \preceq \omega_2$ for $\omega_2 \in Mod(\beta)$. The second case implies that the relations are preserved for the remaining worlds in $Mod(\beta)$ and $\overline{\beta}$-worlds that are less plausible than $\overline{\alpha}$. But all worlds $\omega \in Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta})$, marked in the hatched area in Figure 1 are shifted up to the level of plausibility where $\overline{\alpha}$ resides, thus $\overline{\alpha} \preceq^*_{\alpha,\beta} \overline{\beta}$. This follows from the first case in (4), since $\alpha \Rightarrow (\beta\omega) \preceq \beta\omega'$ is always true if ω' is some minimal $\overline{\alpha}$-world and conversely, $\alpha \Rightarrow (\beta\omega') \preceq \beta\omega \equiv \bot$ holds. This shifting of worlds in $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta})$ results in the fact that worlds in $Mod_{\overline{\alpha}}^{\preceq}(\beta)$ become the most plausible ones in the posterior TPO. This is illustrated in Figure 1 and in Example 2. We provide semantic constraints that depict the mechanism of RbC in a more comprehensible way than (4):

Theorem 1. *Let $\preceq$ be a plausibilistic TPO and $\preceq *_\alpha \beta = \preceq^*_{\alpha,\beta}$ be the posterior TPO according to* (4)*. Then it holds that*

$$\omega \preceq^*_{\alpha,\beta} \omega' \textit{ iff } \begin{cases} \omega \preceq \omega' \textit{ and } \overline{\alpha} \preceq \omega \\ \omega \preceq \omega' \textit{ and } \omega \in Mod_{\overline{\alpha}}^{\preceq}(\beta), \omega' \models \beta \\ \overline{\alpha} \preceq \omega' \textit{ and } \omega \in Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta}), \omega' \models \beta \\ \omega \preceq \overline{\alpha},\ \omega' \not\models \beta \end{cases} \tag{5}$$

Theorem 1 follows from (4) via equivalent reformulations and strict case distinctions between worlds in $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta})$ and $Mod_{\overline{\alpha}}^{\preceq}(\beta)$. For the contradictory case $\alpha \equiv \bot$, it holds that $\overline{\alpha}$ is a logical truth and we fix by convention that $\overline{\alpha}$ is represented by an arbitrary minimal world in the prior ordering. Then the constraints in (4) and (5) lead to no changes in the prior ordering because the first case in (5) holds for all $\omega \in \Omega$ and it holds that $Mod_{\overline{\alpha}}^{\preceq}(\beta), Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta}) = \emptyset$.

The characterization in (5) makes clear that the change of the prior belief state depends on the relation between the reference and the input sentence, something which was already observed in [Fermé and Rott, 2004] and thus RbC displays a hybrid belief change operation. We can further specify this statement: the possible worlds in $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta})$ and $Mod_{\overline{\alpha}}^{\preceq}(\beta)$ are crucial for the posterior ordering and the kind of change performed by RbC is solely dependent on whether these sets are empty or not. It follows from (4) and Theorem 1 that the hybrid character of RbC transfers immediately to our characterization (5). There are three cases of prior orderings $\preceq$ which correspond to the cases defined in [Fermé and Rott, 2004] and have different effects on the posterior state $\preceq^*_{\alpha,\beta}$:

β**-revision:** If $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta}), Mod_{\overline{\alpha}}^{\preceq}(\beta) \neq \emptyset$, s.t. $\overline{\beta}, \beta \prec \overline{\alpha}$, then $\preceq^*_{\alpha,\beta} \models \beta$

Vacuous case: If $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta}) = \emptyset$, s.t. $\overline{\alpha} \preceq \overline{\beta}$, then $\preceq^*_{\alpha,\beta} = \preceq$

α**-contraction:** If $Mod_{\overline{\alpha}}^{\preceq}(\beta) = \emptyset$, s.t. $\overline{\alpha} \preceq \beta$, then $\preceq^*_{\alpha,\beta} \not\models \alpha$

The β-revision corresponds to the *intended case* in [Fermé and Rott, 2004], in which the negation of the input $\overline{\beta}$ is more plausible than the negation of the reference $\overline{\alpha}$. Then it holds that $\beta \in Bel(\preceq^*_{\alpha,\beta})$ and RbC displays a revision with β. If $\overline{\beta}$ is already more or equally plausible as $\overline{\alpha}$, then it follows from (1) that **(Success)** is already satisfied and RbC does not change anything. We speak of the vacuous case. Yet, if $Mod_{\overline{\alpha}}^{\preceq}(\beta) = \emptyset$, then $\overline{\alpha}$ is more plausible than the input β and RbC performs a contraction of α such that $\alpha \notin Bel(\preceq^*_{\alpha,\beta})$. This case is called the *unsuccessful case* in [Fermé and Rott, 2004] and it corresponds to the idea that the reference sentence displays the source β is coming from, so that the plausibility of α can be interpreted as the reliability of this source, which decreases in the case of a sufficiently implausible input. Note that if $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta}), Mod_{\overline{\alpha}}^{\preceq}(\beta) = \emptyset$ then the vacuous case and α-contraction coincide. In this case RbC does not change the prior ordering $\preceq$ and $\alpha \notin Bel(\preceq^*_{\alpha,\beta})$ follows from $\alpha \notin Bel(\preceq)$.

$\overline{\alpha}\overline{\beta}$	$\overline{\alpha}\overline{\beta}$	$\alpha\overline{\beta}$	$\alpha\overline{\beta}$	$\overline{\alpha}\beta, \alpha\beta$	
$\overline{\alpha}\beta$	$\overline{\alpha}\beta, \alpha\overline{\beta}$	$\alpha\beta$	$\alpha\beta$	$\overline{\alpha}\overline{\beta}$	$\overline{\alpha}\beta, \alpha\beta$
$\alpha\beta, \alpha\overline{\beta}$	$\alpha\beta$	$\overline{\alpha}\overline{\beta}, \overline{\alpha}\beta$	$\overline{\alpha}\overline{\beta}, \overline{\alpha}\beta$	$\alpha\overline{\beta}$	$\overline{\alpha}\overline{\beta}, \alpha\overline{\beta}$
$\preceq_1$	$\preceq_1^*$	$\preceq_2$	$\preceq_2^*$	$\preceq_3$	$\preceq_3^*$

Figure 2: Examples for Revision by Comparison for plausibilistic TPOs $\preceq_1, \preceq_2$ and $\preceq_3$. The worlds are ordered in ascending order of plausibility with the most plausible worlds at the bottom.

Example 2. *In Figure 2 we give three examples of RbC-revised TPOs s.t. $\preceq_1^* = (\preceq^*_{\alpha,\beta})_1$, $\preceq_2^* = (\preceq^*_{\alpha,\beta})_2$ and $\preceq_3^* = (\preceq^*_{\alpha,\beta})_3$. The first revised TPO $\preceq_1^*$ corresponds to a β-revision with $Mod_{\overline{\alpha}}^{\preceq_1}(\overline{\beta}) = \{\alpha\overline{\beta}\}$ and $Mod_{\overline{\alpha}}^{\preceq_1}(\beta) = \{\alpha\beta\}$. As we can see it holds that $\preceq_1^* \models \beta$. For $\preceq_2$ it holds that $Mod_{\overline{\alpha}}^{\preceq_2}(\overline{\beta}) = \emptyset$ and the operation does not change anything s.t. $\preceq_2 = \preceq_2^*$, this corresponds to the vacuous case. For the third TPO RbC performs an α-contraction, since $Mod_{\overline{\alpha}}^{\preceq_3}(\beta) = \emptyset$ and it holds that $\preceq_3^* \not\models \alpha$.*

In the following subsection, we will discuss the three above mentioned cases and their effect on the belief change induced by RbC in the framework of ranking function, but for now we complete the examination of RbC $*_\alpha \beta$ in the possible-worlds semantic and transfer the properties of RbC as follows:

(Success)$_{tpo}$ $\overline{\alpha}$ is at least as plausible as $\overline{\beta}$: $\overline{\alpha} \preceq^*_{\alpha,\beta} \overline{\beta}$

(Lifting)$_{tpo}$ It does no extra lifting: $\overline{\alpha} \prec^*_{\alpha,\beta} \overline{\beta}$ iff $\overline{\alpha} \prec \overline{\beta}$

(α**-level**)$_{tpo}$ For any sentence γ, it holds that $\overline{\alpha} \prec^*_{\alpha,\beta} \gamma$ iff $\overline{\alpha} \prec \gamma$, and $\gamma \preceq^*_{\alpha,\beta} \overline{\beta}$ iff $\gamma \preceq \overline{\alpha}$ or $\gamma \preceq \overline{\beta}$

(Min)$_{tpo}$ If $\overline{\alpha} \preceq \gamma$ and $\overline{\alpha} \preceq \delta$ then: $\gamma \preceq^*_{\alpha,\beta} \delta$ iff $\gamma \preceq \delta$

The properties **(Success)**, **(Lifting)**, (α**-level)** and **(Min)** hold for $\leqslant_E^{*_{\alpha,\beta}}$ [Fermé and Rott, 2004] and thus, we can conclude for each TPO $\preceq^*_{\alpha,\beta}$ satisfying (4) that $\preceq^*_{\alpha,\beta}$ satisfies **(Success)**$_{tpo}$, **(Lifting)**$_{tpo}$, (α**-level**)$_{tpo}$ and **(Min)**$_{tpo}$. Note that in the contradictory case ($\alpha \equiv \bot$), **(Success)**$_{tpo}$ holds trivially, since $\top \preceq \overline{\beta}$ is always true. Moreover, the posterior epistemic states of RbC tend to make relations among worlds much coarser, since we loose plausibility distinctions between worlds in $Mod_{\overline{\alpha}}^{\preceq}(\overline{\beta})$. Hence RbC does not satisfy the Darwiche and Pearl postulates [Darwiche and Pearl, 1997].

3.3 Revision by Comparison for Ranking Functions

Next, we present a semi-quantitative version of RbC that implements the semantical recipe of RbC given in (5) for ranking functions κ in a simple, yet elegant way. Furthermore, the use of ranking functions makes the change on the prior belief state and the dependence on the relation between input and reference information in RbC more explicit.

First, we define the set of β- resp. $\overline{\beta}$-worlds that are more plausible than $\overline{\alpha}$ according to κ analog to $Mod^{\preceq}_{\overline{\alpha}}(\beta)$ resp. $Mod^{\preceq}_{\overline{\alpha}}(\overline{\beta})$ as follows:

$$Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa) = \{\omega \in \Omega \,|\, \omega \models \overline{\beta} \text{ and } \kappa(\omega) < \kappa(\overline{\alpha})\},$$
$$Mod^{\leq}_{\overline{\alpha}}(\beta,\kappa) = \{\omega \in \Omega \,|\, \omega \models \beta \text{ and } \kappa(\omega) < \kappa(\overline{\alpha})\}.$$

Definition 3 (RbC for OCFs)**.** *Let κ be a ranking function. We define the* Revision by Comparison $*_\alpha\beta$ with input sentence β and reference sentence α for ranking functions *as follows*

$$\kappa *_\alpha \beta(\omega) = \kappa^{*_{\alpha,\beta}}(\omega) = \kappa_0 + \begin{cases}\kappa(\overline{\alpha}), & \omega \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)\\ \kappa(\omega), & \textit{otherwise}\end{cases} \quad (6)$$

where $\kappa_0 = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\}$ is a normalization constant.

The normalization constant κ_0 can be computed as follows: We have $\kappa_0 = -min_{\omega \notin Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)}\{\kappa(\overline{\alpha}), \kappa(\omega)\}$. For $\omega \notin Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)$, it must hold that $\omega \models \beta$ or $\kappa(\omega) \geq \kappa(\overline{\alpha})$. Hence, only $\omega \models \beta$ can be relevant for the minimum. Thus, we can conclude $\kappa_0 = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\}$. For the contradictory case $\alpha \equiv \bot$, it holds that $Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa) = \emptyset$ and therefore RbC does not modify the prior ranking function and $\kappa^{*_{\alpha,\beta}}(\omega) = \kappa(\omega)$ for all worlds and $\kappa_0 = 0$.

The methodology of RbC depicted in Figure 1 and defined via the constraints in (5) can be seen directly from Definition 3. Now, it is easy to see that the worlds in $Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)$ are shifted to the $\overline{\alpha}$-level of plausibility, s.t. $\kappa^{*_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa^{*_{\alpha,\beta}}(\overline{\beta})$ and that the plausibility relations among the worlds $\Omega \setminus Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)$ remain the same. Their ranks only change according to the normalization of the posterior ranking function. We get the following theorem:

Theorem 2. *Let κ be a ranking function and $\kappa^{*_{\alpha,\beta}} = \kappa *_\alpha \beta$ be the RbC-revised ranking function from* (6)*. Then the corresponding TPO $\preceq_{\kappa^{*_{\alpha,\beta}}}$ can be defined from $\preceq_\kappa$ by* (5)*.*

Since $\kappa^{*_{\alpha,\beta}}$ satisfies (5), it follows from Theorem 1 that it also satisfies (4) and therefore the properties of RbC follow via (2) for $\kappa^{*_{\alpha,\beta}}$ and we get that:

Theorem 3. *Let κ be a ranking function and $\kappa^{*_{\alpha,\beta}} = \kappa *_\alpha \beta$. The associated TPO $\preceq_{\kappa^{*_{\alpha,\beta}}}$ satisfies **(Success)**$_{tpo}$, **(Lifting)**$_{tpo}$, **(α-level)**$_{tpo}$and **(Min)**$_{tpo}$.*

Theorem 2 and 3 together show that Definition 3 is a suitable definition of RbC for ranking functions. In the following, we will show that RbC can also be defined as c-change operation via a designated set of conditionals. This conditional revision also reflects the hybrid character of RbC which we discussed earlier, since it displays a revision or a contraction depending on the prior belief state.

3.4 Revision by Comparison as a c-Revision

C-revisions, introduced by Kern-Isberner [Kern-Isberner, 2001], provide a highly general framework for revising ranking functions by sets of conditionals with respect to conditional interactions, while preserving conditional beliefs in the former belief state as far as possible. In this section, we will use c-revisions to revise a ranking function κ with a suitable set of weak conditionals to obtain $\kappa *_\alpha \beta$ as given in (6).

In general a *c-revision with a set of weak conditionals* $\Delta = \{(|\delta_i|\gamma_i|)\}_{i=1}^n$ is an OCF $\kappa_c = \kappa * \Delta$ of the form

$$\kappa_c(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\omega \models \gamma_i\overline{\delta_i}} \nu_i^-$$

with a normalization constant κ_0 that ensures that κ_c is a ranking function and non-negative impact factors ν_i^- that satisfy the following inequalities:

$$\nu_i^- \geq \min_{\omega \models \gamma_i\delta_i}\Big(\kappa(\omega) + \sum_{\omega \models \gamma_j\overline{\delta_j}}^{j \neq i} \nu_j^-\Big) - \min_{\omega \models \gamma_i\overline{\delta_i}}\Big(\kappa(\omega) + \sum_{\omega \models \gamma_j\overline{\delta_j}}^{j \neq i} \nu_j^-\Big) \quad (7)$$

These inequalities ensure that $\kappa_c \models \Delta$. Note that, (7) is satisfiable for each Δ, since sets that consist solely of weak conditionals are always consistent [Sezgin and Kern-Isberner, 2021]. The c-revision with Δ corresponds to the c-contraction with a set of conditionals that consists of all negated conditionals from Δ, for more details see [Kern-Isberner *et al.*, 2017].

Consider

$$\Delta^{*_{\alpha,\beta}} = \{(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \,|\, \omega \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)\}$$

expressing the constraints that $\kappa^{*_{\alpha,\beta}}(\overline{\alpha}(\overline{\alpha} \vee \omega)) = \kappa^{*_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa^{*_{\alpha,\beta}}(\alpha(\overline{\alpha} \vee \omega))$ should hold in the revised ranking function for $\omega \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)$. For contradictory reference sentences $\alpha \equiv \bot$, we get that $\Delta^{*_{\alpha,\beta}} = \emptyset$ since $Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa) = \emptyset$ and therefore the c-revision does not change anything. The minimal c-revision with $\Delta^{*_{\alpha,\beta}}$ can be computed as follows: First, it holds that the falsification of conditionals $(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \in \Delta^{*_{\alpha,\beta}}$ is

$$(\overline{\alpha} \vee \omega) \wedge \alpha \equiv \omega \wedge \alpha \equiv \omega, \quad (8)$$

because $\omega \in Mod(\alpha)$ holds for $\omega \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)$, since $\kappa(\omega) < \kappa(\overline{\alpha})$ and due to minimality of rank $\kappa(\overline{\alpha})$. Thus, for each $\omega \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)$ only a single conditional is falsified and a c-revision with $\Delta^{*_{\alpha,\beta}}$, $\kappa_c^{*_{\alpha,\beta}} = \kappa * \Delta^{*_{\alpha,\beta}}$, is given by:

$$\kappa_c^{*_{\alpha,\beta}}(\omega) = \kappa_0 + \kappa(\omega) + \begin{cases}\nu_\omega^-, & \omega \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa)\\ 0, & \text{othw.}\end{cases} \quad (9)$$

with κ_0 a normalization constant and ν_ω^- as non-negative impact factors for each weak conditional in $\Delta^{*_{\alpha,\beta}}$ satisfying the constraints

$$\begin{aligned}\nu_\omega^- \geq &\min_{\omega' \models \overline{\alpha}}\{\kappa(\omega') + \sum_{\omega' \models \tilde{\omega} \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa),\ \tilde{\omega} \neq \omega} \nu_{\tilde{\omega}}^-\}\\ &- \min_{\omega' \models \omega \wedge \alpha}\{\kappa(\omega') + \sum_{\omega' \models \tilde{\omega} \in Mod^{\leq}_{\overline{\alpha}}(\overline{\beta},\kappa),\ \tilde{\omega} \neq \omega} \nu_{\tilde{\omega}}^-\}\\ =&\kappa(\overline{\alpha}) - \kappa(\omega).\end{aligned} \quad (10)$$

Note that, in both minima the sums equal zero and the second minimum in (10) ranges only over ω (see (8)). Thus, to obtain a minimal c-revision, we can choose

$$\nu_\omega^- = \kappa(\overline{\alpha}) - \kappa(\omega). \quad (11)$$

Hence,

$$\kappa_c^{*_{\alpha,\beta}}(\omega) = \kappa_0 + \begin{cases} \kappa(\overline{\alpha}), & \omega \in Mod^{\leqslant}_{\overline{\alpha}}(\overline{\beta}, \kappa) \\ \kappa(\omega), & \text{otherwise} \end{cases}$$

The normalization constant κ_0 can be substantiated as $\kappa_0 = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\}$. following the same argumentation as for $\kappa^{*_{\alpha,\beta}}$. So, we get the same posterior ranking function as the quantitative RbC and state the following Theorem:

Theorem 4. *Let κ be a ranking function. Let $\kappa^{*_{\alpha,\beta}}$ be the posterior ranking function after quantitative RbC as defined in Definition 3 and $\kappa_c^{*_{\alpha,\beta}} = \kappa * \Delta^{*_{\alpha,\beta}}$ be the posterior ranking function after the minimal c-revision. Then it holds for all $\omega \in \Omega$ that $\kappa^{*_{\alpha,\beta}}(\omega) = \kappa_c^{*_{\alpha,\beta}}(\omega)$, i.e., both revision mechanisms yield the same results.*

The following corollary follows from Theorem 3 and 4 and shows that RbC defined as a c-revision satisfies all properties of RbC:

Corollary 2. *Let κ be a ranking function and $\kappa_c^{*_{\alpha,\beta}} = \kappa * \Delta^{*_{\alpha,\beta}}$ a minimal c-revision. Then it holds that the corresponding TPO $\preceq_{\kappa_c^{*_{\alpha,\beta}}}$ defined via* (2) *satisfies* (4) *resp.* (5) *and **(Success)**$_{tpo}$, **(Lifting)**$_{tpo}$ and **(α-level)**and **(Min)**$_{tpo}$.*

The advantage of implementing RbC for ranking functions by a c-revision using sets of weak conditionals is that the core of the required change process can be made very explicit in $\Delta^{*_{\alpha,\beta}}$, also showing clearly what does not change, namely the relations among worlds in $\Omega \setminus Mod^{\leqslant}_{\overline{\alpha}}(\overline{\beta}, \kappa)$. Thus, weak conditionals support the hybrid belief change character of RbC. The dependence on the prior ranking function is made visible and we discuss the three cases that we introduced in Section 3.2 and their impact on the RbC-revised ranking function $\kappa_c^{*_{\alpha,\beta}}$:

β-revision: Since $Mod^{\preceq}_{\overline{\alpha}}(\overline{\beta}) \neq \emptyset$ it holds that $\Delta^{*_{\alpha,\beta}} \neq \emptyset$ and we can conclude from $Mod^{\leqslant}_{\overline{\alpha}}(\beta, \kappa) \neq \emptyset$ that $\kappa(\beta) < \kappa(\overline{\alpha})$. Hence, $\kappa(\beta) = \kappa_c^{*_{\alpha,\beta}}(\beta) < \kappa_c^{*_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa_c^{*_{\alpha,\beta}}(\overline{\beta})$. Thus $\kappa_c^{*_{\alpha,\beta}}(\overline{\beta}) > 0$, s.t. $\kappa_c^{*_{\alpha,\beta}} \models \beta$ and RbC for ranking functions performs a revision with β

Vacuous case: Since $Mod^{\preceq}_{\overline{\alpha}}(\overline{\beta}) = \emptyset$ it holds that $\Delta^{*_{\alpha,\beta}} = \emptyset$. Thus, $\kappa_c^{*_{\alpha,\beta}}(\omega) = \kappa(\omega)$ for all $\omega \in \Omega$

α-contraction: Since $Mod^{\preceq}_{\overline{\alpha}}(\beta) = \emptyset$ it holds that $\kappa(\overline{\alpha}) \leqslant \kappa(\beta)$, i.e. $\kappa_c^{*_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa_c^{*_{\alpha,\beta}}(\beta)$. Since $\kappa_c^{*_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa_c^{*_{\alpha,\beta}}(\overline{\beta})$ holds for all $\kappa_c^{*_{\alpha,\beta}}$ and due to the minimality of ranks, we get that $\kappa_c^{*_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa_c^{*_{\alpha,\beta}}(\omega)$ for all $\omega \in \Omega$, thus $\kappa_c^{*_{\alpha,\beta}} \not\models \alpha$ and RbC for ranking functions results in a contraction of α

It is clear from Theorem 4 that these results also hold for $\kappa^{*_{\alpha,\beta}}$ as defined in Definition 3. We have seen that RbC can be displayed as a revision with negative information, therefore the Darwiche and Pearl postulates [Darwiche and Pearl, 1997] are not the right framework to capture the dynamic belief change induced by RbC. We illustrate RbC with input sentence β and α as reference as a c-revision with $\Delta^{*_{\alpha,\beta}}$ in the following example:

Example 3. *In Figure 3 a ranking function κ is depicted with $Mod^{\leqslant}_{\overline{\alpha}}(\overline{\beta}, \kappa) = \{\alpha\overline{\beta}\gamma\} \neq \emptyset$ and $Mod^{\leqslant}_{\overline{\alpha}}(\beta, \kappa) = \{\alpha\beta\gamma, \alpha\beta\overline{\gamma}\} \neq \emptyset$, i.e. RbC $\kappa *_\alpha \beta$ with input β and reference α performs a β-contraction. It holds that $\Delta^{*_{\alpha,\beta}} = \{(|\overline{\alpha}|\overline{\alpha} \vee (\alpha\overline{\beta}\gamma)|)\}$ with impact factor $\nu^-_{\alpha\overline{\beta}\gamma} = 2$ and the c-revised ranking function $\kappa_c^{*_{\alpha,\beta}} = \kappa * \Delta^{*_{\alpha,\beta}} = \kappa *_\alpha \beta$ can be found in Figure 3 and it holds that $\kappa_c^{*_{\alpha,\beta}} \models \beta$.*

	κ	$\kappa_c^{*_{\alpha,\beta}}$
3	$\alpha\overline{\beta}\overline{\gamma}$, $\overline{\alpha}\beta\overline{\gamma}$ $\overline{\alpha}\overline{\beta}\gamma$, $\overline{\alpha}\overline{\beta}\overline{\gamma}$	$\alpha\overline{\beta}\overline{\gamma}$, $\overline{\alpha}\beta\overline{\gamma}$ $\overline{\alpha}\overline{\beta}\gamma$, $\overline{\alpha}\overline{\beta}\overline{\gamma}$
2	$\overline{\alpha}\beta\gamma$	$\overline{\alpha}\beta\gamma$, $\alpha\overline{\beta}\gamma$
1	$\alpha\beta\overline{\gamma}$	$\alpha\beta\overline{\gamma}$
0	$\alpha\beta\gamma, \alpha\overline{\beta}\gamma$	$\alpha\beta\gamma$

Figure 3: Example of Revision by Comparion $\kappa *_\alpha \beta$ performed via a c-revision with $\Delta^{*_{\alpha,\beta}} = \{(|\overline{\alpha}|\overline{\alpha} \vee (\alpha\beta\overline{\gamma})|)\}$

4 Related Work

The dynamic belief change performed by RbC moves it in the vicinity of *credibility-limited revision operators (CL-operators)* $\circ$ [Hansson *et al.*, 2001], where a revision on a general epistemic state Ψ is solely performed if the input information is part of a set of credible formulas $Bel(\Psi) \subseteq \mathcal{C}$ otherwise the prior belief state is kept. In [Booth *et al.*, 2012], these operators are defined for epistemic states. Using RbC we clarify the methodology of choosing a set of credible worlds. We have seen that RbC with input β and reference α performs a revision with β in the case that $\beta, \overline{\beta}$ are more plausible than $\overline{\alpha}$, i.e. if $Mod^{\preceq}_{\overline{\alpha}}(\overline{\beta}), Mod^{\preceq}_{\overline{\alpha}}(\beta) \neq \emptyset$. Thus RbC generates a restricted set of credible worlds $\mathcal{C}_{\alpha,\beta} = Mod^{\preceq}_{\overline{\alpha}}(\beta)$ which is dependent from the inputs β and α with $Bel(\preceq) \cap Mod(\beta) \subseteq \mathcal{C}_{\alpha,\beta}$. The set is restricted since not the whole belief set $Bel(\preceq)$ is part of $\mathcal{C}_{\alpha,\beta}$. We define the belief set of an *RbC-based CL-operator* $\circ^{RbC}_{\alpha}$ as follows:

$$Bel(\preceq \circ^{RbC}_{\alpha} \beta) = \begin{cases} Bel(\preceq *_\alpha \beta), \text{if } Mod(\beta) \cap Mod^{\preceq}_{\overline{\alpha}}(\beta) \neq \emptyset \\ Bel(\preceq), \quad \text{otherwise} \end{cases}$$

RbC is appealing as basis for a CL-operator since it allows for a flexible and reasonable determination of the set of credible worlds via choosing suitable reference sentences. Exploring extensions of CL-operators and their connection to RbC more thoroughly is part of our future work.

5 Conclusion

We have clarified the strategy underlying Revision by Comparison and transferred the versatile belief change mechanism to the semi-quantitative framework of ranking functions. The implementation of the RbC-revised ranking function inherits all characteristics of Revision by Comparison as a hybrid belief change mechanism. Using negated conditional information, we defined RbC as a conditional revision and therefore paved the way to the implementation of interesting operators capable of dealing with hybrid belief change operators. At last, we discussed the relationship between RbC and credibility-limited belief revision operators.

References

[Alchourrón *et al.*, 1985] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 06 1985.

[Booth *et al.*, 2012] Richard Booth, Eduardo Fermé, Sébastien Konieczny, and Ramón Pino Pérez. Credibility-limited revision operators in propositional logic. *13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012*, 01 2012.

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

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

[Fermé and Rott, 2004] Eduardo L. Fermé and Hans Rott. Revision by comparison. *Artif. Intell.*, 157(1-2):5–47, 2004.

[Gärdenfors and Makinson, 1988] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In Moshe Y. Vardi, editor, *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, USA, March 1988*, pages 83–95. Morgan Kaufmann, 1988.

[Grove, 1988] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.

[Hansson *et al.*, 2001] Sven Ove Hansson, Eduardo Leopoldo Fermé, John Cantwell, and Marcelo Alejandro Falappa. Credibility limited revision. *The Journal of Symbolic Logic*, 66(4):1581–1596, 2001.

[Hansson, 2008] Sven Ove Hansson. What's new isn't always best. *Theoria*, 63:1–13, 2008.

[Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artif. Intell.*, 52(3):263–294, 1992.

[Kern-Isberner *et al.*, 2017] Gabriele Kern-Isberner, Tanja Bock, Kai Sauerwald, and Christoph Beierle. Iterated contraction of propositions and conditionals under the principle of conditional preservation. In Christoph Benzmüller, Christine Lisetti, and Martin Theobald, editors, *GCAI 2017. 3rd Global Conference on Artificial Intelligence*, volume 50 of *EPiC Series in Computing*, pages 78–92. EasyChair, 2017.

[Kern-Isberner, 2001] Gabriele Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*, volume 2087 of *Lecture Notes in Computer Science*. Springer, Berlin, 2001.

[Lehmann and Magidor, 1992] Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55(1):1–60, 1992.

[Lewis, 1973] David K. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.

[Nayak, 1994] Abhaya C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41(3):353–390, 1994.

[Peppas and Williams, 1995] Pavlos Peppas and Mary-Anne Williams. Constructive Modelings for Theory Change. *Notre Dame Journal of Formal Logic*, 36(1):120 – 133, 1995.

[Schwind *et al.*, 2018] Nicolas Schwind, Sébastien Konieczny, and Pierre Marquis. On belief promotion. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, pages 297–307. AAAI Press, 2018.

[Sezgin and Kern-Isberner, 2021] Meliha Sezgin and Gabriele Kern-Isberner. System Z for conditional belief bases with positive and negative information. In Eric Bell and Fazel Keshtkar, editors, *Proceedings of the Thirty-Fourth International Florida Artificial Intelligence Research Society Conference, North Miami Beach, Florida, USA, May 17-19, 2021*, 2021.

[Spohn, 1988] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, pages 105–134. Springer, Dordrecht, 1988.