

Interval Orders, Biorders and Credibility-limited Belief Revision

Richard Booth¹, Ivan Varzinczak^{2,3,4}

¹Cardiff University, United Kingdom

²Université Sorbonne Paris Nord, Inserm, Sorbonne Université, Limics, 93017 Bobigny, France

³CAIR, University of Cape Town, South Africa

⁴ISTI-CNR, Pisa, Italy

boothr2@cardiff.ac.uk, ivan.varzinczak@sorbonne-paris-nord.fr

Abstract

Rational belief revision is commonly viewed as being based on a preference order between possible worlds, with the resulting new belief set being those sentences true in all the most preferred models of the incoming new information. Usually, such a preference order is taken to be a total preorder. Nevertheless, there are other, more general classes of ordering that can also be employed. In this paper, we explore two such classes that have been studied within the theory of rational choice but have seen limited or no application in belief revision. We begin with *interval orders*, introduced by Fishburn in the '80s, which associate with each possible world a nonnegative 'interval' of plausibility. We then move on to *biorders*, studied by Aleskerov, Bouyssou, and Monjardet, which generalise interval orders by allowing the intervals to have negative lengths, a feature that can be used to capture a notion of dissonance or instability. We provide axiomatic characterisations of these two resulting families of belief revision operators, as well as of two further families of interest that lie between interval orders and biorders. We show that while biorder-based revisions satisfy the Success postulate, they do not always yield consistent outputs. By modifying their definition to discard inputs that lead to inconsistency as 'incredible', we derive new families of so-called non-prioritised revision that satisfy the Consistency postulate, but not the Success one. These families are linked to credibility-limited revision operators of Hansson et al., but for which the set of credible sentences does not satisfy the single-sentence closure condition. We argue that the biorder-based approach is well-suited for scenarios where an agent might initially reject new information, but may accept it when presented with additional explanation.

1 Introduction and Motivation

Belief revision is concerned with the incorporation of any new item of information into an agent's belief set in a principled way. The so-called AGM-approach to belief revision (Alchourrón, Gärdenfors, and Makinson 1985) does so by formulating a formal account of rationality in this setting, usually characterised by a set of postulates capturing the expected behaviour of revision. Traditionally, two pillars of the AGM approach are the postulates called Success and Consistency, which ensure that the new incoming information is always accepted and the result is a consistent belief set (whenever the incoming information is consistent).

As it turns out, the AGM approach assumes an idealised setting, which is not suitable for modelling the behaviour of realistic rational agents. As a result, the past decades have seen many attempts to bring more realism into it. This is one of the motivations underlying the present paper.

Here, we argue that the requirement for a consistent revision result could be dropped. Realistic agents may face new information that is *destabilising*, represented as sentences that just do not cohere with the agent's current beliefs. As an example, adapted from Fermé (1999), upon learning "My five-year-old nephew had lunch with King Charles today", the agent's beliefs get destabilised and switch to 'maintenance mode', where some form of 'repair', which can take the form of further clarification, is needed.

Alternatively, if we prize consistency, we may drop the blind acceptance of new incoming information, specifically of problematic sentences. This amounts to giving up the Success postulate and is the stance in non-prioritised revision (NPR). In the example above, the sentence is just ignored or 'put on hold' until further clarification or explanation, if any, might make it become acceptable. For instance, by learning that "King Charles visited a local school today", the agent is now keen to accept the first claim above.

In the present paper, we tackle both issues above and their interplay, investigating the consequences but mainly the fruitfulness of dropping either Success or Consistency. Our main tools are a family of orderings on possible worlds, more general than the total preorders usually assumed in the AGM approach. More specifically, we look at *interval orders* (Fishburn 1985), which have already been utilised in the context of belief change, namely for belief contraction (Rott 2014). But there is a more general type of ordering that has been looked at in economics, and which can result in an empty choice set, namely *biorders*, as studied most extensively by Aleskerov et al. (2007). We bring these into the landscape of belief change, which is part of the novelty of the present work.

The remainder of the paper is structured as follows: In Section 2, we give the formal background for the upcoming sections. In Section 3, we present our first contribution by completing the picture initiated by Rott (2014), namely shifting the focus from total preorders in the definition of belief change operators to interval orders. We then generalise our constructions to the broader class of biorders, showing

in particular that intervals of negative length can be used to capture an alternative form of impossibility, which we refer to as dissonance or instability (Section 4). One of the consequences of this generalisation is the loss of the Consistency postulate. Instead of throwing in the towel at this stage, we try to ‘remedy’ the situation by exploring two subclasses of biorders satisfying a form of transitivity and that lie between interval orders and biorders. These strengthen the constructions and results for the general biorder case, but do not bring back Consistency. This serves as a motivation for an investigation of non-prioritised revision (NPR) grounded on the operators previously defined (Section 5), in which inputs leading to inconsistency are discarded as ‘incredible’. Given that, in Section 6, we show how our two subclasses of NPR can be formulated as credibility-limited revision operators. Before concluding the paper, we discuss some relevant related work (Section 7). Proofs of all stated results are available in a longer version of the paper at <https://arxiv.org/abs/2604.27156>.

2 Formal Preliminaries and Conventions

In this section, we provide the required formal background for the remainder of the paper.

2.1 Propositional Notation

Let $\mathcal{P}$ be a finite and non-empty set of propositional *atoms*. We use $p, q, \dots$ as meta-variables for atoms. Propositional sentences are denoted by $\alpha, \beta, \dots$, and are recursively defined in the usual way. With $\mathcal{L}$ we denote the set of all propositional sentences.

With $\mathcal{U} \stackrel{\text{def}}{=} \{0, 1\}^{\mathcal{P}}$, we denote the set of all propositional *valuations* on $\mathcal{P}$, where 1 represents truth and 0 falsity. We use $v, u, \dots$, possibly with primes, to denote valuations. Whenever it eases the presentation, we shall represent valuations as sequences of atoms (e.g., p) and barred atoms (e.g., $\bar{p}$), in alphabetical order, with the understanding that the presence of a non-barred atom indicates that the atom is ‘true’ in the valuation, while the presence of a barred atom indicates that the atom is ‘false’ in the valuation.

Satisfaction of $\alpha \in \mathcal{L}$ by $v \in \mathcal{U}$ is defined in the usual truth-functional way and is denoted by $v \models \alpha$. The set of *models* of a sentence α is defined as $\llbracket \alpha \rrbracket \stackrel{\text{def}}{=} \{v \in \mathcal{U} \mid v \models \alpha\}$. We say sentences α and β are *logically equivalent*, denoted $\alpha \equiv \beta$, if $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket$. We say α is *consistent* if $\llbracket \alpha \rrbracket \neq \emptyset$. We call α a *tautology*, denoted $\models \alpha$, if $\llbracket \alpha \rrbracket = \mathcal{U}$.

The notions above are extended to (possibly infinite) sets $X \subseteq \mathcal{L}$ in the usual way, given $\llbracket X \rrbracket \stackrel{\text{def}}{=} \bigcap_{\alpha \in X} \llbracket \alpha \rrbracket$. We say X (classically) *entails* $\alpha \in \mathcal{L}$, denoted $X \models \alpha$, if $\llbracket X \rrbracket \subseteq \llbracket \alpha \rrbracket$. Given $X \subseteq \mathcal{L}$, with $Cn(X) \stackrel{\text{def}}{=} \{\alpha \mid X \models \alpha\}$ we denote the set of all *logical consequences* of X .

Finally, in the remainder of the paper, we will often be interested in selecting the models of a sentence or those of a set of sentences according to some well-defined criteria (e.g. a notion of closeness w.r.t. the models one started off with, such as those of an initial knowledge base). Let $\mathcal{V} \subseteq \mathcal{U}$, and let $\leq \subseteq \mathcal{U} \times \mathcal{U}$ be an ordering (not necessarily a strict one). The *optimal elements* of $\mathcal{V}$ w.r.t. $\leq$ are defined as:

$$\text{opt}(\mathcal{V}, \leq) \stackrel{\text{def}}{=} \{v \in \mathcal{V} \mid v \leq u \text{ for all } u \in \mathcal{V}\} \quad (1)$$

We shall say that the ordering $\leq$ is *anchored on* $\mathcal{V}$ if $\mathcal{V} = \text{opt}(\mathcal{U}, \leq)$, i.e., $\mathcal{V}$ is the set of optimal elements overall, and that $\leq$ is *anchored* if furthermore $\mathcal{V} \neq \emptyset$.

2.2 Belief Revision

Belief change is concerned with the problem of modifying an agent’s beliefs in a principled way. The AGM paradigm for belief change (Alchourrón, Gärdenfors, and Makinson 1985) has become the gold standard in the area. It assumes an underlying logic with a propositional language, as $\mathcal{L}$, and the above consequence operator $Cn(\cdot)$.

A *belief set* $\mathcal{KB}$ is a *deductively closed* set of sentences, i.e., $\mathcal{KB} = Cn(\mathcal{KB})$. The *revision* of $\mathcal{KB}$ with α corresponds to a consistent resulting belief set $\mathcal{KB} * \alpha$ from which α follows. In the AGM paradigm, the expected behaviour of the revision operator is characterised in terms of the postulates below (we assume w.l.o.g. that the input $\mathcal{KB}$ is *consistent*)¹:

- P0** $\mathcal{KB} * \top = \mathcal{KB}$ (Tautology)
- P1** $\mathcal{KB} * \alpha = Cn(\mathcal{KB} * \alpha)$ (Closure)
- P2** $\alpha \in \mathcal{KB} * \alpha$ (Success)
- P3** $\mathcal{KB} * \alpha \subseteq Cn(\mathcal{KB} \cup \{\alpha\})$ (Inclusion)
- P4** If $\neg\alpha \notin \mathcal{KB}$, then $Cn(\mathcal{KB} \cup \{\alpha\}) \subseteq \mathcal{KB} * \alpha$ (Vacuity)
- P5** If $\alpha \equiv \beta$, then $\mathcal{KB} * \alpha = \mathcal{KB} * \beta$ (Extensionality)
- P6** If $\alpha \not\models \perp$, then $\mathcal{KB} * \alpha \not\models \perp$ (Consistency)
- P7** $\mathcal{KB} * \alpha \cap \mathcal{KB} * \beta \subseteq \mathcal{KB} * (\alpha \vee \beta)$ (Disjunctive Overlap)
- P8** If $\neg\alpha \notin \mathcal{KB} * \beta$, then $Cn(\mathcal{KB} * \beta \cup \{\alpha\}) \subseteq \mathcal{KB} * (\alpha \wedge \beta)$ (Subexpansion)

Let $*$ be a standard AGM revision operator, which can be defined in terms of a total preorder (i.e., a transitive and complete relation) on $\mathcal{U}$ as follows. To a belief set $\mathcal{KB}$, we associate a total preorder $\preceq_{\mathcal{KB}}^{\text{tpo}}$ s.t. $\preceq_{\mathcal{KB}}^{\text{tpo}}$ is anchored on $\llbracket \mathcal{KB} \rrbracket$ (see paragraph following (1) above), and the models of the new belief set $\mathcal{KB} * \alpha$ are selected from $\llbracket \alpha \rrbracket$ by choosing the optimal ones w.r.t. to $\preceq_{\mathcal{KB}}^{\text{tpo}}$, i.e.,

$$\llbracket \mathcal{KB} * \alpha \rrbracket \stackrel{\text{def}}{=} \text{opt}(\llbracket \alpha \rrbracket, \preceq_{\mathcal{KB}}^{\text{tpo}}) \quad (2)$$

where the function $\text{opt}(\cdot)$ is defined as in (1) for $\preceq_{\mathcal{KB}}^{\text{tpo}}$.

The representation result below (Grove 1988; Katsuno and Mendelzon 1991) establishes the correspondence between the postulates above and the definition of $*$.

Theorem 1. *The revision operator $*$ satisfies Postulates P1-P8 iff for each $\mathcal{KB}$ there is a total preorder $\preceq_{\mathcal{KB}}^{\text{tpo}}$ anchored on $\llbracket \mathcal{KB} \rrbracket$ s.t., for each $\alpha \in \mathcal{L}$, $\mathcal{KB} * \alpha$ is determined from $\preceq_{\mathcal{KB}}^{\text{tpo}}$ as in (2).*

In the remainder of the paper, we explore some variations of the definition in (2), informed by different orderings in $\text{opt}(\cdot)$, and their respective impact on the behaviour of the corresponding belief revision operators. Regarding the notation for these, we adopt the following conventions: the operator $*$ is used in contexts in which it is assumed both

¹Note P0 actually follows from the Vacuity postulate P4, but is included for convenience, since Vacuity will not always hold for the operators explored later.

Success and Consistency hold, $\bullet$ in contexts in which Success is assumed but Consistency is not, e.g. in Section 4, and $\circ$ in contexts in which Consistency is assumed but not necessarily Success, e.g. in Sections 5 and 6.

3 Interval Order-based Revision

Interval orders, which, roughly, allow for the comparison of intervals, form a well-known subclass of ordering relations generalising total preorders (Fishburn 1985; Myers 1999). They have proven particularly fruitful in the recent literature on non-monotonic reasoning, with Rott's interval-based belief contraction (Rott 2014) and Booth and Varzinczak's weakening of rational closure for defeasible inference (Booth and Varzinczak 2021). In this section, we complete the picture initiated by Rott by exploring the role of interval orders in defining a belief *revision* operator.

3.1 Interval Orders

An *interval order* over $\mathcal{U}$ is a binary relation $\preceq^{\text{int}}$ satisfying reflexivity, i.e., $v \preceq^{\text{int}} v$ for all $v \in \mathcal{U}$, and the Ferrers condition, i.e., for all $v, u, v', u' \in \mathcal{U}$, we have:²

$$\text{If } v \preceq^{\text{int}} u \text{ and } v' \preceq^{\text{int}} u', \text{ then } v \preceq^{\text{int}} u' \text{ or } v' \preceq^{\text{int}} u \quad (3)$$

There are alternative ways to characterise interval orders. A particularly useful one is in terms of interval-based interpretations, as done by Booth and Varzinczak (2021) in a preferential semantics for defeasible conditionals.

Definition 1 (Interval-Based Interpretation). An *interval-based interpretation* is a pair $\mathcal{I} \stackrel{\text{def}}{=} \langle \mathcal{L}, \mathcal{U} \rangle$, where $\mathcal{L}$ and $\mathcal{U}$ are functions from $\mathcal{U}$ to $\mathbb{N}$ s.t. for all $v \in \mathcal{U}$, $\mathcal{L}(v) \leq \mathcal{U}(v)$. Given $\mathcal{I} = \langle \mathcal{L}, \mathcal{U} \rangle$ and $v \in \mathcal{U}$, $\mathcal{L}(v)$ is the **lower rank of v in $\mathcal{I}$** , and $\mathcal{U}(v)$ is the **upper rank of v in $\mathcal{I}$** . Hence, the pair $(\mathcal{L}(v), \mathcal{U}(v))$ is the **interval of v in $\mathcal{I}$** .

We point out that this definition differs slightly from the one by Booth and Varzinczak (2021) in that the co-domain of both $\mathcal{L}$ and $\mathcal{U}$ is taken to be $\mathbb{N}$ rather than $\mathbb{N} \cup \{\infty\}$. There, the infinite rank ∞ is interpreted as representing *impossibility*, but, as we shall soon see, our framework will be extendable to express different kinds of incoherence naturally, without the use of infinite ranks.

Example 1. Let $\mathcal{P} = \{p, q\}$, and let $\mathcal{L}$ and $\mathcal{U}$ be given by $\mathcal{L}(pq) = 0$, $\mathcal{L}(p\bar{q}) = 2$, $\mathcal{L}(\bar{p}q) = 1$, and $\mathcal{L}(\bar{p}\bar{q}) = 0$; $\mathcal{U}(pq) = 1$, $\mathcal{U}(p\bar{q}) = 3$, $\mathcal{U}(\bar{p}q) = 1$, and $\mathcal{U}(\bar{p}\bar{q}) = 3$. Then $\mathcal{I} = \langle \mathcal{L}, \mathcal{U} \rangle$ is an example of an interval-based interpretation for $\mathcal{P}$.

Figure 1 provides an easier to visualise representation of the interval-based interpretation in Example 1.

Given an interval-based interpretation $\mathcal{I} = \langle \mathcal{L}, \mathcal{U} \rangle$, the relation $\preceq^{\text{int}}_{\mathcal{I}}$ on $\mathcal{U}$, defined by $v \preceq^{\text{int}}_{\mathcal{I}} v'$ if and only if $\mathcal{L}(v) \leq \mathcal{U}(v')$, for every $v, v' \in \mathcal{U}$, is an interval order on $\mathcal{U}$. In terms of the intervals, that means we have $v \preceq^{\text{int}}_{\mathcal{I}} v'$ if and only if either v 's interval overlaps with v' 's one or is strictly 'below' it. For $\mathcal{I}$ as in Example 1, we have,

²Here we focus on the *weak* reading of the relation rather than the strict one, which was adopted by Rott (2014) and Booth and Varzinczak (2021).

	0	1	2	3
$\mathcal{L}$	$pq, \bar{p}\bar{q}$	$\bar{p}q$	$p\bar{q}$	
$\mathcal{U}$		$pq, \bar{p}q$		$p\bar{q}, \bar{p}\bar{q}$

Figure 1: An interval-based interpretation for $\mathcal{P} = \{p, q\}$.

for instance, $(pq, p\bar{q}), (\bar{p}\bar{q}, p\bar{q}) \in \preceq^{\text{int}}_{\mathcal{I}}$, but $(p\bar{q}, pq) \notin \preceq^{\text{int}}_{\mathcal{I}}$. Furthermore, it can be checked that for every interval order $\preceq^{\text{int}}$ over $\mathcal{U}$, there is an interval-based interpretation $\mathcal{I}$ such that $\preceq^{\text{int}} = \preceq^{\text{int}}_{\mathcal{I}}$.

3.2 IOB Revision

An interval order yields a (successful) belief revision operator $*$ in the same way as it is traditionally done for total preorders. To each belief set $\mathcal{KB}$, we can associate an interval order $\preceq^{\text{int}}_{\mathcal{KB}}$ such that the optimal models are the ones in $\llbracket \mathcal{KB} \rrbracket$, and then the models of the resulting new belief set $\mathcal{KB} * \alpha$ are selected from $\llbracket \alpha \rrbracket$ by picking the optimal ones according to $\preceq^{\text{int}}_{\mathcal{KB}}$:

$$\llbracket \mathcal{KB} * \alpha \rrbracket \stackrel{\text{def}}{=} \text{opt}(\llbracket \alpha \rrbracket, \preceq^{\text{int}}_{\mathcal{KB}}) \quad (4)$$

(We remind the reader that the function $\text{opt}(\cdot)$ is defined as in (1) above, with $\preceq^{\text{int}}_{\mathcal{KB}}$ as the informing order here.)

Definition 2 (Interval Order-Based Revision). $*$ is an *interval order-based (IOB) revision operator* if, to each $\mathcal{KB}$ we can associate an interval order $\preceq^{\text{int}}_{\mathcal{KB}}$ anchored on $\llbracket \mathcal{KB} \rrbracket$ s.t., for each α , $\mathcal{KB} * \alpha$ is determined from $\preceq^{\text{int}}_{\mathcal{KB}}$ as in (4) above.

Not surprisingly, the definition of revision in terms of interval orders comes with immediate consequences regarding the postulates the corresponding operator satisfies. Indeed, with the optimal models now informed by an underlying interval order, the revision operator from (4) above fails both Subexpansion (P8) and Vacuity (P4) in general. This is witnessed by the following example.

Example 2. Let $\mathcal{KB} = \text{Cn}(\{p \rightarrow q\})$ and let $\preceq^{\text{int}}_{\mathcal{KB}} = \preceq^{\text{int}}_{\mathcal{I}}$, with $\preceq^{\text{int}}_{\mathcal{I}}$ as in Example 1. It is easy to verify that $\preceq^{\text{int}}_{\mathcal{KB}}$ is anchored on $\llbracket \mathcal{KB} \rrbracket$, as $\llbracket \mathcal{KB} \rrbracket = \text{opt}(\llbracket \mathcal{KB} \rrbracket, \preceq^{\text{int}}_{\mathcal{KB}})$. For the failure of (P4), assume $\alpha = \neg q$. Then we have $q \notin \mathcal{KB}$ and $\llbracket \mathcal{KB} * \neg q \rrbracket = \text{opt}(\llbracket \neg q \rrbracket, \preceq^{\text{int}}_{\mathcal{KB}}) = \{p\bar{q}, \bar{p}\bar{q}\}$. But we have $\neg p \in \text{Cn}(\mathcal{KB} \cup \{\neg q\})$. For the failure of (P8), let $\alpha = \neg q$ and $\beta = \top$. We have $q \notin \mathcal{KB} * \top$ and $\neg p \in \text{Cn}(\mathcal{KB} * \top \cup \{\neg q\})$, but $\neg p \notin \mathcal{KB} * (\neg q \wedge \top)$.

We consider the following Disjunctive Rationality (DR) postulate, which is satisfied by the AGM revision operator based on total preorders $*$ above:

$$\text{P9 } \mathcal{KB} * (\alpha \vee \beta) \subseteq \mathcal{KB} * \alpha \cup \mathcal{KB} * \beta \quad (\text{DR})$$

This gives us a representation result for IOB revision operators (whose proof can actually be straightforwardly based on the proof characterising the more general BOB revision operators in the next section):

Theorem 2. An operator $*$ is an IOB revision operator iff it satisfies Tautology, Closure, Success, Inclusion, Extensionality, Consistency, Disjunctive Overlap, and Disjunctive Rationality.

Even though Vacuity does not hold in general, we note that, in the case one wants it to be enforced, it is sufficient to impose the following condition on $\preceq_{\mathcal{KB}}^{\text{int}}$ (where $\prec_{\mathcal{KB}}^{\text{int}}$ below is the *strict* version of $\preceq_{\mathcal{KB}}$):

$$\text{If } v \in \llbracket \mathcal{KB} \rrbracket \text{ and } v' \notin \llbracket \mathcal{KB} \rrbracket, \text{ then } v \prec_{\mathcal{KB}}^{\text{int}} v' \quad (5)$$

We end this section by noting that another, equivalent, way to define $\llbracket \mathcal{KB} * \alpha \rrbracket$ is directly in terms of an interval-based interpretation (see Definition 1). Let $\mathcal{I} = \langle \mathcal{L}, \mathcal{U} \rangle$, and, for every $\gamma \in \mathcal{L}$, let $\mathcal{L}(\gamma) \stackrel{\text{def}}{=} \min\{\mathcal{L}(v) \mid v \in \llbracket \gamma \rrbracket\}$ and $\mathcal{U}(\gamma) \stackrel{\text{def}}{=} \min\{\mathcal{U}(v) \mid v \in \llbracket \gamma \rrbracket\}$. Then:

$$\llbracket \mathcal{KB} * \alpha \rrbracket \stackrel{\text{def}}{=} \{v \in \llbracket \alpha \rrbracket \mid \mathcal{L}(v) \leq \mathcal{U}(\alpha)\} \quad (6)$$

4 Biorder-based Revision

We now switch to a class of orderings that have been studied within the theory of rational choice (Aleskerov, Bouyssou, and Monjardet 2007), namely biorders, and explore their fruitfulness in belief revision.

4.1 Biorders

A *biorder* over $\mathcal{U}$ is a binary relation $\preceq^{\text{bi}}$ that satisfies the Ferrers condition in (3) above, but that is not necessarily reflexive. Hence, they are a generalisation of the interval orders we have considered in the previous section, in which there are allowed to be valuations $z \in \mathcal{U}$ such that $z \not\preceq^{\text{bi}} z$.

Definition 3 (Dissonant valuation). *Given a biorder $\preceq^{\text{bi}}$ over $\mathcal{U}$ and $z \in \mathcal{U}$, we call z **dissonant** (w.r.t. $\preceq^{\text{bi}}$) if $z \not\preceq^{\text{bi}} z$.*

We will usually reserve z (possibly with subscripts) to denote dissonant valuations.

There are various ways to characterise biorders (Aleskerov, Bouyssou, and Monjardet 2007). A particularly useful, and more visual, one is in terms of two ranking functions, the lower and the upper ones, as in the interval-order case. The main difference is that here the intervals assigned to valuations can be of ‘negative’ length.

More formally, biorders can be characterised via a generalisation of interval-based interpretations:

Definition 4 (Biorder-Based Interpretation). *A **biorder-based interpretation** is a pair $\mathcal{B} \stackrel{\text{def}}{=} \langle \mathcal{L}, \mathcal{U} \rangle$, where $\mathcal{L}$ and $\mathcal{U}$ are functions from $\mathcal{U}$ to $\mathbb{N}$. Given $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ and $v \in \mathcal{U}$, $\mathcal{L}(v)$ is the **lower rank** of v in $\mathcal{B}$, and $\mathcal{U}(v)$ is the **upper rank** of v in $\mathcal{B}$. Hence, the pair $(\mathcal{L}(v), \mathcal{U}(v))$ is the **interval** of v in $\mathcal{B}$.*

Example 3. Let $\mathcal{P} = \{p, q\}$, and let $\mathcal{L}$ and $\mathcal{U}$ be such that $\mathcal{L}(pq) = 2$, $\mathcal{L}(p\bar{q}) = 3$, $\mathcal{L}(\bar{p}q) = 4$, and $\mathcal{L}(\bar{p}\bar{q}) = 0$; $\mathcal{U}(pq) = 3$, $\mathcal{U}(p\bar{q}) = 1$, $\mathcal{U}(\bar{p}q) = 0$, and $\mathcal{U}(\bar{p}\bar{q}) = 4$. Then $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ is an example of a biorder-based interpretation for $\mathcal{P}$.

Figure 2 provides an easier to visualise representation of the biorder-based interpretation in Example 3.

If by the *length* of v ’s interval in $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ we understand the result of $\mathcal{U}(v) - \mathcal{L}(v)$, then v having a negative interval corresponds to the case where $\mathcal{U}(v) < \mathcal{L}(v)$. In the biorder-based interpretation in Example 3, pq ’s interval has length 1, $p\bar{q}$ ’s interval has length -2 , $\bar{p}q$ ’s interval has length -4 , and $\bar{p}\bar{q}$ ’s interval has length 4.

	0	1	2	3	4
$\mathcal{L}$	$\bar{p}\bar{q}$		pq	$p\bar{q}$	$\bar{p}q$
$\mathcal{U}$	$\bar{p}q$	$\bar{p}\bar{q}$		pq	$\bar{p}\bar{q}$

Figure 2: A biorder-based interpretation for $\mathcal{P} = \{p, q\}$.

The intuition behind the assignment of a negative interval to a valuation is that it captures the fact the valuation in question bears some form of ‘impossibility’. Here, this stands as a simple alternative to the adoption of infinite ranks, as in the preferential approaches to defeasible reasoning.

As with interval-based interpretations, each biorder-based interpretation $\mathcal{B}$ yields a relation $\preceq_{\mathcal{B}}^{\text{bi}}$ over $\mathcal{U}$ given by $v \preceq_{\mathcal{B}}^{\text{bi}} v'$ if and only if $\mathcal{L}(v) \leq \mathcal{U}(v')$. Clearly, the dissonant valuations (see Definition 3) for $\preceq_{\mathcal{B}}^{\text{bi}}$ are precisely the valuations that get assigned negative intervals in $\mathcal{B}$. The following result is essentially shown by Aleskerov et al. (2007).

Proposition 1. *Let $\preceq^{\text{bi}} \subseteq \mathcal{U} \times \mathcal{U}$. Then $\preceq^{\text{bi}}$ is a biorder over $\mathcal{U}$ iff there is a biorder-based interpretation $\mathcal{B}$ s.t. $\preceq^{\text{bi}} = \preceq_{\mathcal{B}}^{\text{bi}}$.*

Given $\mathcal{B}$ from Example 3 above, we have $\preceq_{\mathcal{B}}^{\text{bi}} = \{(\bar{p}\bar{q}, \bar{p}\bar{q}), (\bar{p}\bar{q}, p\bar{q}), (\bar{p}\bar{q}, pq), (\bar{p}\bar{q}, p\bar{q}), (pq, pq), (pq, \bar{p}\bar{q}), (p\bar{q}, pq), (p\bar{q}, \bar{p}\bar{q}), (\bar{p}q, \bar{p}\bar{q})\}$.

We observe that the mapping $\mathcal{B} \mapsto \preceq_{\mathcal{B}}^{\text{bi}}$ is not necessarily one-to-one, as several distinct pairs of ranking functions may yield the same ordering. However, the proof of Proposition 1 shows that every biorder can be represented by a *compressed* biorder-based interpretation.

Definition 5 (Compressed Biorder-Based Interpretation). *A biorder-based interpretation $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ is **compressed** if, (a) for all $1 \leq i < j$, $\mathcal{L}^{-1}(j) \neq \emptyset$ implies $\mathcal{L}^{-1}(i) \neq \emptyset$, (b) for all $0 \leq i < j$, $\mathcal{U}^{-1}(j) \neq \emptyset$ implies $\mathcal{U}^{-1}(i) \neq \emptyset$, and (c) $0 \leq \max \mathcal{L} - \max \mathcal{U} \leq 1$, where $\max \mathcal{L} = \max\{i \mid \mathcal{L}^{-1}(i) \neq \emptyset\}$ (and similarly for $\max \mathcal{U}$).*

In words, $\mathcal{B}$ is compressed if there are no ‘gaps’ in either the lower or the upper rankings (except, possibly, $\mathcal{L}^{-1}(0)$, which will be empty exactly when $\preceq^{\text{bi}}$ has no optimal elements overall), and the largest lower rank is either equal to the largest upper rank or is bigger by one. For example Figure 3 shows a compressed biorder-based interpretation that gives the same biorder as the one in Figure 2.

	0	1	2
$\mathcal{L}$	$\bar{p}\bar{q}$	$pq, p\bar{q}$	$\bar{p}q$
$\mathcal{U}$	$\bar{p}q, p\bar{q}$	pq	$\bar{p}\bar{q}$

Figure 3: A compressed biorder-based interpretation.

For convenience, we assume that every biorder-based interpretation $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ is *normal* in that there is at least one valuation v such that $\mathcal{L}(v) = 0$. This ensures that the set of worlds that are optimal overall in $\mathcal{U}$ is non-empty.

4.2 BOB Revision

As we have done for interval orders, we can define a belief revision operator yielded by a biorder. To each (consistent)

belief set $\mathcal{KB}$, we therefore associate a biorder $\preceq_{\mathcal{KB}}^{\text{bi}}$ such that the optimal models overall are the ones in $\llbracket \mathcal{KB} \rrbracket$. The models of the resulting new belief set $\mathcal{KB} \bullet \alpha$ are selected from $\llbracket \alpha \rrbracket$ by choosing the optimal ones according to $\preceq_{\mathcal{KB}}^{\text{bi}}$:

$$\llbracket \mathcal{KB} \bullet \alpha \rrbracket \stackrel{\text{def}}{=} \text{opt}(\llbracket \alpha \rrbracket, \preceq_{\mathcal{KB}}^{\text{bi}}) \quad (7)$$

Definition 6 (Biorder-Based Revision). *$\bullet$ is a biorder-based (BOB) revision operator if, to each $\mathcal{KB}$ we can associate a biorder $\preceq_{\mathcal{KB}}^{\text{bi}}$ anchored on $\llbracket \mathcal{KB} \rrbracket$ s.t., for each α , $\mathcal{KB} \bullet \alpha$ is determined from $\preceq_{\mathcal{KB}}^{\text{bi}}$ as in (7) above.*

Just as for IOB revision (cf. end of Section 3.2), we can define $\llbracket \mathcal{KB} \bullet \alpha \rrbracket$ directly in terms of $\mathcal{L}$ and $\mathcal{U}$ as follows:

$$\llbracket \mathcal{KB} \bullet \alpha \rrbracket \stackrel{\text{def}}{=} \{v \in \llbracket \alpha \rrbracket \mid \mathcal{L}(v) \leq \mathcal{U}(\alpha)\} \quad (8)$$

where $\mathcal{L}(\gamma) = \min\{\mathcal{L}(v) \mid v \in \llbracket \gamma \rrbracket\}$ and $\mathcal{U}(\gamma) = \min\{\mathcal{U}(v) \mid v \in \llbracket \gamma \rrbracket\}$. Thus, we can see that $\bullet$ will not satisfy Consistency in general. Indeed, if $\mathcal{U}(\alpha) < \mathcal{L}(\alpha)$, then $\llbracket \mathcal{KB} \bullet \alpha \rrbracket$ will be empty, even if $\llbracket \alpha \rrbracket$ is non-empty. For an example of such a situation, let $\mathcal{KB} = \text{Cn}(\{-p \wedge \neg q\})$, let $\alpha = p$, and let $\preceq_{\mathcal{KB}}^{\text{bi}} = \preceq_{\mathcal{B}}^{\text{bi}}$, where $\mathcal{B}$ is as in Example 3.

Definition 7 (Destabilising Sentences). *Given a BOB revision operator $\bullet$ and a belief set $\mathcal{KB}$, we say that α is **destabilising** for $\bullet$ and $\mathcal{KB}$ if $\mathcal{KB} \bullet \alpha$ is inconsistent.*

A rational agent may have sentences that are destabilising for it. Intuitively, a destabilising sentence is an incoming piece of information that would make an agent's beliefs 'collapse' if accepted. It turns out that one can have α as destabilising, while $\alpha \wedge \beta$, for some β , is not. For instance, for $\mathcal{KB}$ as above, we have that p is destabilising while $p \wedge q$ is not. This corresponds to situations in which some extra information, or context, as called by Schwind (2025) (see also Section 7), makes the acceptance of the sentence compatible with the agent's current beliefs. We shall give a special name to those destabilising sentences that cannot be rendered stable even with the addition of more information.

Definition 8 (Irreconcilable Sentences). *Given a BOB revision operator $\bullet$ and a belief set $\mathcal{KB}$, we say that α is **irreconcilable** for $\bullet$ and $\mathcal{KB}$ if $\mathcal{KB} \bullet \beta$ is inconsistent for all $\beta \in \mathcal{L}$ such that $\beta \models \alpha$.*

For an example of an irreconcilable sentence, let $\mathcal{KB} = \text{Cn}(\{-p \wedge \neg q\})$ and let $\preceq_{\mathcal{KB}}^{\text{bi}} = \preceq_{\mathcal{B}}^{\text{bi}}$, where $\mathcal{B}$ is as in Example 3. Then $\neg p \wedge q$ is irreconcilable.

We have the following representation result for BOB revision operators, which shows that BOB revision is characterised by the same list of postulates as IOB revision, just without Consistency:

Theorem 3. *An operator $\bullet$ is a BOB revision operator iff it satisfies Tautology, Closure, Success, Inclusion, Extensionality, Disjunctive Overlap, and Disjunctive Rationality.*

For the completeness part of the proof of this result, we construct, from a given operator $\bullet$ satisfying the stated postulates, and for each $\mathcal{KB}$, a biorder $\preceq_{\mathcal{KB}}^{\text{bi}}$ via a compressed biorder-based interpretation $\mathcal{B}_{\mathcal{KB}}^{\bullet} = \langle \mathcal{L}_{\mathcal{KB}}^{\bullet}, \mathcal{U}_{\mathcal{KB}}^{\bullet} \rangle$ as follows. We inductively build two sequences of subsets of $\mathcal{U}$, an increasing and a decreasing one:

$$U_0 \subseteq U_1 \subseteq \dots \subseteq U_n \quad \text{and} \quad V_0 \supseteq V_1 \supseteq \dots \supseteq V_n$$

Initially $U_{-1} = \emptyset$ and $V_{-1} = \mathcal{U}$. And then, for each $i \geq -1$:³

$$U_{i+1} \stackrel{\text{def}}{=} U_i \cup \llbracket \mathcal{KB} \bullet \bigvee V_i \rrbracket$$

$$V_{i+1} \stackrel{\text{def}}{=} \bigcup \{Y \subseteq \mathcal{U} \mid \llbracket \mathcal{KB} \bullet \bigvee Y \rrbracket \not\subseteq U_{i+1}\}$$

Let n be minimal s.t. $n \geq 1$ and $U_{n+1} = U_n$. (Note n is guaranteed to exist since $\mathcal{U}$ is finite.) If, after the procedure terminates, we have $U_n \neq \mathcal{U}$, then we also add $\mathcal{U}$ to the end of the list to make sure every member of $\mathcal{U}$ appears in at least one U_i and thus that $\mathcal{L}_{\mathcal{KB}}^{\bullet}$ defined below assigns a value to each valuation. Finally, from these two sequences we read off $\mathcal{L}_{\mathcal{KB}}^{\bullet}$ and $\mathcal{U}_{\mathcal{KB}}^{\bullet}$:

$$\mathcal{L}_{\mathcal{KB}}^{\bullet}(u) \stackrel{\text{def}}{=} \min\{i \mid u \in U_i\}, \quad \mathcal{U}_{\mathcal{KB}}^{\bullet}(u) \stackrel{\text{def}}{=} \min\{i \mid u \notin V_i\}$$

It can then be shown that $\mathcal{U}_{\mathcal{KB}}^{\bullet}$ also assigns a value to each valuation, that $\preceq_{\mathcal{KB}}^{\text{bi}, \bullet}$ is anchored on $\llbracket \mathcal{KB} \rrbracket$ and that $\text{opt}(\llbracket \alpha \rrbracket, \preceq_{\mathcal{KB}}^{\text{bi}, \bullet}) = \llbracket \mathcal{KB} \bullet \alpha \rrbracket$ for each α .

Definition 9 (Canonical Biorder-based Interpretation). *Given a BOB revision operator $\bullet$ and belief set $\mathcal{KB}$, the biorder-based interpretation $\mathcal{B}_{\mathcal{KB}}^{\bullet}$ defined above is the **canonical biorder-based interpretation associated to $\bullet$ and $\mathcal{KB}$** .*

We observe that two different biorders may yield the same revision behaviour. In particular, for every dissonant valuation z , the value of $\mathcal{L}(z)$ is irrelevant, i.e., we can increase $\mathcal{L}(z)$ without changing the resulting revision behaviour. This has essentially been pointed out by Aleskerov et al. (2007) as well. To put it another way, if, for a given biorder-based interpretation $\mathcal{B}$, we have $z \not\preceq_{\mathcal{B}}^{\text{bi}} z$, then, for all $v \in \mathcal{U}$, the relation given by $\preceq_{\mathcal{B}}^{\text{bi}} \setminus \{(z, v)\}$, which is a biorder, will yield the same revision behaviour. As a result of this, every BOB revision operator $\bullet$ can be generated by a function that assigns to each $\mathcal{KB}$ a biorder satisfying the following additional property (which does not play any role in the soundness proof of Theorem 3 above):

$$\text{If } v \preceq^{\text{bi}} v', \text{ then } v \preceq^{\text{bi}} v \quad (9)$$

That is, no dissonant valuation is as good as any other. This fact will be useful in proving some results in later sections. (We note that $\preceq_{\mathcal{KB}}^{\text{bi}}$ constructed in the completeness proof for Theorem 3 satisfies Property (9), as the elements in U_n are precisely those that are not dissonant in $\preceq_{\mathcal{KB}}^{\text{bi}}$.)

Even though we do not have Consistency in general, given the assumption that $\mathcal{KB}$ is consistent, we do get the following restricted version:

P6' If $\alpha \in \mathcal{KB}$, then $\mathcal{KB} \bullet \alpha \not\models \perp$ (Relative Consistency)

We will return to the consistency issue later. We also have:

If $\alpha \in \mathcal{KB}$, then $\mathcal{KB} \bullet \alpha \subseteq \mathcal{KB}$ (Endogenous Inclusion)

Nevertheless, BOB revision does not satisfy the following property, in general:

If $\alpha \in \mathcal{KB}$, then $\mathcal{KB} \subseteq \mathcal{KB} \bullet \alpha$ (Positive Vacuity)

³Here whenever a valuation appears within the scope of a propositional connective, we are taking it to stand for any sentence which has that valuation as its only model.

That is, it is possible that, upon being presented with a sentence that is already a consequence of the agent's beliefs, the agent may in fact remove some beliefs. This can be seen in Example 3, where we have $\mathcal{KB} = \text{Cn}(\{\neg p \wedge \neg q\})$ and hence $p \leftrightarrow q \in \mathcal{KB}$, but $\mathcal{KB} \bullet (p \leftrightarrow q) = \text{Cn}(\{p \leftrightarrow q\})$, i.e., we lose belief in $\neg p \wedge \neg q$.

It is possible to capture Positive Vacuity by enforcing a simple condition on biorder-based interpretations, namely $\mathcal{L}^{-1}(0) \cap \mathcal{U}^{-1}(0) \neq \emptyset$. However, doing that would still not necessarily give us the following generalisation of it, which corresponds to the rule of Cautious Monotony from conditional reasoning:

If $\beta \in \mathcal{KB} \bullet \alpha$, then $\mathcal{KB} \bullet \alpha \subseteq \mathcal{KB} \bullet (\alpha \wedge \beta)$

That is, acceptance of α may still leave us open to some kind of instability. We have the following counterexample.

Example 4. Suppose $\mathcal{KB} = \text{Cn}(\{p \leftrightarrow q\})$ and $\bullet$ associates the biorder $\preceq_{\mathcal{B}}^{\text{bi}}$ to $\mathcal{KB}$, where $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ is a compressed interpretation specified by: $\mathcal{L}(pq) = \mathcal{L}(\bar{p}\bar{q}) = 0$, $\mathcal{L}(p\bar{q}) = 1$ and $\mathcal{L}(\bar{p}q) = 2$; $\mathcal{U}(\bar{p}\bar{q}) = \mathcal{U}(\bar{p}q) = 0$, $\mathcal{U}(p\bar{q}) = 1$, and $\mathcal{U}(pq) = 2$. Then $p \in \mathcal{KB} \bullet (p \vee q)$ but $q \in \mathcal{KB} \bullet (p \vee q) \setminus \mathcal{KB} \bullet ((p \vee q) \wedge p)$.

In the above example, p is an example of what we call a *precarious* input:

Definition 10 (Precarious Sentences). Given a BOB revision operator $\bullet$ and a belief set $\mathcal{KB}$, we say that α is *precarious* for $\bullet$ and $\mathcal{KB}$ if there is $\beta \in \mathcal{L}$ s.t. $\beta \in \mathcal{KB} \bullet \alpha$ and $\mathcal{KB} \bullet \alpha \not\subseteq \mathcal{KB} \bullet (\alpha \wedge \beta)$

We have the following easy result concerning the relationship between destabilising, irreconcilable and precarious sentences:

Proposition 2. Let $\alpha \in \mathcal{L}$. If α is destabilising for $\bullet$ and $\mathcal{KB}$, then it is either irreconcilable or precarious for $\bullet$ and $\mathcal{KB}$.

4.3 ZTBOB Revision

We now turn our attention to a special subclass of biorders, namely those also satisfying some forms of transitivity.

We start by noting that the converse to Proposition 2 does not hold, in general, as can easily be seen in Example 4, where p is precarious but not destabilising for $\bullet$ and $\mathcal{KB}$. Since clearly irreconcilability implies destabilising, this boils down to requiring that precarious implies destabilising, which can be expressed as the following Consistent Cautious Monotony (CCM) postulate on the $\bullet$ operator:

P10 If $\beta \in \mathcal{KB} \bullet \alpha$ and $\mathcal{KB} \bullet \alpha \not\models \perp$, then $\mathcal{KB} \bullet \alpha \subseteq \mathcal{KB} \bullet (\alpha \wedge \beta)$ (Consistent Cautious Monotony)

We can satisfy this postulate by forcing a further condition on $\mathcal{B}$ and $\preceq_{\mathcal{B}}^{\text{bi}}$ (remember z denotes a dissonant valuation):

BZT If $\mathcal{L}(u) \leq \mathcal{U}(z) < \mathcal{U}(u)$, then $\mathcal{L}(z) \leq \mathcal{U}(z)$

In words, there is no dissonant valuation z such that $\mathcal{L}(u) \leq \mathcal{U}(z) < \mathcal{U}(u)$. This property fails for the $\mathcal{B}$ in Example 4, as can be seen by taking $u = pq$ and $z = \bar{p}q$. If $\mathcal{B}$ satisfies BZT, then $\preceq_{\mathcal{B}}^{\text{bi}}$ also satisfies the following weakening of transitivity, which we call *z-transitivity*:

$$\text{If } u \preceq_{\mathcal{B}}^{\text{bi}} v, v \preceq_{\mathcal{B}}^{\text{bi}} z \text{ and } z \not\models^{\text{bi}} z, \text{ then } u \preceq_{\mathcal{B}}^{\text{bi}} z \quad (10)$$

We call any biorder $\preceq^{\text{bi}}$ satisfying Property (10) a *z-transitive* biorder.

Proposition 3. (i) Let $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ be a biorder-based interpretation satisfying BZT. Then $\preceq_{\mathcal{B}}^{\text{bi}}$ is *z-transitive*. (ii) Let $\preceq^{\text{bi}}$ be a *z-transitive* biorder. Then every **compressed** biorder-based interpretation $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ s.t. $\preceq_{\mathcal{B}}^{\text{bi}} = \preceq^{\text{bi}}$ satisfies BZT.

The ordering $\preceq_{\mathcal{B}}^{\text{bi}}$ in Example 4 is not *z-transitive*, as witnessed by the fact that $p\bar{q} \preceq_{\mathcal{B}}^{\text{bi}} pq$, $pq \preceq_{\mathcal{B}}^{\text{bi}} \bar{p}q$, $\bar{p}q \not\models_{\mathcal{B}}^{\text{bi}} \bar{p}q$ but $p\bar{q} \not\models_{\mathcal{B}}^{\text{bi}} \bar{p}q$.

We thus arrive at our first subclass of BOB revision, and its associated axiomatic characterisation.

Definition 11 (ZTBOB Revision). $\bullet$ is a *z-transitive* biorder-based (ZTBOB) revision operator if, to each $\mathcal{KB}$ we can associate a *z-transitive* biorder $\preceq_{\mathcal{KB}}^{\text{bi}}$ anchored on $\llbracket \mathcal{KB} \rrbracket$ s.t., for each α , $\mathcal{KB} \bullet \alpha$ is determined from $\preceq_{\mathcal{KB}}^{\text{bi}}$ as in (7).

Theorem 4. An operator $\bullet$ is a ZTBOB revision operator iff it satisfies all the postulates from Theorem 3, i.e., Tautology, Closure, Success, Inclusion, Extensionality, Disjunctive Overlap and Disjunctive Rationality, **plus** CCM.

4.4 TBOB Revision

As discussed in the previous section, *z-transitivity* is a weakened form of transitivity. A natural question to ask is “what if we enforce full transitivity?” In terms of compressed biorder-based interpretations (Definition 5), this corresponds to a particularly simple property, namely:

$$\text{BT } \mathcal{U}(v) \leq \mathcal{L}(v)$$

Notice that we still allow for valuations to be assigned negative intervals here, but we disallow intervals of non-zero positive length. As a result, transitive biorders are still not necessarily reflexive.

Proposition 4. (i) For every biorder-based interpretation $\mathcal{B}$ satisfying BT, $\preceq_{\mathcal{B}}^{\text{bi}}$ is transitive. (ii) Let $\preceq^{\text{bi}}$ be a transitive biorder; then every **compressed** biorder-based interpretation $\mathcal{B} = \langle \mathcal{L}, \mathcal{U} \rangle$ s.t. $\preceq_{\mathcal{B}}^{\text{bi}} = \preceq^{\text{bi}}$ satisfies BT.

One can see that in the biorder-based interpretation from Example 3, the condition BT does not hold as $\mathcal{U}(\bar{p}\bar{q}) > \mathcal{L}(\bar{p}\bar{q})$. Furthermore, the associated biorder $\preceq_{\mathcal{B}}^{\text{bi}}$ is not transitive as $\{(p\bar{q}, \bar{p}\bar{q}), (\bar{p}\bar{q}, p\bar{q})\} \subseteq \preceq_{\mathcal{B}}^{\text{bi}}$, but $(p\bar{q}, p\bar{q}) \notin \preceq_{\mathcal{B}}^{\text{bi}}$.

Definition 12 (TBOB Revision). $\bullet$ is a transitive biorder-based (TBOB) revision operator if, to each $\mathcal{KB}$ we can associate a transitive biorder $\preceq_{\mathcal{KB}}^{\text{bi}}$ anchored on $\llbracket \mathcal{KB} \rrbracket$ s.t., for each α , $\mathcal{KB} \bullet \alpha$ is determined from $\preceq_{\mathcal{KB}}^{\text{bi}}$ as in (7) above.

As it turns out, TBOB revision brings us a step closer to AGM revision, in that this family of operators captures the second AGM supplementary postulate, Subexpansion. We have the following representation result:

Theorem 5. An operator $\bullet$ is a TBOB revision operator iff it satisfies all postulates from Theorem 3, i.e., Tautology, Closure, Success, Inclusion, Extensionality, Disjunctive Overlap and Disjunctive Rationality, **plus** Subexpansion.




Figure 4: System-of-spheres style representation of a TBOB revision operator. Grey rings indicate “impossible” ranks.

In other words, transitive biorders characterise the revision operators that are characterised exactly by the AGM postulates *minus* Consistency *plus* Disjunctive Rationality. (Note that Vacuity is recaptured since it follows from Subexpansion, Tautology and Extensionality). In the literature on rational choice theory, the family of choice functions corresponding to a restricted version of transitive biorders has been characterised by Gerasimou (2018).

We note that Subexpansion is equivalent to the following Disjunctive Monotony (DM) postulate, given the other postulates for TBOB revision:

$$\text{If } \neg\beta \notin \mathcal{KB} \bullet \alpha, \text{ then } \mathcal{KB} \bullet (\alpha \vee \beta) \subseteq \mathcal{KB} \bullet \beta \quad (\text{DM})$$

This postulate is essentially the belief revision version of the (γ^+) postulate from rational choice of Salant and Rubinstein (2008), which has been considered by Gerasimou (2018), and also by Grossi et al. (2022), from where we have taken the name Disjunctive Monotony. We could substitute it for Subexpansion in Theorem 5 above, but would still need Disjunctive Rationality as it does not follow from the other postulates mentioned in Theorem 5. As we will see, Disjunctive Monotony comes in handy for the representation results we state in the next section.

The property $\mathcal{U}(u) \leq \mathcal{L}(u)$ for all u , combined with the property that $\mathcal{L}(u)$ is irrelevant for all those u such that $\mathcal{U}(u) < \mathcal{L}(u)$ roughly means that we only need to keep $\mathcal{U}$. Furthermore any dissonant valuations z s.t. $\mathcal{U}(z) = i$ can be safely placed in a newly created upper rank in between i and $i + 1$ without changing the revision behaviour. All this means that we can specify an interpretation $\langle \mathcal{L}, \mathcal{U} \rangle$ using a single ranking $\mathcal{R}$ **plus** a specification of which, if any, ranks are “impossible”. See Figure 4 above for an illustration.

$$\llbracket \mathcal{KB} \bullet \alpha \rrbracket = \begin{cases} \{u \in \llbracket \alpha \rrbracket \mid \mathcal{R}(u) \leq \mathcal{R}(\alpha)\}, & \text{if } \mathcal{R}(\alpha) \text{ is not impossible} \\ \emptyset, & \text{otherwise} \end{cases}$$

To sum it up, while both BOB and TBOB revision operators satisfy many important AGM postulates, in particular the Success one, their significance for standard belief revision is, at least at first sight, diminished by the fact that they do not always yield consistent outputs. Indeed, the Consistency postulate is non-negotiable in ‘orthodox’ belief revision. We could, in principle, circumvent this by enforcing

consistency. Note, however, that this would lead us back to interval orders or even total preorders, which is essentially where we started from. In the next section, we step back and look at the possibly incoherent revision inputs as still modelling something useful about an agent’s belief state.

5 Non-prioritised Revision

We shall now look at our previously defined operators from a different angle, namely that of non-prioritised belief revision. This entitles us to modify their definition in order to discard inputs leading to inconsistency as ‘incredible’. This captures the intuition that whenever an agent receives as input a sentence α that is destabilising, they will just ignore it, sacrificing Success for Consistency, which is precisely the non-prioritised revision approach.

5.1 BOB NPR Operators

We start by noting that the last three operators we have defined, namely general biorder-based, z -transitive biorder-based, and transitive biorder-based revisions, give rise to corresponding classes of non-prioritised revision operators, with associated axiomatics.

Definition 13 (BOB Non-Prioritised Revision). *Let $\bullet$ be a BOB revision operator. The BOB non-prioritised revision operator (BOB NPR, for short) $\circ$ generated by $\bullet$ is the operator defined by setting, for each belief set $\mathcal{KB}$ and $\alpha \in \mathcal{L}$,*

$$\mathcal{KB} \circ \alpha \stackrel{\text{def}}{=} \begin{cases} \mathcal{KB} \bullet \alpha, & \text{if this is consistent} \\ \mathcal{KB}, & \text{otherwise} \end{cases}$$

Equivalently, $\circ$ is a BOB NPR operator if to each $\mathcal{KB}$ we can associate a biorder $\preceq_{\mathcal{KB}}^{\text{bi}}$ anchored on $\llbracket \mathcal{KB} \rrbracket$ s.t. $\llbracket \mathcal{KB} \circ \alpha \rrbracket$ is $\text{opt}(\llbracket \alpha \rrbracket, \preceq_{\mathcal{KB}}^{\text{bi}})$ if this is non-empty, and is $\llbracket \mathcal{KB} \rrbracket$ otherwise. We similarly obtain ZTBOB NPR operators and TBOB NPR operators, when $\bullet$ is taken to be a ZTBOB, respectively TBOB, revision operator.

Definition 14 (Credible Sentences). *We say α is credible w.r.t. $\mathcal{KB}$ and $\circ$ (denoted $\alpha \in \mathcal{C}_{\mathcal{KB}}^\circ$) if $\alpha \in \mathcal{KB} \circ \alpha$, and incredible otherwise.*

It is easy to see that if $\circ$ is defined from $\bullet$ as above, then α is credible w.r.t. $\mathcal{KB}$ and $\circ$ if and only if α is not destabilising w.r.t. $\mathcal{KB}$ and $\bullet$ (cf. Definition 7).

Definition 13 leads us to the following representation result:

Theorem 6. *An operator $\circ$ is a BOB NPR operator iff it satisfies Tautology, Closure, Inclusion, Extensionality, Consistency, Disjunctive Overlap and the list of postulates below.*

P2' $\alpha \in \mathcal{KB} \circ \alpha$ or $\mathcal{KB} \circ \alpha = \mathcal{KB}$ (Relative Success)

P11 If $\alpha \vee \beta \in \mathcal{KB} \circ (\alpha \vee \beta)$, then $\alpha \in \mathcal{KB} \circ \alpha$ or $\beta \in \mathcal{KB} \circ \beta$

P12 If $\alpha \vee \beta \in \mathcal{KB} \circ (\alpha \vee \beta)$ and $\beta \notin \mathcal{KB} \circ \beta$, then $\mathcal{KB} \circ \alpha \subseteq \mathcal{KB} \circ (\alpha \vee \beta)$

P13 If $\alpha \in \mathcal{KB} \circ \alpha$ and $\beta \in \mathcal{KB} \circ \beta$, then $\mathcal{KB} \circ (\alpha \vee \beta) \subseteq \mathcal{KB} \circ \alpha \cup \mathcal{KB} \circ \beta$

Some of these postulates are familiar from the literature on credibility-limited revision (Hansson et al. 2001). P2' is a hallmark property of such revision operators. It says that, at least as far as their effects on the belief set are concerned,

incredible input sentences are plain ignored. P11 says a disjunction is credible only if at least one of its disjuncts is. P12 says that if one of the disjuncts β of a credible sentence $\alpha \vee \beta$ is incredible, then $\mathcal{KB} \circ (\alpha \vee \beta)$ should contain everything that would be believed from revising with α alone. P13 is a restricted form of Disjunctive Rationality (DR), saying that it should hold whenever α and β are both credible.

Proposition 5. *Disjunctive Overlap implies the following postulate in the presence of Closure.*

P14 *If $\alpha \in \mathcal{KB} \circ \alpha$ and $\beta \in \mathcal{KB} \circ \beta$, then $\alpha \vee \beta \in \mathcal{KB} \circ (\alpha \vee \beta)$*

The above property says the disjunction of two credible sentences is itself credible.

5.2 ZTBOB and TBOB NPR Operators

One of the characteristic properties for the general biorders-based NPR $\circ$ is P12. However, a case can be made for a version of P12 with equality in the consequent, i.e., if a credible disjunction is received, then any incredible disjunct is essentially ignored.

P12+ *If $\alpha \vee \beta \in \mathcal{KB} \circ (\alpha \vee \beta)$ and $\beta \notin \mathcal{KB} \circ \beta$, then $\mathcal{KB} \circ \alpha = \mathcal{KB} \circ (\alpha \vee \beta)$*

P12+ does not hold, in general, for BOB NPR operators:

Example 5. *Let $\mathcal{KB} = \text{Cn}(\{p \leftrightarrow q\})$ and take the biorder $\preceq_{\mathcal{B}}^{\text{bi}}$ anchored on $\mathcal{KB}$ given by $\mathcal{B}$ from Example 4. We have $p \vee (\neg p \wedge q) \in \mathcal{KB} \circ (p \vee (\neg p \wedge q)) = \text{Cn}(\{p \wedge q\})$ and $\neg p \wedge q \notin \mathcal{KB} \circ (\neg p \wedge q)$. But $\mathcal{KB} \circ p = \text{Cn}(\{p\})$ and so $\mathcal{KB} \circ (p \vee (\neg p \wedge q)) \not\subseteq \mathcal{KB} \circ p$.*

It may be noted in the previous example that BZT fails, i.e., the ordering is not z -transitive. It turns out that the BOB NPR operators satisfying P12+ are *precisely* the ZTBOB NPR operators.

Theorem 7. *An operator $\circ$ is a ZTBOB NPR operator iff it satisfies all the postulates from Theorem 6, i.e., Tautology, Closure, Inclusion, Extensionality, Consistency, Disjunctive Overlap, Relative Success, P11, P13, but with P12 replaced by P12+.*

Notice that P12+ causes some redundancy in the list: Besides replacing P12, it causes P13 to reduce to Disjunctive Rationality. So the official characteristic list of properties for ZTBOB NPR operators becomes: Tautology, Closure, Relative Success, Inclusion, Extensionality, Consistency, Disjunctive Overlap, Disjunctive Rationality, P11 and P12+.

When a TBOB operator $\bullet$ is assumed in the definition of $\circ$, the following representation result is warranted:

Theorem 8. *An operator $\circ$ is a TBOB NPR operator iff it satisfies all the postulates satisfied by ZTBOB NPR operators, i.e., Tautology, Closure, Relative Success, Inclusion, Extensionality, Consistency, Disjunctive Overlap, Disjunctive Rationality, P11 and P12+, **plus** Disjunctive Monotony.*

Note that, although Disjunctive Monotony is sound for TBOB NPR operators, Subexpansion (which is sound for BOB revision operators) is not.

6 BOB NPR as Credibility-limited Revision

In this section, we show how the two subclasses of BOB non-prioritised revision from the previous section can be formulated as credibility-limited revision operators.

Credibility-limited revision has been put forward by Hansson et al. (2001) and its general schema is as follows. We assume a (successful and consistent) revision operator $*$ and a function $\mathcal{C}$ that, to each belief set $\mathcal{KB}$, assigns a *credible set* $\mathcal{C}_{\mathcal{KB}} \subseteq \mathcal{L}$, which, intuitively, contains sentences deemed acceptable. Then we have the following operator:

$$\mathcal{KB} \circ_{*, \mathcal{C}} \alpha \stackrel{\text{def}}{=} \begin{cases} \mathcal{KB} * \alpha, & \text{if } \alpha \in \mathcal{C}_{\mathcal{KB}} \\ \mathcal{KB}, & \text{otherwise} \end{cases}$$

The operator $*$ is to be interpreted as the revision operator the agent would use if they *had* to accept the new sentence (even though at first glance they would choose to reject it). The sentences not in $\mathcal{C}_{\mathcal{KB}}$ represent inputs that are in some sense too remote in the agent's belief state to be entertained.

By assuming different properties of $*$ and $\mathcal{C}$, and possibly some properties governing the interplay between these two, we obtain different properties of the operator $\circ_{*, \mathcal{C}}$. Regarding the underlying revision operator $*$, Hansson et al. consider the case where it is a basic or a full AGM revision operator. Regarding $\mathcal{C}$, they consider several properties, including the following:

C1 $\mathcal{KB} \subseteq \mathcal{C}_{\mathcal{KB}}$

C2 If $\alpha \in \mathcal{C}_{\mathcal{KB}}$ and $\alpha \equiv \beta$, then $\beta \in \mathcal{C}_{\mathcal{KB}}$

C3 If $\alpha \vee \beta \in \mathcal{C}_{\mathcal{KB}}$, then $\alpha \in \mathcal{C}_{\mathcal{KB}}$ or $\beta \in \mathcal{C}_{\mathcal{KB}}$

C4 If $\alpha \in \mathcal{C}_{\mathcal{KB}}$ and $\alpha \models \beta$, then $\beta \in \mathcal{C}_{\mathcal{KB}}$

C1 says if a sentence is already supported by our belief set, then we have no reason to reject it, C2 is a syntax irrelevance property on credible sentences, while C3 says a disjunction is credible only if at least one of its disjuncts is (cf. Postulate P14). We will adopt C1-C3 here. C4 is the Single Sentence Closure property (Hansson et al. 2001), which we consider too strong. The reason is that it could be rational for an agent to reject α at first glance (i.e., $\alpha \notin \mathcal{C}_{\mathcal{KB}}$), but to accept it when strengthened with additional ‘explanation’ $\alpha \wedge \beta$ (i.e., $\alpha \wedge \beta \in \mathcal{C}$). We have already encountered a concrete example of this situation in the Introduction, where α is “my five-year-old nephew had lunch with King Charles today” and β is “King Charles visited a local school today.” A further discussion on the reasons to reject C4 is provided by Schwind (2025). We are interested in weakening Single Sentence Closure to the following:

C4' If $\alpha, \beta \in \mathcal{C}_{\mathcal{KB}}$, then $\alpha \vee \beta \in \mathcal{C}_{\mathcal{KB}}$

In the presence of C2 the above property implies the following:

If $\beta \wedge \alpha, \neg \beta \wedge \alpha \in \mathcal{C}_{\mathcal{KB}}$ then $\alpha \in \mathcal{C}_{\mathcal{KB}}$.

This seems to be plausible, after all, if we do not hold “My five year old nephew had lunch with King Charles today” to be credible, but we do hold “King Charles visited a local school today and my five year old nephew had lunch with King Charles today” to be credible, then presumably we would hold “King Charles did not visit a local school today

and my five year old nephew had lunch with King Charles today” to be incredible.

We also require some minimal property of coherence between $*$ and $\mathcal{C}$. Intuitively, if a sentence α is acceptable but the addition of β renders it unacceptable, then β is exceptional within α . This is formalised by the following:

$$\text{If } \alpha \in \mathcal{KB} \text{ and } \alpha \wedge \beta \notin \mathcal{KB} \text{ then } \neg\beta \in \mathcal{KB} * \alpha \quad (11)$$

Definition 15 (CL Structure). We say the pair $\langle *, \mathcal{C} \rangle$ is a **CL structure** if (i) $*$ is a revision operator satisfying Success and Consistency, (ii) $\mathcal{C}$ satisfies C1-C3 and C4', and (iii) $*$ and $\mathcal{C}$ jointly satisfy (11) above.

First we turn to ZTBOB non-prioritised revision.

Theorem 9. An operator $\circ$ is a ZTBOB NPR operator iff $\circ = \circ_{*, \mathcal{C}}$ for some CL structure $\langle *, \mathcal{C} \rangle$ s.t. $*$ is an IOB revision operator.

The ‘only if’ part of the proof of this result requires us to construct a CL structure $\langle *, \mathcal{C}^\circ \rangle$ from a given BOB NPR operator $\circ$. The definition of $\mathcal{C}^\circ$ is the expected one, viz.,

$$\alpha \in \mathcal{C}_{\mathcal{KB}}^\circ \text{ iff } \alpha \in \mathcal{KB} \circ \alpha$$

For $*$ it essentially boils down to finding the appropriate way to transform a biorder into an interval order. For suppose $\circ$ associates $\preceq_{\mathcal{KB}}^{\text{bi}}$ to $\mathcal{KB}$. We define an ordering $\preceq_{\mathcal{KB}}^{\text{int}}$ as follows, for any $u, v \in \mathcal{U}$:

$$u \preceq_{\mathcal{KB}}^{\text{int}} v \quad \text{iff} \quad [u, v \notin \text{Diss}(\preceq_{\mathcal{KB}}^{\text{bi}}) \text{ and } u \preceq_{\mathcal{KB}}^{\text{bi}} v] \\ \text{or } v \in \text{Diss}(\preceq_{\mathcal{KB}}^{\text{bi}})$$

where $\text{Diss}(\preceq_{\mathcal{KB}}^{\text{bi}})$ is the set of valuations that are dissonant for $\preceq_{\mathcal{KB}}^{\text{bi}}$. It can then be shown that $\preceq_{\mathcal{KB}}^{\text{int}}$ indeed forms an interval order that, provided $\preceq_{\mathcal{KB}}^{\text{bi}}$ is z -transitive (which it will be if $\circ$ is a ZTBOB NPR operator), preserves all non-empty optimal sets. Hence, by associating each $\mathcal{KB}$ with $\preceq_{\mathcal{KB}}^{\text{int}}$, we obtain an IOB revision operator $*$ and we get $\circ = \circ_{*, \mathcal{C}}$ for $* = *$ and $\mathcal{C} = \mathcal{C}^\circ$.

If $\preceq_{\mathcal{KB}}^{\text{bi}}$ is not merely z -transitive but fully transitive, then $\preceq_{\mathcal{KB}}^{\text{int}}$ above becomes a total preorder, which provides the key extra step towards the following characterisation of TBOB NPR operators.

Theorem 10. An operator $\circ$ is a TBOB NPR operator iff $\circ = \circ_{*, \mathcal{C}}$ for some CL structure $\langle *, \mathcal{C} \rangle$ s.t. $*$ is an AGM revision operator.

7 Related Work

As already mentioned, Rott (2014) was the first to have considered (strict) interval orders in belief change to define belief contraction operators. Booth and Varzinczak (2021; 2025) have utilised such orderings as the semantic foundation for defeasible conditionals, leading to the identification of a relevant notion of non-monotonic consequence relation lying between preferential closure and rational closure (Lehmann and Magidor 1992).

An investigation of the interdefinability of Rott’s contraction operators and our IOB revision via the Levi and Harper identities is an exercise we shall leave for future work.

In addition to the work of Rott, there have been several other studies in belief change that deploy kinds of ordering that are different from total preorders, starting with Kat-suno and Mendelzon (1991), who considered strict partial orders with an extra restriction (essentially the condition (5) from Section 3.2) to ensure Vacuity. Their work was generalised by Benferhat et al. (2005), who used unrestricted strict partial orders. Semiorders, which are a specific kind of (strict) interval order which correspond to interval-based interpretations in which all non-zero length intervals have the same length, have been considered for belief contraction by Rott (2014) and for both revision and contraction by Pappas and Williams (2014). The latter again considered a restricted family of semiorders designed to ensure the Vacuity postulate held, a restriction that was lifted by León and Pino Pérez (2021). Finally, Gaigne et al. (2024) recently looked at a novel class of orderings called partitioned linear orders. All of the above papers offer representation results for the resulting classes of revision operators.

With regard to credibility-limited revision, to the best of our knowledge, Schwind (2025) was the first to have challenged the single-sentence closure condition in this setting. Building on Booth et al.’s (2012) work, he motivates and introduces a family of *context-based* (CB) revision operators. These capture the intuition that a sentence may be deemed credible on the grounds that it strengthens (or provides extra context for) one of its non-credible consequences.

In terms of postulates, Schwind’s proposal consists of replacing Success Monotonicity in Booth et al.’s original proposal, which formalises the single-sentence closure condition, with a set of weaker postulates. Semantically, the belief revision operators are defined in terms of total preorders, as in the standard AGM approach to belief change, with the difference that the orderings are informed by a further refinement of the set of possible valuations.

Besides the fact that here we consider alternatives to total preorders in belief revision, our approach to credibility-limited revision differs from Schwind’s in other key aspects. First, the CB revision with an input α that is consistent with the agent’s initial beliefs systematically reduces to belief *expansion*, i.e., the result of revision is $Cn(\mathcal{KB} \cup \{\alpha\})$. This corresponds to the Vacuity (P4) postulate, which, as we have seen in Example 2, does not hold in general in our approach, but which can be enforced if desired, via the condition in (5). In that sense, our construction offers the user a choice.

Second, one of the core postulates in Schwind’s proposal is the Contextual Monotonicity (CM) principle, which states that if strengthening a credible sentence results in a non-credible one, then every further strengthening thereof remains incredible. We believe this is a rigid stricture to impose on the behaviour of the revision operator, as the following example, extending the one by Schwind (2025), shows.

Example 6. Let p be “Bob had lunch with the prime minister yesterday” (initially non-credible), let q be “the prime minister visited Bob’s institute” (and then $p \wedge q$ becomes credible), and let r be “Bob had dinner with the prime minister yesterday” (so that $p \wedge q \wedge r$ is non-credible). In Schwind’s approach, every $\alpha \in \mathcal{L}$ s.t. $\alpha \models p \wedge q \wedge r$ remains non-credible. However, let s be “they’ve discussed dinner plans

over lunch”, so that $p \wedge q \wedge r \wedge s$ ought to be credible.

It can be shown that the principles underlying Schwind’s Contextual Monotonicity do not hold in our framework, in general, and that our CL Structures from Definition 15 can handle cases such as the one in Example 6 correctly, as already suggested by Figure 4.

8 Concluding Remarks

Recapitulating the main contributions of the present paper, they can be summarised as follows: (i) the definition of a family of (successful) belief revision operators, each one informed by a special type of *biorder*, something that had not been sufficiently explored in the belief-change literature before, (ii) the provision of axiomatic characterisations *à la* AGM for the respective operators, along with an analysis of the postulates characterising each of them, in particular of the conditions sufficient for enforcing those that do not hold in each of the new settings, (iii) the deployment of the biorder-based revision operators in non-prioritised revision for recovering consistency and a provision of the respective representation results, and (iv) the formulation of biorder-based non-prioritised revision operators as credibility-limited revision with an accompanying interdefinability result, delivering an intuitive and versatile relaxation of the Single Sentence Closure principle.

Among the consequences of the above contributions to knowledge representation and reasoning is the fact that potential users now have a larger set of (justifiable) choices for belief revision at their disposal. Furthermore, the utilisation of biorders in belief revision opens up a broad avenue for exploration in closely-related areas, e.g. in reasoning with defeasible conditionals and in normative reasoning.

Immediate topics for further investigation stemming from the results in this paper include: (i) the identification of a meaningful strengthening as well as of an interesting weakening of the postulates on the credible set $\mathcal{C}$, (ii) a formulation of *general* BOB revision, besides the ZTBOB and TBOB ones, into the framework of credibility-limited revision, and (iii) the definition and study of BOB, ZTBOB and TBOB *contraction* and of how these stand in comparison to Rott’s contraction operators.

A further research question relates to the use of the full expressivity of biorders. As we have seen, different biorders can yield the same revision behaviour. We conjecture the answer to where this extra structure comes from may emerge from an investigation of the iterated revision case.

Acknowledgments

Thanks are due to Xavier Parent for interesting discussions and to the anonymous reviewers for their useful comments.

AI Declaration

The authors used Copilot to suggest some of the terminology used in the paper, and to check for any mistakes in the proofs once they had been written up.

References

- Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.
- Aleskerov, F.; Bouyssou, D.; and Monjardet, B. 2007. *Utility maximization, choice and preference*, volume 16. Springer Science & Business Media.
- Benferhat, S.; Lagrue, S.; Papini, O.; et al. 2005. Revision of partially ordered information: Axiomatization, semantics and iteration. In *IJCAI*, volume 5, 376–381.
- Booth, R., and Varzinczak, I. 2021. Conditional inference under disjunctive rationality. In Leyton-Brown, K., and Mausam., eds., *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 6227–6234. AAAI.
- Booth, R., and Varzinczak, I. 2025. On the disjunctive rational closure of a conditional knowledge base. *Artificial Intelligence* 348.
- Booth, R.; Fermé, E.; Konieczny, S.; and Pérez, R. P. 2012. Credibility-limited revision operators in propositional logic. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012*. AAAI Press.
- Fermé, E. 1999. *Revising the AGM Postulates*. Ph.D. Dissertation, Universidad de Buenos Aires.
- Fishburn, P. 1985. *Interval Orders and Interval Graphs: A Study of Partially Ordered Sets*. Wiley.
- Gaigne, J.; Belahcène, K.; and Lagrue, S. 2024. Partitioned linear orders and belief revision. In *36th International Conference on Tools with Artificial Intelligence (ICTAI 2024)*, 49–56. IEEE.
- Gerasimou, G. 2018. Indecisiveness, undesirability and overload revealed through rational choice deferral. *The Economic Journal* 128(614):2450–2479.
- Grossi, D.; van der Hoek, W.; and Kuijer, L. B. 2022. Reasoning about general preference relations. *Artif. Intell.* 313:103793.
- Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17:157–170.
- Hansson, S. O.; Fermé, E. L.; Cantwell, J.; and Falappa, M. A. 2001. Credibility limited revision. *The Journal of Symbolic Logic* 66(4):1581–1596.
- Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 3(52):263–294.
- Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.
- León, M. V., and Pérez, R. P. 2021. Orders and belief revision. In *19th International Workshop on Non-Monotonic Reasoning*, 11–20.
- Myers, A. 1999. Basic interval orders. *Order* 16:261–275.
- Peppas, P., and Williams, M.-A. 2014. Belief change and semiorders. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR 2014)*, 160–167.

- Rott, H. 2014. Four floors for the theory of theory change: The case of imperfect discrimination. In *Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA)*, number 8761 in LNCS, 368–382. Springer.
- Salant, Y., and Rubinstein, A. 2008. (A, f): Choice with frames. *The Review of Economic Studies* 75(4):1287–1296.
- Schwind, N. 2025. Context-based belief revision. In *Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning (KR'25)*, 534–543.