Ranking Revision with Conditional Knowledge Bases

Emil Weydert

University of Luxembourg

Abstract

We introduce a new model for belief dynamics where the belief states are ranking measures and the informational inputs are finite sets of parametrized conditionals interpreted by ranking constraints. The approach is inspired by the minimal information paradigm and generalizes ranking construction strategies developed for default reasoning. We show its handling of principles for conditional and parallel revision.

1 Introduction

Since the seminal work of Alchourron, Gärdenfors, and Makinson [AGM 85], many formal models of belief revision have seen the day. In a nutshell, for a given belief state and informational input, they try to identify the most appropriate revised state(s). In the context of iterated belief revision, this typically includes the revision of higher-order information, like doxastic preferences or plausibility measures. It is useful to see belief change as a two-step process. First, the input is evaluated in the light of the prior belief structure and translated into a constraint over successor states. Secondly, a revised state is chosen relative to this goal condition, the prior state, and general rationality considerations. For instance, in standard belief revision, where the input is a sentence φ, the constraint asks for plain belief in φ, resp. the absence of plain belief if we seek belief contraction. But we may also consider more general constraints, e.g. those expressed by a finite set of belief conditionals originating from a reliable introspective source.

Iterated propositional revision goes back to Spohn's seminal work [Spo 88,90,08]. Iterated conditional belief revision started with Boutilier and Goldszmidt [BG 93, Bou 96]. Unfortunately, their rudimentary qualitative minimal change strategy prevented an appropriate handling of independence and seemed to radical. Iterated multiple propositional revision was investigated among others by Zhang [ZF 01, Zha 04], and Delgrande, Jin [DJ 08], who offered specific algorithms and rationality postulates. But they did not discuss conditional belief change. This issue was taken up by Kern-Isberner, who identified and instantiated nine requirements for iterated revision with single conditionals, and also proposed general principles for multiple conditional revision [KI 99,02,04]. Independently, Weydert proposed two accounts for iterated multiple conditional belief change in the context of the ranking measure framework [Wey 99,05], where belief states are modeled by ranking measures [Wey 94], and conditionals are interpreted by constraints over these. Ranking measures, which generalize Spohn's κ-functions (NCFs), are well-behaved (im)plausibility valuations able to deal with graded plain belief and independence information in a coherent way.

However, all these revision formalisms face serious conceptual problems, e.g. linked to their ad hoc character, or the naive virtual conditionalization method. In the present paper we want to tackle these issues by introducing a new direct ranking construction algorithm for multiple graded conditional revision. It generalizes a technique we developed for default reasoning and is based on a two-step procedure inspired by the Levi identity. In particular, we will show how our approach handles old and new postulates for iterated multiple/conditional revision.

2 Ranking measure epistemology

Our belief semantics is based on ranking measures [Wey 94]. These are quasi-probabilistic plausibility valuations expressing the degree of disbelief/surprise of propositions. They generalize Spohn's integer-valued κ-ranking functions (or natural conditional functions), introduced to model iterated revision of graded plain belief [Spo 88,90,08], as well as multiplicative real-valued possibility valuations [DuP 98]. Their value range carries a total order with endpoints and an additive structure for expressing conditionalization. We have exploited this concept in default reasoning [Wey 96,98,03], and belief revision [Wey 99,05].

Definition 2.1 (Ranking measures) *$R : \mathcal{B} \rightarrow \mathcal{V}$ is called a ranking measure (or $R \in \mathcal{R}_{\mathcal{V}}^{\mathcal{B}}$) iff*
1. $\mathcal{B} = (\mathbb{B}, \top, \bot, -, \cup, \cap)$ is a boolean algebra,
2. $\mathcal{V} = (V, +, 0, \infty, \leq)$ is a ranking algebra, i.e. the positive half of a totally ordered commutative group $\mathcal{G} = (G, +, 0, \leq)$ extended by ∞ s.t. for all $v \in V$, $v \leq v + \infty = \infty + v = \infty$ ($V = G^+ \cup \{\infty\}$),
3. $R(\top) = 0$, $R(\bot) = \infty$, $R(A \cup B) = \min_{\leq}\{R(A), R(B)\}$,
4. $R(\cup_{i \in I} A_i) = \infty$ if $\cup_{i \in I} A_i \in \mathbb{B}$ and $R(A_i) = \infty$ for all $i \in I$.
The conditional ranking measure associated with R is $R(.|.) : \mathbb{B} \times \mathbb{B} \rightarrow \mathcal{V}$, where $R(B|A) = R(B \cap A) - R(A)$ if $R(A) \neq \infty$, otherwise $R(B|A) = \infty$.

If $\mathcal{V}$ is non-trivial, i.e. $V \neq \{0, \infty\}$, which is required for modeling graded belief and iterated revision, then total orderedness implies that $0 < v < v + v < v + v + v < ... < \infty$ holds for each $0 < v < \infty$ ($v < w$ iff $v \leq w$ and $v \neq w$). The ranking algebra $\mathcal{V}_\kappa = (\mathbb{N} \cup \{\infty\}, +, 0, \infty, \leq)$ for Spohn's κ-ranking functions is the smallest non-trivial instance. $\mathcal{V}$ is said to be divisible iff it is non-trivial and each $v \in V$ can be written as the sum of $n + 1$ equal terms for every integer n. The smallest divisible ranking algebra is therefore $\mathcal{V}_{\kappa\pi} = (Rat^+ \cup \{\infty\}, +, 0, \infty, \leq)$ (Rat^+: set of the positive rational numbers). The real-valued possibilistic ranking algebra $\mathcal{V}_\pi^{real} = ([0, 1], \times, 1, 0, \geq)$ is also divisible. But its rational-valued substructure is not because Rat^+ is not closed under roots. However, every ranking algebra can be embedded into a divisible one. We can obtain a probabilistic interpretation by associating, e.g., each rank $0 < r \in \mathcal{V}_{\kappa\pi}$ to the non-standard probability ε^r, where $\varepsilon \neq 0$ is an arbitrary but fixed infinitesimal. This link allows the transfer of some powerful probabilistic tools, like entropy maximization.

Let R_0 be the uniform ranking measure with $R_0(A) = 0$ for all $A \neq \bot$. More generally, for $A \in \mathbb{B} - \{\top\}$, let $[A]_r : \mathbb{B} \rightarrow \mathcal{V}$ be the ranking measure given by $[A]_r(X) = r$ for $X \cap A = X \neq \bot$, and $[A]_r(X) = 0$ for $X \cap A \neq X$. Intuitively speaking, we obtain $[A]_r$ from R_0 by uniformly shifting A, or the A-worlds, by the amount r. We drop the index if $r = 1$. $A_1, \ldots, A_n$ are called independent w.r.t. R iff, for all $\vec{X}_i$ with $X_i \in \{A_i, -A_i\}$, $R(X_1 \cap \ldots \cap X_n) = R(X_1) + \ldots + R(X_n)$. A pseudo-ranking measure over $\mathcal{B}, \mathcal{V}$ is a function F on $\mathcal{B}$ verifying $F(A) = R(A) - v$, where $R \in \mathcal{R}_{\mathcal{V}}^{\mathcal{B}}$ is a ranking measure and $v \in V - \{\infty\}$. Let $\mathcal{V}^{\pm}$ be the extension of $\mathcal{V}$ to $] - \infty, \infty]_{\mathcal{G}}$. We define normalization (mapping pseudo-ranking measures to ranking measures) by $||F||(A) = F(A) - F(\top)$. Shifting is an important (pseudo-)ranking measure transformation for specifying Jeffrey/Spohn-conditionalization in the ranking framework. Given a pseudo-ranking measure R, shifting a proposition $A \in \mathbb{B}$ by the amount $r \in] - \infty, \infty]_{\mathcal{G}}$ means passing from R to $R + r[A]$, where

- $(R + r[A])(B) = min_{\leq}\{R(B \cap A) + r, R(B \cap \neg A)\}$, for all $B \in \mathbb{B}$.

In what follows, let $\mathcal{V}$ be a fixed divisible ranking algebra. Let L be a propositional language closed under the usual logical connectives, $\models$ a classical satisfaction relation for L, $\vdash$ the corresponding monotonic entailment relation, $[\![\varphi]\!] = \{m \mid m \models \varphi\}$, and $\mathcal{B}$ the associated boolean model set algebra with $\mathbb{B} = \{[\![\varphi]\!] \mid \varphi \in L\}$. On top of $(L, \models)$ we introduce a flat implicational language $L(\Rightarrow) = \{\varphi \Rightarrow_r \psi \mid 0 < r \in V, \varphi, \psi \in L\}$ to express graded conditional belief. Because ranking measures are well-suited to model belief states and belief change, we use them to interpret the belief conditionals with a minimization-friendly truth condition.

- $R \models_{rk} \varphi \Rightarrow_r \psi$ iff $R(\varphi \wedge \psi) + r \leq R(\varphi \wedge \neg\psi)$ (abbreviating $R(\chi) := R([\![\chi]\!])$).

Let $[\![\Delta]\!]_{rk} = \{R \in \mathcal{R} \mid R \models_{rk} \Delta\}$ be the model set of $\Delta \subseteq L(\Rightarrow)$. It follows from folklore that the corresponding monotonic entailment relation $\vdash_{rk}$ on $L(\Rightarrow)$ validates the rules of preferential logic and disjunctive rationality. But it violates rational monotony if $\mathcal{V}$ is divisible. We say that a proposition φ is believed to the degree r (at least) iff the rank/degree of surprise of $\neg\varphi$ is at least r. To specify plain belief we can fix a default threshold $0 < r_o < \infty$ and represent $Bel(\varphi)$ by $T \Rightarrow_{r_o} \varphi$. That is, it is possible to attribute different degrees of belief to φ and $\neg\varphi$ without supporting plain belief in φ or $\neg\varphi$. In $\mathcal{V}_{\kappa\pi}$, w.l.o.g., we may set $r_o = 1$ because all the values $v \neq 0, \infty$ are structurally indiscernible.

3 Ranking measure dynamics

A ranking revision system over $\mathcal{B}, \mathcal{V}$ is a quadruple $(\mathcal{R}, \mathcal{I}, \iota, \star)$ where $\mathcal{R} \subseteq \mathcal{R}_{\mathcal{V}}^{\mathcal{B}}$ is a set of ranking measures, $\mathcal{I}$ the collection of possible inputs, ι the input evaluation function mapping each pair $(R, i) \in \mathcal{R} \times \mathcal{I}$ to the set of admissible revision candidates $\iota(R,i) \subseteq \mathcal{R}$, and $\star : \mathcal{R} \times \mathcal{I} \rightarrow \mathcal{R}$ the revision function with $R \star i \in \iota(R,i)$. Because every ranking algebra can be embedded into a divisible one without affecting $\mathcal{R}$, we may assume w.l.o.g. that $\mathcal{V}$ is divisible. $\mathcal{B}$ is typically the boolean model set algebra of a classical background logic $(L, \models)$ closed under the usual propositional connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$.

Spohn's proposal for iterated propositional ranking revision is based on J(effrey)-conditionalization for κ-ranking functions and considers strength-parametrized propositional inputs $(\varphi, \alpha) \in L \times V$ [Spo 88,90]. Here we may set $\mathcal{R} = \mathcal{R}^{\mathcal{B}}_{\mathcal{V}_\kappa}$ and $\iota_J(\varphi, \alpha) = \{R \mid R(\varphi) = \infty \text{ or } R(\neg\varphi) = \alpha\}$. The revision step is realized by the uniform shifting of φ and/or $\neg\varphi$. In [Wey 96] we have proposed what we call minimal Spohn revision (only shifting as far as necessary), a more liberal variant in line with the minimal information change philosophy. Here the doxastic goal is only to believe φ to the degree α (at least), to be realized by minimal uniform shifting of $\varphi, \neg\varphi$. That is, $\iota_{msp}(\varphi, \alpha) = \{R \in \mathcal{R} \mid R(\varphi) = \infty \text{ or } R(\neg\varphi) \geq \alpha\}$. If φ is already believed to some degree α, weaker inputs (φ, β) with $\beta \leq \alpha$ are ignored, i.e. we assume redundancy by default.

The task of conditional ranking revision is to determine for each prior $R \in \mathcal{R}^{\mathcal{B}}_{\mathcal{V}}$, and any finite collection Δ of $\varphi \Rightarrow_r \psi$, a revised ranking measure $R \star \Delta \in \mathcal{R}$. Let $\Delta^{\rightarrow} = \{\varphi \rightarrow \psi \mid \varphi \Rightarrow_r \psi \in \Delta\}$ be the collection of material implications corresponding to the belief conditionals in Δ. If $R(\wedge\Delta^{\rightarrow}) \neq \infty$, i.e. if Δ is considered doxastically possible by R, then we should have $R \star \Delta \models_{rk} \Delta$. Accordingly, we set $\iota(R, \Delta) = [\![\Delta]\!]_{rk}$. If $R(\wedge\Delta^{\rightarrow}) = \infty$, we may stick to the prior and stipulate $\iota(R, \Delta) = \mathcal{R}$.

Our starting point is the minimal ranking construction philosophy which has been applied in default reasoning to obtain canonical preferred ranking measure models of default conditionals [Wey 98,03]. It tries to adapt the minimal information paradigm from probabilistic reasoning to the specificities of the ranking measure framework while keeping the flavour of System Z [Pea 90]. For conditional revision, it translates into the

Revision construction principle: For each prior R and input base $\Delta = \{\varphi_i \Rightarrow_{r_i} \psi_i \mid i \leq n\} \subseteq L(\Rightarrow)$, $R \star \Delta$ is obtained by iterated parametrized propositional revision with $\varphi_i \rightarrow \psi_i$ and $\varphi_i \rightarrow \neg\psi_i$, i.e. there are $x_i^-, x_i^+ \in V$ such that $R \star \Delta = R + \Sigma_{i \leq n} x_i^+ [\varphi_i \wedge \neg\psi_i] - \Sigma_{i \leq n} x_i^- [\varphi_i \wedge \psi_i]$ (where $\Sigma_{i \leq n} v_i[\chi_i]$ is an abbreviation for $v_0[[\![\chi_0]\!]] + \ldots + v_n[[\![\chi_n]\!]]$).

However, if we seek a revision function, it is not enough to uniformly minimize the x_i^-, x_i^+, because there may be infitely many minima. Furthermore, different shifting moves may have different informational costs or impacts. This suggests a hierarchical shifting procedure guided by the ranking constraints from Δ and the wish to minimize the shifting efforts. It should privilege the most relevant moves, e.g. those aiming at the most plausible target ranks, and proceed step by step, e.g. trying to maximize – lexicographically and bottom-up – the most plausible ranking layers. To determine the largest possible extension of a ranking layer, the most natural tool is

Relative ranking minimization: $R \star_{min} \Delta = Max_{\leq}\{Min_{\leq}\{R' \mid R \leq R', R' \models_{rk} \Delta\}, R\}$.

If there are models of Δ above R, $R \star_{min} \Delta$ is the unique least surprising one, otherwise, we may stay with R. This condition holds iff $R(\wedge\Delta^{\rightarrow}) = 0$. $(R \star_{min} \Delta)(\neg \wedge \Delta^{\rightarrow})$ is then the lowest rank which can be affected by Δ. We call it the top-active rank for R, Δ. The corresponding layer is the initial target of the ranking construction of $R \star \Delta$ from R.

Which propositions are we going to shift? The semantics of $\varphi \Rightarrow_r \psi$ invites us to shift upwards $\varphi \wedge \neg\psi$, and/or to shift downwards $\varphi \wedge \psi$. But upwards moves are preferable insofar as their information costs are lower in the probabilistic translation. Actually, if $R(\wedge\Delta^{\rightarrow}) = 0$, there is no need for contraction, i.e. downwards shifting. Then we may set $x_i^- = 0$ and only shift the $\varphi_i \wedge \neg\psi_i$.

The goal is to minimize the shifting efforts pushing the relevant shiftable propositions to their target rank. The idea here is to prefer local shifting constructions which minimize the longer moves, i.e. those carrying higher generic informational costs. This suggests the use of a lexicographic preference relation $\prec_{msh}$ comparing tuples of shifting lengths $\vec{x} = (x_i \mid i \leq m) \in V^{m+1}$. That is, over a given collection of propositions, one set of shifting moves should be preferable to another one iff, at the maximal r where their subsets of shifts of length $\geq r$ diverge, the subset of the first one is strictly included in that of the second one, reflecting a lower effort. Let $\vec{x}^{\geq r} = \{j \leq m \mid x_j \geq r\}$.

Definition 3.1 (Relative shifting effort) $\vec{x} \prec_{msh} \vec{y}$ *iff for the largest* r *with* $\vec{x}^{\geq r} \neq \vec{y}^{\geq r}$, $\vec{x}^{\geq r} \subset \vec{y}^{\geq r}$.

$(6,9,7,6) \prec_{msh} (7,9,7,5)$ because $\vec{x}^{\geq 9}, \vec{x}^{\geq 8} = \{1\} = \vec{y}^{\geq 8}, \vec{y}^{\geq 9}$, but $\vec{x}^{\geq 7} = \{1,2\} \subset \{0,1,2\} = \vec{y}^{\geq 7}$.

Theorem 3.2 (Shifting minimization) *If $R(\neg\delta_0 \wedge \ldots \wedge \neg\delta_m) = 0$ and $r < \infty$, then there is a unique $\prec_{msh}$-minimal $\vec{a_i}$ s.t. $(R + a_0[\delta_0] + \ldots + a_m[\delta_m])(\delta_j) \geq r$ for all $j \leq m$.*

If $r = \infty$, we set $a_i = \infty$. Uniqueness holds because for $\vec{a} \neq \vec{a'}$, $(\vec{a} + \vec{a'})/2 \prec_{msh} \vec{a}, \vec{a'}$. Another important minimal shifting feature is what we call justifiability. If proper shifting ($a_i > 0$) of $\varphi \wedge \neg\psi$ occurs to validate the ranking constraint $R(\varphi_i \wedge \psi_i) + r \leq R(\varphi_i \wedge \neg\psi_i)$, then it should be satisfied as an equality constraint. That is, the shifting should be minimal in the sense that the constraint is not over-satisfied.

We have now the ingredients to specify a hierarchical minimal construction strategy for building a ranking model of Δ. First we consider ranking expansion, i.e. conditional ranking revision for $R(\wedge\Delta^{\rightarrow}) = 0$. Our algorithm extends the JLZ-procedure for default reasoning [Wey 03] to non-uniform priors.

Minimal ranking expansion: $R, \Delta \mapsto R +_{mrr} \Delta$

Let $\Delta = \{\varphi_i \Rightarrow_{r_i} \psi_i \mid i \in I\}$ be finite and $\mathcal{V}$ be a divisible ranking algebra. To ensure local syntax independence while keeping the definitions transparent, w.l.o.g., we assume that $([\![\varphi_i \wedge \psi_i]\!], [\![\varphi_i \wedge \neg\psi_i]\!]) = ([\![\varphi_j \wedge \psi_j]\!], [\![\varphi_j \wedge \neg\psi_j]\!])$ implies $i = j$.

If $R(\wedge\Delta^{\rightarrow}) \neq 0$, we stipulate $R +_{mrr} \Delta = R$, observing that the result then necessarily fails to verify Δ. If $R(\wedge\Delta^{\rightarrow}) = 0$, the algorithm is based on an inductive bottom-up construction starting at the prior R and proceeding from more plausible to less plausible ranks, rank by rank, trying to approximate relative ranking minimization by local ranking constructions while minimizing the shifting efforts for each target rank. We begin with the top-active rank $(R \star_{min} \Delta)(\neg \wedge \Delta^{\rightarrow})$ for R and Δ. At the induction step, we consider the top-active rank for the ranking measure resulting from the preceding partial ranking construction and the collection of those conditionals which have not yet been settled, i.e. realized as an equality constraint. This means building two increasing sequences of ranking measures $(R_i)_{0<i\leq h}$ and $(R_i^*)_{0<i\leq h}$, with $R_i \leq R_{i+1}, R_i^* \leq R_{i+1}^*$, $R_i \leq R_i^*$, which are to converge to a ranking model $R^\star = R_h = R_h^*$ of Δ. We write:

- R_j : current ranking measure construction,
- R_j^* : current approximative ranking model of Δ: $R_j \star_{min} \Delta$,
- s_j : current target rank,
- I_j : indices of the shiftable propositions $[\![\varphi_i \wedge \neg\psi_i]\!]$ considered at level j,
- I'_j : indices of the $\varphi_i \Rightarrow_{r_i} \psi_i$ settled at level j,

Procedure for computing $R +_{mrr} \Delta$:

Induction start ($j = 1$) : $s_1 = 0,\ I_1 = I'_1 = \emptyset,\ R_1 = R,\ R_1^* = R_1 \star_{min} \Delta$ (with $R(\wedge\Delta^{\rightarrow}) = 0$).

Induction step ($j \rightarrow j+1$) :

- s_{j+1} smallest $s > s_j$ of the form $s = R_j^*(\varphi_i \wedge \neg\psi_i)$ for $i \in I - (I'_1 \cup \ldots \cup I'_j)$,
- $I_{j+1} = \{i \in I \mid R_j^*(\varphi_i \wedge \neg\psi_i) = s_{j+1}\} \subseteq I - (I'_1 \cup \ldots \cup I'_j)$,
- $R_{j+1} = R_j + \Sigma_{i \in I_{j+1}} a_i[\varphi_i \wedge \neg\psi_i]$ where $\vec{a}$ is the $\prec_{msh}$-minimal construction s.t., for all $i \in I_{j+1}$, $(R_j + \Sigma_{h \in I_{j+1}} a_i[\varphi_h \wedge \neg\psi_h] + \Sigma_{h \notin I_{\leq j+1}} \infty[\varphi_h \wedge \neg\psi_h])(\varphi_i \wedge \neg\psi_i) \geq s_{j+1}$, i.e. reaching s_{j+1} while ignoring the shiftable propositions with $R_j^*(\varphi_h \wedge \neg\psi_h) > s_{j+1}$,
- $R_{j+1}^* = R_{j+1} \star_{min} \Delta$,
- $I'_{j+1} = \{i \in I_{j+1} \mid R_{j+1}^*(\varphi_i \wedge \psi_i) + r_i = R_{j+1}^*(\varphi_i \wedge \neg\psi_i)\}$ ($\neq \emptyset$ if $I_{j+1} \neq \emptyset$)

Induction stop ($j \rightarrow$ *stop*) : If s_j does not exist, then $R +_{mrr} \Delta = R_j$ $(= R_{j-1})$.

Example: $R_0 +_{mrr} \{T \Rightarrow_1 \varphi\} +_{mrr} \{\psi \Rightarrow_2 \neg\varphi\} = 1[\neg\varphi] +_{mrr} \{\psi \Rightarrow_2 \neg\varphi\} = 1[\neg\varphi] + 3[\psi \wedge \varphi]$.

4 Minimal ranking revision

Proper conditional ranking revision is concerned with the passage from a prior R to a revised $R^* = R \star \Delta \models_{rk} \Delta$ when $R(\wedge\Delta^{\rightarrow}) \neq 0$. To implement the minimal ranking construction philosophy in a coherent way, we propose a two-step procedure inspired by the Levi identity. It starts with an auxiliary minimal contraction step to construct a suitable R^-_Δ verifying $R^-_\Delta(\wedge\Delta^\rightarrow) = 0$, and then applies minimal expansion to arrive at R^*.

The simplest proper revision task is to determine $R\star\{\varphi \Rightarrow_r \psi\}$ for $R(\varphi \rightarrow \psi) > 0$. The naive approach (virtual conditionalization) would be to proceed initially as for minimal ranking expansion, shifting upwards $\varphi\wedge\neg\psi$ as far as necessary, followed by normalization. As we will see, this may produce questionable results. An alternative is to try first to realize the precondition $R(\wedge\Delta^\rightarrow) = R(\varphi \rightarrow \psi) = 0$. If $R(\varphi \wedge \psi) \neq \infty$, the most parsimonous strategy may be to shift downwards $\varphi \wedge \psi$ until the precondition is met (contraction step), and then to apply minimal expansion to satisfy the ranking constraint. There are three scenarios.

If $0 < R(\varphi \rightarrow \psi) < \infty$ with $R(\varphi \wedge \psi) \neq \infty$, then $R(\varphi \wedge \neg\psi) = 0$, and we set

$$R \star \{\varphi \Rightarrow_r \psi\} = R - R(\varphi \wedge \psi)[\varphi \wedge \psi] + r[\varphi \wedge \neg\psi].$$

If $0 < R(\varphi \rightarrow \psi) < \infty$ with $R(\varphi \wedge \psi) = \infty$, we can only shift $\varphi \wedge \neg\psi$ to ∞, and normalize T to 0.

$$R \star \{\varphi \Rightarrow_r \psi\} = R - R(\varphi \rightarrow \psi)[T] + \infty[\varphi \wedge \neg\psi].$$

If $0 < R(\varphi \rightarrow \psi) = \infty$, actual revision is blocked, inviting us to stipulate $R \star \{\varphi \Rightarrow_r \psi\} = R$.

We can transfer this strategy also to proper revision with multiple conditionals. Let $\Delta = \{\varphi_i \Rightarrow_{r_i} \psi_i \mid i \leq n\}$. If $R(\wedge\Delta^\rightarrow) = \infty$, the simplest approach is to ignore the impossible input and set $R \star \Delta = R$. If $0 < R(\wedge\Delta^\rightarrow) < \infty$, the idea is again to first transform R into an appropriately contracted (pseudo-)ranking measure R^-_Δ verifying $R^-_\Delta(\wedge\Delta^\rightarrow) = 0$, and then to use $+_{mrr}$ to expand R^-_Δ with Δ. Contraction, aimed at lowering support for the $\varphi_i \Rightarrow_{r_i} \neg\psi_i$ (opposing the $\varphi_i \Rightarrow_{r_i} \psi_i$), proceeds most naturally by uniformly shifting downwards the $\varphi_i \wedge \psi_i$ (i.e. $\varphi_i \wedge \neg\neg\psi_i$) until some $\wedge\Delta^\rightarrow \wedge (\varphi_i \wedge \psi_i)$ hits 0. A problem may just arise if $R(\wedge\Delta^\rightarrow \wedge (\varphi_i \wedge \psi_i)) = \infty$ for all $i \leq n$. Then the only admissible solution is to shift the $\varphi_i \wedge \neg\psi_i$ to ∞, and normalize T to 0.

$$R \star \Delta = R - R(\wedge\Delta^\rightarrow)[T] + \infty[\neg \wedge \Delta^\rightarrow].$$

Without this degeneration, the uniform parallel shifting of those $\varphi_i \wedge \psi_i$ with $R(\wedge\Delta^\rightarrow \wedge (\varphi_i \wedge \psi_i)) \neq \infty$ to obtain R^-_Δ from R prevents biasedness and minimizes the maximal necessary shifting length. In the second step, we can now again apply minimal ranking expansion with Δ to transform R^-_Δ into $R \star \Delta$. Although R^-_Δ may well be a non-normalized pseudo-ranking measure, $R \star \Delta = R^-_\Delta +_{mrr} \Delta$ will always be a ranking measure. In fact, by construction, $\wedge\Delta^\rightarrow$ will get and keep rank 0, and $\neg \wedge \Delta^\rightarrow$ is just the disjunction of the $\varphi_i \wedge \neg\psi_i$, which Δ forces to be shifted above 0.

Auxiliary ranking contraction: $R, \Delta \mapsto R -_{mrr} \Delta = R^-_\Delta$

Let $\Delta = \{\varphi_i \Rightarrow_{r_i} \psi_i \mid i \in I\}$ be finite and $R(\wedge\Delta^\rightarrow) < \infty$. The role of $-_{mrr}$ is to realize in a minimal way the precondition $R(\wedge\Delta^\rightarrow) = 0$ for applying $+_{mrr}$. Its algorithm is based on a top-down construction which aims at transforming R with minimal shifting efforts into a contracted (pseudo-)ranking measure R^-_Δ validating $R^-_\Delta(\wedge\Delta^\rightarrow) = 0$, and thereby paves the way for $+_{mrr}$. Let Δ^-_R be the collection of conditionals from Δ which are doxastically consistent with $\wedge\Delta^\rightarrow$ in R:

- $\Delta^-_R = \{\varphi_i \Rightarrow_{r_i} \psi_i \in \Delta \mid R(\wedge\Delta^\rightarrow \wedge (\varphi_i \wedge \psi_i)) \neq \infty\}$

Procedure for computing $R -_{mrr} \Delta$**:**

- If $\Delta^-_R = \emptyset$, then $R -_{mrr} \Delta = R - R(\wedge\Delta^\rightarrow)[T] + \infty[\neg \wedge \Delta^\rightarrow]$
- If $\Delta^-_R \neq \emptyset$, then $R -_{mrr} \Delta = R - \Sigma_{\varphi_i \Rightarrow_{r_i} \psi_i \in \Delta^-_R}\ \alpha[\varphi_i \wedge \psi_i]$

  where α is minimal such that $(R - \Sigma_{\varphi_i \Rightarrow_{r_i} \psi_i \in \Delta^-_R}\ \alpha[\varphi_i \wedge \psi_i])(\wedge\Delta^\rightarrow) = 0$

It follows from the finiteness of Δ, the divisibility of $\mathcal{V}$, and the character of the transformations, that $-_{mrr}$ is well-defined. We can now specify our proposal for multiple conditional ranking revision. It extends propositional minimal ranking revision, i.e. $R \star_{mrr} \{T \Rightarrow_r \varphi\} = R \star_{msp} (\varphi, r)$.

Definition 4.1 (Minimal ranking revision) *Let $R \in \mathcal{R}^{\mathcal{B}}_{\mathcal{V}}$, $\mathcal{V}$ be divisible, and $\Delta \subseteq L(\Rightarrow)$ be finite.*

- *If $R(\wedge\Delta^\rightarrow) \neq \infty$, then* $R \star_{mrr} \Delta = R -_{mrr} \Delta +_{mrr} \Delta$

- *If* $R(\wedge\Delta^{\rightarrow}) = \infty$, *then* $R \star_{mrr} \Delta = R$

Some illustrative examples

1. $\Delta_1 = \{T \Rightarrow_2 \neg\varphi, T \Rightarrow_1 \neg\psi\}$ and $\Delta_2 = \{T \Rightarrow_1 \varphi, T \Rightarrow_1 \psi\}$, for logically independent $\varphi, \psi \in L$.

- $R_1 = R_0 \star_{mrr} \Delta_1 = R_0 +_{mrr} \Delta_1 = 2[\varphi] + 1[\psi]$
- $R_1(\wedge\Delta_2^{\rightarrow}) = R_1((T \rightarrow \varphi) \wedge (T \rightarrow \psi)) = R_1(\varphi \wedge \psi) = 3 > 0$
- $R_1 -_{mrr} \Delta_2 = R_1 - 3/2[\varphi] - 3/2[\psi] = 1/2[\varphi \wedge \neg\psi] - 1/2[\psi \wedge \neg\varphi]$, with $(R_1 -_{mrr} \Delta_2)(\varphi \wedge \psi) = 0$
- $R_1 \star_{mrr} \Delta_2 = R_1 -_{mrr} \Delta_2 +_{mrr} \Delta_2 =$

$-1/2[\psi \wedge \neg\varphi] + 1/2[\varphi \wedge \neg\psi] + 3/2[\neg\varphi] + 1/2[\neg\psi] = 1[\neg\varphi] + 1[\neg\psi]$

2. $\Delta_1 = \{T \Rightarrow_1 \varphi, \varphi \Rightarrow_1 \psi\}$ and $\Delta_2 = \{\varphi \Rightarrow_1 \neg\psi\}$. Difference with virtual conditionalization.

- $R_1 = R_0 \star_{mrr} \Delta_1 = R_0 +_{mrr} \Delta_1 = 1[\neg\varphi] + 1[\varphi \wedge \neg\psi]$
- $R_1 -_{mrr} \Delta_2 = R_1 - 1[\varphi \wedge \neg\psi] = 1[\neg\varphi]$
- $R_1 \star_{mrr} \Delta_2 = R_1 -_{mrr} \Delta_2 +_{mrr} \Delta_2 = 1[\neg\varphi] + 1[\varphi \wedge \psi]$

The auxiliary contraction strategy is illustrated by the following revision configurations.

3. Let $\Delta_1 = \{\varphi \Rightarrow_2 \psi, \neg\varphi \Rightarrow_1 \psi\}$ and $\Delta_2 = \{\varphi \Rightarrow_1 \neg\psi, \neg\varphi \Rightarrow_1 \neg\psi\}$.

- $R_1 = R_0 \star_{mrr} \Delta_1 = R_0 +_{mrr} \Delta_1 = 2[\varphi \wedge \neg\psi] + 1[\neg\varphi \wedge \neg\psi]$
- $R_1 -_{mrr} \Delta_2 = 1[\varphi \wedge \neg\psi]$ where $\wedge\Delta_2^{\rightarrow} \dashv\vdash \neg\psi$
- $R_1 \star_{mrr} \Delta_2 = 1[\varphi \wedge \neg\psi] + 2[\varphi \wedge \psi] + 1[\neg\varphi \wedge \psi]$

4. Let $\Delta_1 = \{T \Rightarrow_1 \varphi, \neg\varphi \Rightarrow_4 \psi\}$ and $\Delta_2 = \{(\psi \rightarrow \varphi) \Rightarrow_1 \neg\varphi, \neg\varphi \Rightarrow_1 \psi\}$.

- $R_1 = R_0 \star_{mrr} \Delta_1 = R_0 +_{mrr} \Delta_1 = 1[\neg\varphi] + 4[\neg\varphi \wedge \neg\psi]$
- $R_1 -_{mrr} \Delta_2 = 5[\neg\varphi \wedge \neg\psi]$ where $\wedge\Delta_2^{\rightarrow} \dashv\vdash \neg\varphi \wedge \psi$
- $R_1 \star_{mrr} \Delta_2 = 6[\varphi] + 5[\neg\varphi \wedge \neg\psi]$

The iteration of conflicting conditional evidence may increase disbelief in the condition itself.

5. Let $\Delta_1 = \{\varphi \Rightarrow_1 \psi\}, \Delta_2 = \{\varphi \Rightarrow_1 \neg\psi\}, \Delta_3 = \{\varphi \Rightarrow_1 \psi\}$.

- $R_0 \star_{mrr} \Delta_1 = 1[\varphi \wedge \neg\psi] \not\models_{rk} T \Rightarrow_1 \neg\varphi$
- $R_0 \star_{mrr} \Delta_1 \star_{mrr} \Delta_2 = 1[\varphi \wedge \neg\psi] + 2[\varphi \wedge \psi] \models_{rk} T \Rightarrow_1 \neg\varphi$
- $R_0 \star_{mrr} \Delta_1 \star_{mrr} \Delta_2 \star_{mrr} \Delta_3 = 3[\varphi \wedge \neg\psi] + 2[\varphi \wedge \psi] \models_{rk} T \Rightarrow_2 \neg\varphi$

5 Properties and relationships

How does $\star_{mrr}$ handle existing and new desiderata for iterated revision[1]? To discuss principles for propositional revision, we may translate $R \star \varphi$ by $R \star_{mrr} T \Rightarrow_1 \varphi$. The classical axioms for iterated revision (in the epistemic state formulation) **C.1 - C.4** [DP 97] are valid, whereas Lehmann's postulates **I.5, I.6** [Leh 95] are not. In fact, $\star_{mrr}$ neither prevents $R \star \varphi \star \varphi \wedge \psi \neq R \star \varphi \wedge \psi$, nor $R \star \varphi \star \psi \neq R \star \varphi \star \varphi \wedge \psi$ if $R \star \varphi \not\models_{rk} T \Rightarrow_1 \neg\psi$. Because we can have $R \not\models_{rk} \psi \Rightarrow_1 \neg\varphi$ and $R \star \varphi \not\models_{rk} \psi \Rightarrow_1 \varphi$, the Ind-axiom [JT 07] also fails.

From Kern-Isberner's principles for iterated revision with single, non-parametrized conditionals[KI 99, 02, 04], $\star_{mrr}$ validates only **CR.0, 4, 6, 7**, whereas the other requirements don't hold. Many failures are linked to our threshold semantics for belief and the doxastic possibility requirement, which e.g. restricts success. On the other hand, $\star_{mrr}$ satisfies our following postulates, which encode shifting minimality for single conditional revision.

CRM.1 If $R(\varphi \rightarrow \psi) \neq \infty$, then $R \star \varphi \Rightarrow_r \psi \models_{rk} \varphi \Rightarrow_r \psi$, else $R \star \varphi \Rightarrow_r \psi = R$ (Cautious success)

CRM.2 If $R \models_{rk} \varphi \Rightarrow_r \psi$, then $R \star \varphi \Rightarrow_r \psi = R$ (Redundancy)

[1] For the sake of readability, we may drop the set parentheses for single inputs.

CRM.3 If $R \not\models_{rk} \varphi \Rightarrow_r \psi$ and $r < \infty$,
then $R \star \varphi \Rightarrow_r \psi \not\models_{rk} \varphi \Rightarrow_{r+x} \psi$ for $x > 0$ (Minimal construction)

CRM.4 If $R(\varphi \wedge \neg\psi) > 0$, or $R(\varphi \wedge \psi) = 0$, or $R(\neg\varphi) = 0$, or $R(\varphi \wedge \psi) = \infty$,
then $(R \star \varphi \Rightarrow_r \psi)(\varphi \wedge \psi) = R(\varphi \wedge \psi)$ (Confirmation stability)

CRM.5 If $R(\varphi \wedge \neg\psi) = 0$, $R(\neg\varphi) > 0$, and $R(\varphi \wedge \psi) < \infty$,
then $(R \star \varphi \Rightarrow_r \psi)(\varphi \wedge \psi) = 0$ (Local contraction)

Zhang [Zha 04] and Delgrande, Jin [DJ 08] have advanced rationality postulates for parallel propositional revision. But their single step principles $\mathbf{K} \otimes P$ and $\mathbf{K} \otimes C$ are invalid for $\star_{mrr}$. The generalization of the DP-axioms to multiple revision holds for **C1**$^{\otimes}$**, C3**$^{\otimes}$**, C4**$^{\otimes}$, but fails for **C2**$^{\otimes}$. There are also counterexamples for the alternatives **Ind**$^{\otimes}$ and **Ret**$^{\otimes}$ [DJ 08]. Kern-Isberner has proposed five axioms **CSR.1 - CSR.5** for conditional set revision [KI 08].

- **CSR.1 (Success)** $R \star \Delta \models_{rk} \Delta$
- **CSR.2 (Stability)** $R \models_{rk} \Delta$ implies $R \star \Delta = R$
- **CSR.3 (Semantic Equivalence)** $\Delta \dashv\vdash_{rk} \Delta'$ implies $R \star \Delta = R \star \Delta'$
- **CSR.4 (Reciprocity)** $R \star \Delta \models_{rk} \Delta'$ and $R \star \Delta' \models_{rk} \Delta$ implies $R \star \Delta = R \star \Delta'$
- **CSR.5 (Logical coherence)** $R \star (\Delta \cup \Delta') = R \star \Delta \star (\Delta \cup \Delta')$.

For $\star_{mrr}$, Success presupposes $R(\wedge\Delta^{\rightarrow}) \neq \infty$. Stability is obvious. But **CSR.3, 4, 5** fail because they induce global semanticality, which blocks desirable forms of independence reasoning (see exceptional inheritance paradox [Wey 03]). In fact, there are $[\![\Delta]\!]_{rk} = [\![\Delta']\!]_{rk}$ with $R \star \Delta \neq R \star \Delta'$. What we have is

Local semanticality: $\{[\![\delta]\!]_{rk} \mid \delta \in \Delta\} = \{[\![\delta]\!]_{rk} \mid \delta \in \Delta'\}$ implies $R \star \Delta = R \star \Delta'$.

That is, while we take the individual ranking constraints – the local semantic content – seriously, their exact syntactic form should be irrelevant. Furthermore, the only reason for $R \star \Delta$ not to satisfy Δ should be the doxastic impossibility of $\wedge\Delta^{\rightarrow}$.

Cautious success: If $R(\wedge\Delta^{\rightarrow}) \neq \infty$, then $R \star \Delta \models_{rk} \Delta$, else $R \star \Delta = R$.

The following desirable and powerful principles, which are derived from a characterization of cross entropy minimization [SJ 80], a distinguished variant of probabilistic revision, are also valid for $\star_{mrr}$.

Invariance. For any boolean automorphism $\pi : \mathcal{B}_L \rightarrow \mathcal{B}_L$ and any bijection $\pi' : L \rightarrow L$ with $\pi([\![\varphi]\!]) = [\![\pi'(\varphi)]\!]$, if $\Delta^{\pi'} = \{\pi'(\varphi) \Rightarrow_r \pi'(\psi) \mid \varphi \Rightarrow_r \psi \in \Delta\}$ and $R^{\pi} = R \circ \pi^{-1}$, then $(R \star \Delta)^{\pi} = R^{\pi} \star \Delta^{\pi'}$.

System independence. If L_1, L_2 have disjoint non-logical vocabularies, $\Delta_i \subseteq L_i(\Rightarrow)$, $dom(R_i) = \mathcal{B}_{L_i}$, R_0^i is the uniform ranking measure on $\mathcal{B}_{L_i}$, and $R_i(\wedge\Delta_i^{\rightarrow}) \neq \infty$,
then $(R_1 \times R_2) \star (\Delta_1 \cup \Delta_2) = (R_1 \star \Delta_1) \times R_0^2 + R_0^1 \times (R_2 \star \Delta_2)$.

What do these results tell us? First, we can relativize many failures because the corresponding principles are easily adaptable to our more general perspective (threshold belief strength, divisibility, doxastic impossibility). But some violations are essential – they follow from the minimal information philosophy and the wish to handle independence information in an intuitive way. Note that for complex iterated revision, a simple intuitive axiomatic characterization (like for AGM) may well be elusive. What is however important is the overall plausibility of the semantic revision framework.

Our earlier proposals for iterated conditional revision within the ranking framework [Wey 99, Wey 05] are more cumbersome and differ from $\star_{mrr}$ because they rely on virtual conditionalization for revision. However, this naive approach conflicts with the minimal information paradigm. Consider for instance $R = R_0 \star \{T \Rightarrow_1 \varphi, \varphi \Rightarrow_1 \psi\} \star \{\varphi \Rightarrow_1 \neg\psi\}$. Here virtual conditionalization gives us $R = 1[\varphi \wedge \psi]$, whereas $\star_{mrr}$ supports $1[\neg\varphi] + 1[\varphi \wedge \psi]$, which seems more natural because there is no unmotivated impact on φ.

Another possibility is to apply cross-entropy minimization (MCE) by exploiting a translation between ranking measures and non-standard probability [Wey 95b,03]. But this move introduces free parameters whose naive choice may violate invariance properties. Furthermore, there is no ranking construction algorithm to compute the MCE-revised ranking measure for arbitrary ranking constraints. The exisiting proposals, e.g. [BP 03], force the user to fix ranking values in an ad hoc way. Nevertheless, for "non-entangled" examples, the standard cross-entropy-based approach $\star_{mce}$ produces the same results as $\star_{mrr}$.

Kern-Isberner [KI 99,02,04,08] proposes a number of postulates for conditional set revision, but she does not offer a specific proposal comparable to $\star_{mrr}$. Another limitation is her focus on disrete-valued ranking measures – information minimization requires divisibility [Wey 03]. In [KI 99] she suggests a revision function for integer-valued ranking measures and single non-parametrized conditionals. But there are simple examples where her approach diverges from $\star_{mrr}$ and $\star_{mce}$. For instance, these approaches both support $R_0 \star T \Rightarrow_1 \varphi \star \psi \Rightarrow_1 \neg\varphi = R_0 + 1[\neg\varphi] + 2[\varphi \wedge \psi] = R_0 \star \{T \Rightarrow_1 \varphi, \psi \Rightarrow_1 \neg\varphi\} \models_{rk} T \Rightarrow_1 \varphi, \psi \Rightarrow_1 \neg\varphi$. This seems very reasonable given that the two conditionals are consistent. However, $R_0 \star_{ki} T \Rightarrow_1 \varphi \star_{ki} \psi \Rightarrow_1 \neg\varphi = R_0 + 1[\neg\varphi \wedge \neg\psi] + 1[\varphi \wedge \psi] \not\models_{rk} T \Rightarrow_1 \varphi$. This points to $\star_{mrr}$ as a promising proposal for multiple conditional set revision.

[AGM 85] C.E. Alchourron, P. Gr̈denfors, D. Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. Journal of Symboloic Logic 50(2): 510-530, 1985.

[BG 93] C. Boutilier, M. Goldszmidt. Revision by Conditional Beliefs. AAAI 1993: 649-654, 1993.

[Bou 96] C. Boutilier. Iterated revision and minimal revision of conditional beliefs. Journal of Philosophical Logic 25: 262 304, 1996.

[BP 03] R.A. Bourne, S. Parsons: Extending the Maximum Entropy Approach to Variable Strength Defaults. Annals of Mathematics and Artificial Intelligence 39(1-2): 123-146, 2003.

[DP 97] A. Darwiche, J. Pearl. On the logic of iterated belief revision. AIJ 89(1-2):1-29, 1997.

[DJ 08] J.P. Delgrande, Y. Jin: Parallel Belief Revision. AAAI 08: 430-435, 2008.

[DuP 98] D. Dubois, H. Prade. Possibility theory : qualitative and quantitative aspects. Quantified Representation of Uncertainty and Imprecision, Handbook on Defeasible Reasoning and Uncertainty Management Systems - Vol. 1 (ed. P. Smets): 169-226. Kluwer, 1998.

[JT 07] Y. Jin, M. Thielscher: Iterated belief revision, revised. Artificial Intelligence. 171(1): 1-18, 2007.

[KI 99] G. Kern-Isberner. Postulates for conditional belief revision. IJCAI 99: 186 - 191. Morgan Kaufmann, 1999.

[KI 02] G. Kern-Isberner: Handling conditionals adequately in uncertain reasoning and belief revision. Journal of Applied Non-Classical Logics 12(2): 215-237, 2002.

[KI 04] G. Kern-Isberner: A Thorough Axiomatization of a Principle of Conditional Preservation in Belief Revision. Annals of Mathematics and Artificial Intelligence 40(1-2): 127-164, 2004.

[KI 08] G. Kern-Isberner: Linking Iterated Belief Change Operations to Nonmonotonic Reasoning. KR 2008: 166-176, 2008. [Leh 95] D.J. Lehmann: Belief Revision, Revised. IJCAI 1995: 1534-1540, 1995.

[Par 94] J. Paris 94. The uncertain reasoners companion. Cambridge University Press, 1994.

[Pea 90] J. Pearl. System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. TARK 3: 121-135. Morgan Kaufmann, 1990.

[SJ 80] J.E. Shore, R.W. Johnson. Axiomatic derivation of the principle of cross-entropy minimization. IEEE Transactions on Information Theory, IT-26(1):26-37, 1980.

[Spo 88] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. Causation in Decision, Belief Change, and Statistics (eds. W.L. Harper, B. Skyrms): 105-134. Kluwer, 1988.

[Spo 90] W. Spohn. A general non-probabilistic theory of inductive reasoning. Uncertainty in Artificial Intelligence 4 (eds. R.D. Shachter et al.): 149-158. North-Holland, Amsterdam, 1990.

[Spo 08] W. Spohn. A Survey of Ranking Theory. Degrees of Belief. An Anthology (eds. F. Huber, C. Schmidt-Petri): 185.228. Oxford University Press, 2008.

[Wey 94] E. Weydert. General belief measures. UAI 94: 575-582. Morgan Kaufmann, 1994.

[Wey 95b] E. Weydert. Defaults and infinitesimals. Defeasible inference by non-archimdean entropy maximization. UAI 95: 540-547. Morgan Kaufmann, 1995.

[Wey 96] E. Weydert. System J - Revision entailment. FAPR 96: 637-649. Springer, 1996.

[Wey 98] E. Weydert. System JZ - How to build a canonical ranking model of a default knowledge base. KR 98: 190-201. Morgan Kaufmann, 1998.

[Wey 99] E. Weydert. JZBR - Iterated belief change for conditional ranking constraints. Spinning Ideas: Electronic essays dedicated to Peter Gaerdenfors on his 50th Birthday, 1999. http://www.lucs.lu.se/spinning/categories/dynamics/Weydert/index.html.

[Wey 03] E. Weydert. System JLZ - Rational default reasoning by minimal ranking constructions. Journal of Applied Logic 1(3-4):273308. Elsevier, 2003.

[Wey 05] E. Weydert. Projective default epistemology - a first look. Conditionals, Information, and Inference, LNAI 3301: 65-85. Springer, 2005.

[ZF 01] D. Zhang, N. Foo. Infinitary belief revision. Journal of philosophical logic 30(6): 525-570, 2001.

[Zha 04] D. Zhang. Properties of Iterated Multiple Belief Revision. LPNMR 2004: 314-325, 2004.