

How to Revise a Total Preorder

Richard Booth · Thomas Meyer

Received: 24 August 2010 / Accepted: 22 December 2010
© Springer Science+Business Media B.V. 2011

Abstract Most approaches to *iterated belief revision* are accompanied by some motivation for the use of the proposed revision operator (or family of operators), and typically encode enough information in the *epistemic state* of an agent for uniquely determining one-step revision. But in those approaches describing a *family* of operators there is usually little indication of how to proceed uniquely after the first revision step. In this paper we contribute towards addressing that deficiency by providing a formal framework which goes beyond the first revision step in two ways. First, the framework is obtained by enriching the epistemic state of an agent starting from the following intuitive idea: we associate to each world x two abstract objects x^+ and x^- , and we assume that, in addition to preferences over the set of worlds, we are given preferences over this set of objects as well. The latter can be considered

This paper is a combined and extended version of papers which first appeared in the proceedings of KR 2006, the 10th International Conference on Principles of Knowledge Representation and Reasoning [8], and ECSQARU 2007, the 9th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty [10].

R. Booth (✉)

Faculty of Science, Technology and Communication, University of Luxembourg,
6, Rue Richard Coudenhove-Kalergi, Luxembourg 1359, Luxembourg
e-mail: richard.booth@uni.lu

T. Meyer

CSIR Meraka Institute, PO Box 395, Pretoria 0001, South Africa

T. Meyer

School of Computer Science, University of KwaZulu-Natal,
KwaZulu-Natal, South Africa
e-mail: tommie.meyer@meraka.org.za

as meta-information encoded in the epistemic state which enables us to go beyond the first revision step of the revision operator being applied, and to obtain a unique set of preferences over worlds. We then extend this framework to consider, not only the revision of preferences over worlds, but also the revision of this extended structure itself. We look at some desirable properties for revising the structure and prove the consistency of these properties by giving a concrete operator satisfying all of them. Perhaps more importantly, we show that this framework has strong connections with two other types of constructions in related areas. Firstly, it can be seen as a special case of *preference aggregation* which opens up the possibility of extending the framework presented here into a full-fledged framework for preference aggregation and social choice theory. Secondly, it is related to existing work on the use of *interval orderings* in a number of different contexts.

1 Introduction

Total preorders (hereafter *tpos*) are used to represent preferences in many contexts. In particular, they are a common tool in *belief revision* [20, 24, 34]. In that setting they are taken to stand for plausibility orderings on the set of propositional worlds, which are used to represent the *dispositions* for change, or the *conditional beliefs* of an agent, and are encoded as part of the *epistemic state* of the agent. The associated *belief set* is taken to be the set of those sentences true in all the minimal worlds. When new evidence α comes in, the plausibility ordering is used to calculate the new belief set, usually by setting it to be the set of those sentences true in all the minimal models of α . This ensures a unique new belief set, but does not provide enough information to obtain a new tpo which may then serve as the target for the *next* revision input. Thus the question of modelling the dynamics of *tpos* is of critical importance to the problem of *iterated belief revision*.

The past fifteen years has seen a flurry of activity in this area with the work of Darwiche and Pearl [12], Nayak et al. [30], and Booth and Meyer [7] being representative examples. Most approaches devote considerable effort to motivating the use of their proposed revision operator (or family of operators). But in those approaches describing a family of operators there is usually little (or no) indication of how to choose among the available operators. In this paper we make a contribution towards overcoming that deficiency by providing a formal framework which obtains a unique *tpo* following one revision step, thereby going beyond just the belief set resulting from the revision input. This does not allow for choosing a unique operator, but is a step towards such a choice, since it uniquely identifies both the belief set and the tpo. The framework is obtained by enriching the epistemic state of the agent beyond a simple tpo, starting from the following intuitive idea: when we compare two different worlds x and y according to preference, often we are able to imagine different contingencies, according to whether all goes well in x and y or not. Our idea is to associate to each world x two abstract objects x^+

and x^- , with the intuition that x^+ represents x in *positive circumstances*, while x^- represents x in *negative circumstances*, and we assume that, *in addition to* the given tpo \leq over the set of worlds, we are given a tpo \preceq over this set of objects.

This meta-information allows us to uniquely determine the new tpo: when new evidence α comes in it casts a more favourable light on those worlds in which α holds. Thus the evidence signals the use of the positive versions of the worlds satisfying α , and the use of the negative versions of the $\neg\alpha$ -worlds. The revised tpo \leq_α^* is obtained by setting $x \leq_\alpha^* y$ iff $x^\epsilon \preceq y^\delta$, with $\epsilon, \delta \in \{+, -\}$ depending on whether x, y satisfy α or not.

As we will see, one commonly assumed rule from belief revision which will *not* generally hold for our revision operators is that the input α is necessarily an element of the *belief set* associated to \leq_α^* . Thus, at the belief set level, we are in the realm of so-called *non-prioritised* revision [21, 22].

Although the approach described allows us to determine more than just the belief set associated with an epistemic state, there is a problem with this approach regarding iterated tpo-revision. While the extra structure tells us how to determine a new tpo, it tells us nothing about how to determine the new extra structure which is needed to guide the next revision. Clearly the problem of iterated belief revision has simply re-emerged “one level up”. We investigate this problem by considering some desirable properties for revising the extra structure, and prove the consistency of these properties by giving a concrete operator satisfying all of them.

The plan of the paper is as follows. In Section 2 we give a brief introduction to the influential approach to iterated belief revision proposed by Darwiche and Pearl [12]. This is followed, in Section 3, by describing our enriched epistemic state. Then, in Section 4, we show how to use this enrichment to define a unique tpo-revision operator, and we axiomatically characterise the resulting family of operators. Initially we describe the properties of this family on a *semantic* level, i.e., in terms of how the ordering of individual worlds x, y undergo change. We show that the framework presented here can be viewed as a special case of *preference aggregation* or *social choice theory* [3]. This opens up the possibility of extending the framework presented here into a full-fledged framework for preference aggregation and social choice theory. In Section 5 we give an alternative, *sentential* formulation in terms of *conditional beliefs*, and introduce the notion of what it means for one sentence to *overrule* another in the context of a tpo-revision operator. In Section 6 we study notions of strict preference which can be extracted from \preceq and show how these are closely related to the overrules relation. In Section 7 we examine two known special cases of our family and give an example which shows how rigid use of either of these can sometimes lead to counter-intuitive results. In Section 8 we describe and axiomatise an interesting subclass of our family which remains general enough to include the two special cases, while in Section 9 we compare our general family with another family of tpo-revision operators which has recently been proposed, viz. the *improvement operators* of Konieczny et al. [25, 27]. In Section 10 we introduce an alternative way of representing the

\preceq orderings which we call *strict preference hierarchies* (SPHs). We point out the link between this representation and the use of interval orderings in various circumstances [2, 32]. We also show that these are equivalent to the \preceq orderings. In Section 11 we consider a few desirable properties which any good operator for revising SPHs should satisfy, before proving the consistency of these properties in Section 11.1 by providing an example of a concrete operator which is shown to satisfy them all. We conclude and mention ideas for further research in Section 12.

Preliminaries We work in a propositional language L generated by finitely many propositional variables, and with \top being the canonical representative of a tautology. We use \vdash and \equiv to denote classical logical consequence and classical logical equivalence respectively. We sometimes also use Cn to denote the operation of closure under classical logical consequence. W is the set of propositional worlds. Given $\alpha \in L$, we denote the set of worlds which satisfy α by $[\alpha]$. Given any set $S \subseteq W$ of worlds, $Th(S)$ will denote the set of sentences true in all the worlds in S . A tpo is a binary relation \leq which is both transitive and connected (for any x, y either $x \leq y$ or $y \leq x$).

2 Darwiche–Pearl Revision by Way of AGM Revision

Darwiche and Pearl [12] reformulated the AGM postulates [1] for belief revision to be compatible with their suggested approach to iterated revision. This necessitated a move from belief sets to *epistemic states*. Epistemic states, as envisaged by Darwiche and Pearl, are abstract entities containing all the information needed for coherent reasoning including, in particular, the strategy for belief revision which the agent wishes to employ at a given time. Thus, an epistemic state will include the belief set of an agent, a plausibility ordering (formally represented as a tpo on W), as well as any additional structure, which could include the enriched preference information we propose in this paper. Given the abstract nature of such epistemic states, it may well be possible to have different syntactic representations of, what is essentially, the same epistemic state.¹ In Darwiche and Pearl's reformulated postulates $*$ is a belief change operator on epistemic states, not belief sets. We denote by $B(\mathbb{E})$ the belief set extracted from an epistemic state \mathbb{E} . We use $B(\mathbb{E}) + \alpha$ to denote $Cn(B(\mathbb{E}) \cup \{\alpha\})$, i.e., the *expansion* of $B(\mathbb{E})$ by α .

- (\mathbb{E}^*1) $B(\mathbb{E} * \alpha) = Cn(B(\mathbb{E} * \alpha))$
- (\mathbb{E}^*2) $\alpha \in B(\mathbb{E} * \alpha)$
- (\mathbb{E}^*3) $B(\mathbb{E} * \alpha) \subseteq B(\mathbb{E}) + \alpha$
- (\mathbb{E}^*4) If $\neg\alpha \notin B(\mathbb{E})$ then $B(\mathbb{E}) + \alpha \subseteq B(\mathbb{E} * \alpha)$
- (\mathbb{E}^*5) If $\mathbb{E} = \mathbb{F}$ and $\alpha \equiv \beta$ then $B(\mathbb{E} * \alpha) = B(\mathbb{F} * \beta)$

¹Personal communication with Adnan Darwiche.

- (E*6) $\perp \in B(\mathbb{E} * \alpha)$ iff $\vdash \neg\alpha$
- (E*7) $B(\mathbb{E} * (\alpha \wedge \beta)) \subseteq B(\mathbb{E} * \alpha) + \beta$
- (E*8) If $\neg\beta \notin B(\mathbb{E} * \alpha)$ then $B(\mathbb{E} * \alpha) + \beta \subseteq B(\mathbb{E} * (\alpha \wedge \beta))$

Darwiche and Pearl then show, via a representation result similar to that of Katsuno and Mendelzon [24], that their version of revision on epistemic states can be represented in terms of plausibility orderings associated with epistemic states. More specifically, every epistemic state \mathbb{E} has associated with it a tpo $\leq_{\mathbb{E}}$ on W , with elements lower down in the ordering deemed more plausible. Moreover, for any two epistemic states \mathbb{E} and \mathbb{F} which are identical (but may be syntactically different), it has to be the case that $\leq_{\mathbb{E}} = \leq_{\mathbb{F}}$.

The other difference between the original AGM postulates and the Darwiche–Pearl reformulation—first inspired by a critical observation by Freund and Lehmann [17]—occurs in (E*5), which states that revising by logically equivalent sentences results in epistemic states with identical associated belief sets. This is a weakening of the original AGM postulate, phrased in our notation as follows:

- (B*5) If $B(\mathbb{E}) = B(\mathbb{F})$ and $\alpha \equiv \beta$ then $B(\mathbb{E} * \alpha) = B(\mathbb{F} * \beta)$

(B*5) states that two epistemic states with identical associated *belief sets* will, after having been revised by equivalent inputs, produce two epistemic states with identical associated belief sets. This is stronger than (E*5) which requires equivalent associated belief sets only if the original *epistemic states* were identical. As a consequence, (B*5) does *not* follow from the Darwiche–Pearl postulates.

In addition to these differences we introduce a minor modification of our own to the Darwiche–Pearl postulates. Let $\min(\alpha, \leq_{\mathbb{E}})$ denote the minimal models of α under $\leq_{\mathbb{E}}$. The belief set associated with the epistemic state is obtained by considering the minimal models in $\leq_{\mathbb{E}}$ i.e., $[B(\mathbb{E})] = \min(\top, \leq_{\mathbb{E}})$. Observe that this means that $B(\mathbb{E})$ has to be consistent. This requirement enables us to obtain a unique belief set from the total preorder $\leq_{\mathbb{E}}$, but it is incompatible with a successful revision by \perp . This requires that we jettison (E*6) and insist on consistent epistemic inputs only. (The left-to-right direction of (E*6) is rendered superfluous by (E*1) and the assumption that belief sets extracted from all epistemic states have to be consistent.) We shall refer to the reformulated AGM postulates, with (E*6) removed, as DP-AGM.

DP-AGM guarantees a unique extracted belief set when revision by α is performed. It sets $[B(\mathbb{E} * \alpha)]$ equal to $\min(\alpha, \leq_{\mathbb{E}})$ and thereby fixes the most plausible worlds in $\leq_{\mathbb{E} * \alpha}$. However, it places no restriction on the rest of the ordering. The purpose of the Darwiche–Pearl framework is to constrain this remaining part of the new ordering. It is done by way of a set of postulates for iterated revision [12]. (We follow the convention that $*$ is left associative.)

- (C1) If $\beta \vdash \alpha$ then $B(\mathbb{E} * \alpha * \beta) = B(\mathbb{E} * \beta)$
- (C2) If $\beta \vdash \neg\alpha$ then $B(\mathbb{E} * \alpha * \beta) = B(\mathbb{E} * \beta)$

- (C3) If $\alpha \in B(\mathbb{E} * \beta)$ then $\alpha \in B(\mathbb{E} * \alpha * \beta)$
- (C4) If $\neg\alpha \notin B(\mathbb{E} * \beta)$ then $\neg\alpha \notin B(\mathbb{E} * \alpha * \beta)$

The postulate (C1) states that when two pieces of information—one more specific than the other—arrive, the first is made redundant by the second. (C2) says that when two contradictory epistemic inputs arrive, the second one prevails; the second evidence alone yields the same belief set. (C3) says that a piece of evidence α should be retained after accommodating more recent evidence β that entails α given the current belief set. (C4) simply says that no epistemic input can act as its own defeater. The following are the corresponding semantic versions (with $v, w \in W$):

- (CR1) If $v \in [\alpha], w \in [\alpha]$ then $v \preceq_{\mathbb{E}} w$ iff $v \preceq_{\mathbb{E} * \alpha} w$
- (CR2) If $v \in [\neg\alpha], w \in [\neg\alpha]$ then $v \preceq_{\mathbb{E}} w$ iff $v \preceq_{\mathbb{E} * \alpha} w$
- (CR3) If $v \in [\alpha], w \in [\neg\alpha]$ then $v \prec_{\mathbb{E}} w$ only if $v \prec_{\mathbb{E} * \alpha} w$
- (CR4) If $v \in [\alpha], w \in [\neg\alpha]$ then $v \preceq_{\mathbb{E}} w$ only if $v \preceq_{\mathbb{E} * \alpha} w$

(CR1) states that the relative ordering between α -worlds remain unchanged following an α -revision, while (CR2) requires the same for $\neg\alpha$ -worlds. (CR3) requires that, for an α -world strictly more plausible than a $\neg\alpha$ -world, this relationship be retained after an α -revision, and (CR4) requires the same for weak plausibility. Darwiche and Pearl showed that, given DP-AGM, a precise correspondence obtains between (C_i) and (CR_i) above ($i = 1, 2, 3, 4$).

For the rest of the paper we assume a fixed but arbitrary initial tpo \preceq over W which we wish to revise. This tpo plays the role of the plausibility ordering over worlds introduced by Darwiche and Pearl into epistemic states. \prec will denote the strict part of \preceq , and \sim the symmetric closure of \preceq (i.e. $x \sim y$ iff both $x \preceq y$ and $y \preceq x$). We are interested in functions $*$ which, for each $\alpha \in L$, return a new ordering \preceq_{α}^* , and we will refer to any such $*$ as a *revision operator for \preceq* .

3 Enriching the Epistemic State

We let $W^{\pm} = \{x^{\epsilon} \mid x \in W \text{ and } \epsilon \in \{+, -\}\}$. We assume $x^{\epsilon} = y^{\delta}$ only if both $x = y$ and $\epsilon = \delta$. We suppose, along with \preceq , we are given some relation \preceq over W^{\pm} . The relation \preceq contains the additional information to be added to an epistemic state (already containing the plausibility ordering \preceq) when performing revision. We expect some basic conditions on \preceq and its interrelations with \preceq :

- (\preceq 1) \preceq is a tpo over W^{\pm}
- (\preceq 2) $x^+ \preceq y^+$ iff $x \preceq y$
- (\preceq 3) $x^- \preceq y^-$ iff $x \preceq y$
- (\preceq 4) $x^+ \prec x^-$

(\preceq 2) and (\preceq 3) say that the choice between the positive representations (negative representations respectively) of two worlds should be precisely the same as that dictated by \preceq . (\preceq 4) just says that given the choice between the positive and negative representations of x , we should choose the former over the latter.

Definition 1 Let $\preceq \subseteq W^\pm \times W^\pm$. If \preceq satisfies $(\preceq 1)$ – $(\preceq 4)$ we say \preceq is a \preceq -faithful tpo (over W^\pm).

From this definition it is easy to see that \preceq already contains enough information to determine \leq uniquely: simply observe how \preceq behaves when restricted to $\{x^+ \mid x \in W\}$ or $\{x^- \mid x \in W\}$. Strictly speaking, therefore, we need only include \preceq in an epistemic state.

The following result shows that we could equivalently replace $(\preceq 4)$ in this definition by a seemingly stronger property:


Proposition 1 Let $\preceq \subseteq W^\pm \times W^\pm$ be any relation satisfying $(\preceq 1)$ and at least one of $(\preceq 2)$ and $(\preceq 3)$. Then \preceq satisfies $(\preceq 4)$ iff it satisfies the following rule:

$$(\preceq 4') \quad x \preceq y \text{ implies } x^+ \prec y^-$$


Proof Let \preceq be as stated. That $(\preceq 4') \Rightarrow (\preceq 4)$ is clear. For the converse direction suppose $(\preceq 4)$ is satisfied and suppose $x \preceq y$. If \preceq satisfies $(\preceq 2)$ then this gives $x^+ \preceq y^+$. We have $y^+ \prec y^-$ by $(\preceq 4)$, so putting these two together using $(\preceq 1)$ gives the required $x^+ \prec y^-$. If \preceq satisfies $(\preceq 3)$ rather than $(\preceq 2)$ then $x \preceq y$ yields $x^- \preceq y^-$. We know $x^+ \prec x^-$ by $(\preceq 4)$ so putting these two together using $(\preceq 1)$ again gives $x^+ \prec y^-$. □

How can we picture these orderings \preceq ? One way was given by Booth et al. [8], using an assignment of numbers to a $2 \times n$ array, where n is the number of ranks according to the tpo associated to \preceq . In this paper we would like to use the alternative graphical representation introduced in [10] which is perhaps more intuitive, and is easier to work with when trying to construct examples. The idea is, for each $x \in W$, to think of the pair (x^+, x^-) as representing an *abstract interval* assigned to x . We can imagine that to each x we assign a “stick” whose left and right endpoints are x^+ and x^- respectively. Condition $(\preceq 1)$ says the endpoints of all these possible sticks are totally preordered. By $(\preceq 2)$ and $(\preceq 3)$ these sticks may be visualised as all having the same length, which $(\preceq 4)$ demands is non-zero. We may arrange the sticks in an order such as the one shown in Fig. 1, which shows the sticks associated to the five worlds x_1 – x_5 . The further to the left an endpoint is, the lower, i.e., more preferred, it is according to \preceq . Thus we see for example that $x_1^+ \prec x_3^+$ and $x_2^- \sim x_4^+$.

Fig. 1 Example of abstract interval ordering



Note that, although we said that all sticks have equal length, this assumption is not *absolutely* necessary. For example we could have depicted the \preceq in Fig. 1 also as



Here, in order for (\preceq_2) and (\preceq_3) to be satisfied, it is only necessary that the sticks associated to x_1 and x_2 are equal. However this assumption helps to simplify the visualisation, and so we will keep to it in the rest of the paper.

4 Revision Operators Defined from \preceq

Now given a \preceq -faithful tpo \preceq over W^\pm we want to use the information given by \preceq to define a revision operator $* = *_{\preceq}$ for \preceq . The idea is that the evidence α casts a favourable light on those worlds satisfying α . In other words, we consider worlds satisfying α to be associated with their positive representations, and worlds inconsistent with the evidence to be associated with their negative representations'. We set, for any $\alpha \in L$ and $x \in W$:


$$r_\alpha(x) = \begin{cases} x^+ & \text{if } x \in [\alpha] \\ x^- & \text{if } x \in [\neg\alpha] \end{cases}$$

The revised tpo \preceq_α^* is defined by setting, for each $x, y \in W$,

$$x \preceq_\alpha^* y \text{ iff } r_\alpha(x) \preceq r_\alpha(y).$$

Definition 2 For each \preceq -faithful tpo \preceq over W^\pm , we refer to $*_{\preceq}$ as defined above as the revision operator (for \preceq) *generated by* \preceq .

Example 1 In terms of our picture, each world gets mapped to one of the endpoints of the stick associated to it—left if it is an α -world and right if it is a $\neg\alpha$ -world. From this the new tpo \preceq_α^* may be read off. For example in Fig. 1 suppose we revise by α such that $x_4, x_5 \in [\alpha]$ and $x_1, x_2, x_3 \in [\neg\alpha]$. Then \preceq_α^* may be read off by looking at the black circles in the figure below.



So we see $x_1 \sim_\alpha^* x_2 \sim_\alpha^* x_4 <_\alpha^* x_3 <_\alpha^* x_5$.

We point out that if we look at the *belief set* associated to the new tpo \leq_{α}^* in this example then we see it does not contain the new evidence α due to the presence of the $\neg\alpha$ -worlds x_1 and x_2 among the minimal worlds in \leq_{α}^* . Thus we see that, at the level of belief sets, we are in the realm of so-called *non-prioritised* belief revision [22].

What are the properties of $*_{\leq}$? Consider the following list:

- (*1) \leq_{α}^* is a tpo over W
- (*2) $\alpha \equiv \gamma$ implies $\leq_{\alpha}^* = \leq_{\gamma}^*$
- (*3) If $x, y \in [\alpha]$ then $x \leq_{\alpha}^* y$ iff $x \leq y$
- (*4) If $x, y \in [\neg\alpha]$ then $x \leq_{\alpha}^* y$ iff $x \leq y$
- (*5) If $x \in [\alpha], y \in [\neg\alpha]$ and $x \leq y$ then $x <_{\alpha}^* y$
- (*6) If $x \in [\alpha], y \in [\neg\alpha]$ and $y \leq_{\alpha}^* x$ then $y <_{\gamma}^* x$
- (*7) If $x \in [\alpha], y \in [\neg\alpha]$ and $y <_{\alpha}^* x$ then $y <_{\gamma}^* x$

(*1) just says revising a tpo over W should result in another tpo over W , while (*2) is a syntax-irrelevance property. The next three rules are all familiar from the literature on iterated belief change. (*3) and (*4) are respectively identical to (CR1) and (CR2) in Section 2. (*5) was proposed independently by Booth and Meyer [7] and Jin and Thielscher [23]. It is easily seen to be stronger than the other two rules in the Darwiche–Pearl list (which can be obtained by replacing \leq by $<$ (CR3) and $<_{\alpha}^*$ by \leq_{α}^* (CR4) respectively). It says if an α -world x was considered at least as preferred as a $\neg\alpha$ -world y before receiving α , then after revision it should be considered *strictly* more preferred. These three rules were considered characteristic of a family of operators called *admissible* revision operators [7].

So far each of our rules mention only *one* revision input sentence α (modulo logical equivalence). By analogy with the AGM postulates for *belief set* revision [1], we might consider them as the set of *basic* postulates for tpo-revision. One thing largely missing from the literature on iterated belief change is a serious study of *supplementary* rationality properties which bestow a certain amount of coherence on the results of revising \leq by *different* sentences. The last couple of properties do this. First, suppose evidence α is received, and let $x \in [\alpha], y \in [\neg\alpha]$, but suppose $y \leq_{\alpha}^* x$. We propose that if x is not more preferred than y , *even after* receiving evidence which clearly points more to x being the case than it does to y , then there can be *no* evidence which will lead to x being more preferred to y . This is expressed by (*6). Similarly (*7) says if x is deemed *strictly* less preferred than y after receiving α then x must be strictly less preferred after receiving *any* input.

It turns out that these properties provide an exact characterisation of the revision operators we consider.

Theorem 1 *Let $*$ be any revision operator for \leq . Then $*$ is generated from some \leq -faithful tpo \leq over W^{\pm} iff $*$ satisfies (*1)–(*7).*

Proof

Soundness (*1) holds because of (≤ 1). (*2) holds because, as is easily seen, $\alpha \equiv \gamma$ implies $r_\alpha(x) = r_\gamma(x)$ for all $x \in W$. (*3) and (*4) hold as direct consequences of (≤ 2) and (≤ 3) respectively. (*5) holds as a consequence of ($\leq 4'$). For (*6) suppose $x \in [\alpha]$, $y \in [-\alpha]$ and $y \leq_\alpha^* x$. From the first two we know $r_\alpha(x) = x^+$ and $r_\alpha(y) = y^-$. Using these with $y \leq_\alpha^* x$ gives $y^- \leq x^+$. From this and (≤ 4) we have $y^+ < y^- \leq x^+ < x^-$. Thus, we see that for *any* $\gamma \in L$, we will have $r_\gamma(y) \leq r_\gamma(x)$, i.e., $y \leq_\gamma^* x$ as required. (*7) is proved similarly.

Completeness Starting from any revision operator $*$ for \leq we can define an ordering \leq_* over W^\pm as follows. Let $x, y \in W$ and $\delta, \epsilon \in \{+, -\}$. If $\delta = \epsilon$ then we set $x^\delta \leq_* y^\delta$ iff $x \leq y$. If $\delta \neq \epsilon$ we consider two cases: If $x = y$ then we simply set $x^+ <_* x^-$. Otherwise we set $x^+ \leq_* y^-$ iff $x \leq_x^* y$ and $x^- \leq_* y^+$ iff $x \leq_y^* y$. Here, when we use a world x as a subscript in \leq_x^* , we are using it to denote any sentence α such that $[\alpha] = \{x\}$. Likewise, in the proofs which follow, when x appears within the scope of a propositional connective, e.g., $x \vee y$ (note that $*$ satisfies (*2) the precise choice of α is irrelevant).

We need to show two things: (a) \leq_* is a \leq -faithful tpo, and then (b) the revision operator generated from \leq_* is precisely $*$.

- (a) \leq_* is a \leq -faithful tpo To show this we need to show that (≤ 1)–(≤ 4) are satisfied. (≤ 2), (≤ 3), and (≤ 4) obviously hold by construction. It remains to prove (≤ 1), i.e., \leq_* is a tpo.

\leq_ is connected* We need to show, for any $x, y \in W$ and $\epsilon, \delta \in \{+, -\}$, either $x^\epsilon \leq_* y^\delta$ or $y^\delta \leq_* x^\epsilon$. If $\delta = \epsilon$ this reduces to showing either $x \leq y$ or $y \leq x$ by construction of \leq_* , and this clearly holds since \leq is itself connected. So suppose $\delta \neq \epsilon$. Now if $x = y$ then the result holds since our construction ensures that *precisely one* of $x^\delta \leq x^\epsilon$ and $x^\epsilon \leq x^\delta$ holds (the former if $\delta = +$, the latter if $\epsilon = +$). So suppose also $x \neq y$. Then the construction tells us $x^\epsilon \leq y^\delta$ iff $x \leq_A^* y$ and $y^\delta \leq x^\epsilon$ iff $y \leq_A^* x$, where $A = x$ if $\epsilon = +$ while $A = y$ if $\delta = +$. Whatever the value of A we know \leq_A^* is connected by (*1), thus at least one of $x^\epsilon \leq y^\delta$ and $y^\delta \leq x^\epsilon$ must hold as required.

\leq_ is transitive* We need, for any $x, y, z \in W$ and $\delta, \epsilon, \nu \in \{+, -\}$,

$$\text{if } x^\delta \leq_* y^\epsilon \text{ and } y^\epsilon \leq_* z^\nu \text{ then } x^\delta \leq_* z^\nu.$$

For this proof let us denote these three by A, B, C respectively. Proving $A + B \Rightarrow C$ is a tedious matter of individually going through all eight combinations of choices for δ, ϵ, ν . The easiest cases are when $\delta = \epsilon = \nu = +$ or $\delta = \epsilon = \nu = -$, for in these cases showing $A + B \Rightarrow C$ reduces

to showing that $x \leq y$ and $y \leq z$ implies $x \leq z$, which clearly holds since \leq is itself transitive. Now let's go through the other six cases:

- (i) $\delta = \epsilon = +, \nu = -$. Firstly if $x = y$ then B and C reduce to the same thing and so the result holds. Also if $x = z$ then C holds by construction. So we assume $x \neq y$ and $x \neq z$. Then A becomes $x \leq y$. We now split into two subcases according to whether $y = z$. If $y = z$ then the target consequent C becomes $x \leq_x^* y$. But using $x \leq y$ with our assumption $x \neq y$ we may apply (*5) to deduce $x <_x^* y$. Thus C certainly holds. Now suppose $y \neq z$. Then $A + B \Rightarrow C$ reduces to showing $x \leq y + y \leq_y^* z \Rightarrow x \leq_x^* z$. Suppose for contradiction that $A + B$ holds but C does not. If C doesn't hold then $z <_x^* x$ by (*1) so, since we assume $z \neq x$, $z <_{x\nu y}^* x$ by (*7). From $x \leq y$ we get $x \leq_{x\nu y}^* y$ by (*3) and so $z <_{x\nu y}^* y$. Since we also assume $z \neq y$ we may apply (*7) to this to obtain $z <_y^* y$ – contradicting $y \leq_y^* z$. Hence the consequent must hold also in this case.
- (ii) $\delta = +, \epsilon = -, \nu = +$. Now B reduces to $y^- \leq_* z^+$ which means, by the already established (\leq_4), we must have $y \neq z$. Meanwhile C becomes $x^+ \leq_* z^+$, i.e., $x \leq z$. If $x = z$ then C clearly holds. So we assume also $x \neq z$. We now consider two subcases. Subcase 1: If $x = y$ then B becomes $x^- \leq_* z^+$, i.e., $x \leq_z^* z$ (since $x \neq z$). So $B \Rightarrow C$ by (*5). Subcase 2: If $x \neq y$ then A is $x \leq_x^* y$ and B is $y \leq_z^* z$, so we must show $x \leq_x^* y + y \leq_z^* z \Rightarrow x \leq z$. Assume for contradiction $A + B$ holds but C doesn't. From the latter $z < x$, then $z <_{x\nu z}^* x$ by (*3). Meanwhile, since $y \neq z$, the assumption $y \leq_z^* z$ gives $y \leq_{x\nu z}^* z$ by (*6). Hence $y <_{x\nu z}^* x$ using (*1). Since we assume $x \neq y \neq z$ we apply (*7) here to deduce $y <_x^* x$, contradicting $x \leq_x^* y$. Hence the consequent must hold.
- (iii) $\delta = +, \epsilon = -, \nu = -$. If $x = z$ then C becomes $x^+ \leq_* x^-$, which already holds by (\leq_4). Thus we assume $x \neq z$ and so C is $x \leq_x^* z$. Meanwhile B reduces to $y \leq z$. If $x = y$ then this reduces in turn to $x \leq z$ and so in this case we get $B \Rightarrow C$ by (*5). If $x \neq y$ then A is $x \leq_x^* y$ and so $A + B \Rightarrow C$ reduces to $x \leq_x^* y + y \leq z \Rightarrow x \leq_x^* z$. Assume for contradiction A and B hold and C does not. Then $z <_x^* x$ from $\text{not}C$ by (*1). Since we assume $y \neq x \neq z$ we may apply (*4) to $y \leq z$ to obtain $y \leq_x^* z$. Using this with $z <_x^* x$ and (*1) yields $y <_x^* x$, contradicting A . Hence C must follow from A and B .
- (iv) $\delta = -, \epsilon = +, \nu = +$. Here, A is $x^- \leq_* y^+$, which implies $x \neq y$ by (\leq_4) and so gives $x \leq_y^* y$. Meanwhile B is $y^+ \leq_* z^+$, which gives $y \leq z$. We first claim $A + B$ implies $x \neq z$. For if $x = z$ then B would give $y \leq x$. Since $x \neq y$ this then yields $y <_y^* x$ using (*5), contradicting the $x \leq_y^* y$ we obtained from A . Hence $x \neq z$ as claimed, and so given $A + B$, C becomes $x \leq_z^* z$. We must now show $x \leq_y^* y + y \leq z \Rightarrow x \leq_z^* z$. But since $x \neq y$ we may use (*6) to get $x \leq_{y\nu z}^* y$ from $x \leq_y^* y$, while $y \leq_{y\nu z}^* z$ from $y \leq z$ using (*3).

- $x \leq_{y \vee z}^* y$ and $y \leq_{y \vee z}^* z$ together give $x \leq_{y \vee z}^* z$ using (*1). From this, since $y \neq x \neq z$, we may apply (*6) to deduce $x \leq_z^* z$ as required.
- (v) $\delta = -, \epsilon = +, \nu = -$. As in the previous case, A implies $x \neq y$ and $x \leq_y^* y$, while this time C reduces to $x \leq z$. In the case $y = z$ then this in turn becomes $x \leq y$, which is a consequence of $x \leq_y^* y$ by (*5). Thus in this case $A \Rightarrow C$. So suppose instead $y \neq z$. Then B reduces to $y \leq_y^* z$ and so $A + B \Rightarrow C$ reduces to $x \leq_y^* y + y \leq_y^* z \Rightarrow x \leq z$. From $A + B$ and the transitivity of \leq_y^* it follows that $x \leq_y^* z$, from which it follows that $x \leq z$ using (*4) with the assumptions $x \neq y \neq z$.
- (vi) $\delta = -, \epsilon = -, \nu = +$. Now B yields $y \neq z$ (by (≤ 4)) and $y \leq_z^* z$, while A is equivalent to $x \leq y$. If $x = z$ were the case then this latter would become $z \leq y$ which would imply $z <_z^* y$ by (*5) (since $y \neq z$ from B). But this contradicts the $y \leq_z^* z$ we obtained from B and so we must have $x \neq z$. Hence, given $A + B$, C reduces to $x \leq_z^* z$ and so we must show $x \leq y + y \leq_z^* z \Rightarrow x \leq_z^* z$. But since $x \neq z \neq y$ we may use $x \leq y$ to deduce $x \leq_z^* y$ using (*4). From this and $y \leq_z^* z$ we obtain $x \leq_z^* z$ as required.
- (b) *The revision operator generated from \leq_* is precisely $*$.* Now let $*'$ be the revision operator generated from \leq_* . We now need to show $*'$ is precisely $*$, i.e., for any $\alpha \in L$ and $x, y \in W$, $x \leq_{\alpha}^* y$ iff $x \leq_{\alpha}^{\prime} y$. Since this latter is equivalent to $r_{\alpha}(x) \leq_* r_{\alpha}(y)$, this means we need to show that $x \leq_{\alpha}^* y$ iff $r_{\alpha}(x) \leq_* r_{\alpha}(y)$. We split into the three cases $x <^{\alpha} y$, $x \sim^{\alpha} y$ and $y <^{\alpha} x$. (Using the \leq^{α} -notation defined in Definition 3.)

Case $x <^{\alpha} y$ In this case $r_{\alpha}(x) = x^+$ and $r_{\alpha}(y) = y^-$. So we must show $x \leq_{\alpha}^* y$ iff $x^+ \leq_* y^-$. Since $x <^{\alpha} y$ we must have $x \neq y$ so by construction of \leq_* the right-hand side is equivalent to $x \leq_x^* y$. We will show $x \leq_{\alpha}^* y$ iff $x \leq_x^* y$. By (*1) this is equivalent to showing $y <_{\alpha}^* x$ iff $y <_x^* x$. But by (*7) each side of this biconditional is equivalent to $[y <_{\gamma}^* x \text{ for all } \gamma]$. Hence in this case the result holds.

Case $x \sim^{\alpha} y$ In this case we show that both $x \leq_{\alpha}^* y$ and $r_{\alpha}(x) \leq_* r_{\alpha}(y)$ are equivalent to $x \leq y$. That $x \leq_{\alpha}^* y$ iff $x \leq y$ follows from either (*3) or (*4) (depending on whether $x, y \in [\alpha]$ or $x, y \in [-\alpha]$ respectively). Meanwhile we have $r_{\alpha}(x) \leq_* r_{\alpha}(y)$ iff $x^{\delta} \leq_* y^{\delta}$ (where $\delta = +$ if $x, y \in [\alpha]$ and $\delta = -$ otherwise). By construction of \leq_* this latter is equivalent to $x \leq y$ as required.

Case $y <^{\alpha} x$ In this case we show that both $x \leq_{\alpha}^* y$ and $r_{\alpha}(x) \leq_* r_{\alpha}(y)$ are equivalent to saying $x \leq_{\gamma}^* y$ for all γ . For $x \leq_{\alpha}^* y$ this follows from (*6). Meanwhile $r_{\alpha}(x) \leq_* r_{\alpha}(y)$ iff $x^- \leq_* y^+$. Since $y <^{\alpha} x$ we know $x \neq y$ so by construction of \leq_* this latter is equivalent to $x \leq_{\gamma}^* y$. That this is equivalent to $[x \leq_{\gamma}^* y \text{ for all } \gamma]$ follows once more from (*6). □

4.1 The Link with Preference Aggregation and Social Choice Theory

In this subsection we discuss some more properties satisfied by our revision operators. These properties are recognisable as versions of properties familiar from the theory of *social choice*, or *preference aggregation* [3]. The problem of preference aggregation is the problem of finding some function f which, given any list of tpós (over some given set X of *alternatives*) \leq_1, \dots, \leq_n , with the \leq_i s representing the preferences over X of the *individuals* in a group, will return a new single ordering $f(\leq_1, \dots, \leq_n)$ over X which adequately represents the preferences of the *group* as a whole. Now, we can think of our problem of determining \leq_α^* as a highly specialised case of this problem. To do this we need to repackage the new evidence $\alpha \in L$ into tpo-form. The simplest way to do this is as follows.

Definition 3 For any $\alpha \in L$, the tpo \leq^α generated by α is the tpo over W given by $x \leq^\alpha y$ iff $x \in [\alpha]$ or $y \in [\neg\alpha]$.

In other words \leq^α is the tpo over W consisting of (at most) two ranks: the lower one containing all the α -worlds and the upper one containing all the $\neg\alpha$ -worlds. Then we can think of revision of \leq by α as an aggregation of \leq with \leq^α . (This manoeuvre is also carried out by Glaister [19] and Nayak [31]. An alternative way of generating tpós from sentences, based on the Hamming distance between two propositional worlds, is mentioned by Benferhat et al. [4].)

Many properties of preference aggregation operators have been proposed. One well-known property, known as the *Pareto* condition, is that, given two alternatives x and y , if every individual prefers x at least as much as y , and if at least one individual *strictly* prefers x over y , then the group should *strictly* prefer x over y . In our case, this condition translates into the following property:

(Pareto) If $x \leq y$ and $x \leq^\alpha y$, and at least one of these two inequalities is strict, then $x <_\alpha^* y$

The case of the above rule where \leq^α is strict is nothing other than (*5), while the case where $x \sim^\alpha y$ and $x < y$ is easily seen to follow mainly from (*3) and (*4). Thus we have:

Proposition 2 Every revision operator $*$ generated by some \leq -faithful tpo \leq over W^\pm satisfies (Pareto).

Another well-known property from preference aggregation, known as the *Independence of Irrelevant Alternatives*, states that for any two alternatives x and y , the group preference between x and y should depend only on how each individual ranks x and y . More precisely, if we were to replace individual i 's

$\text{tpo } \leq_i$ by any other $\text{tpo } \leq'_i$ which ranks x and y in exactly the same way as \leq , then x and y would be ranked in exactly the same way in $f(\leq_1, \dots, \leq'_i, \dots, \leq_n)$ as in $f(\leq_1, \dots, \leq_i, \dots, \leq_n)$. It turns out that our family of operators satisfy a restricted version of this rule, which we call *Independence of Irrelevant Alternatives in the Input*. We will make use of the following terminology:

Definition 4 Given $\alpha, \gamma \in L$, and $x, y \in W$, we say α and γ agree on x and y iff either both $x <^\alpha y$ and $x <^\gamma y$, or both $x \sim^\alpha y$ and $x \sim^\gamma y$, or both $y <^\alpha x$ and $y <^\gamma x$.

In other words α and γ agree on x and y if they both “say the same thing” regarding the relative plausibility of x and y .

(IIA-Input) If α and γ agree on x and y then $x \leq^*_\alpha y$ iff $x \leq^*_\gamma y$

That this is a property of our family of tpo-revision operators can be straightforwardly shown by considering an arbitrary \leq -faithful $\text{tpo } \leq$ over W^\pm . But in fact we can show the following:

Proposition 3 Let $*$ be any revision operator for \leq which satisfies (*1) and (*3)–(*5). Then $*$ satisfies (IIA-Input) iff $*$ satisfies both (*6) and (*7).

Proof Let $*$ satisfy (*1) and (*3)–(*5).

(IIA-Input) \Rightarrow (*6) + (*7). First we show the following property, which will be useful:

If $x <^\alpha y$, $y \leq^*_\alpha x$ and γ and α do not agree on x and y , then $y <^*_\gamma x$

To see this, first note if $x <^\alpha y$ and $y \leq^*_\alpha x$ then we must have $y < x$ by (*5). Also note if γ and α do not agree on x, y then, since $x <^\alpha y$, we must have either $x \sim^\gamma y$ or $y <^\gamma x$. In the first case we know $x \leq^*_\gamma y$ iff $x \leq y$ and $y \leq^*_\gamma x$ iff $y \leq x$ by (*3) and (*4). Using these with the already established $y < x$ gives $y <^*_\gamma x$ as required. In the case $y <^\gamma x$ we can use the fact $y < x$ to conclude $y <^*_\gamma x$ by (*5).

Now to show (*6) suppose $x <^\alpha y$ and $y \leq^*_\alpha x$. If γ and α do not agree on x, y then $y <^*_\gamma x$ by the above property, so $y \leq^*_\gamma x$ as required. If γ agrees with α on x, y then we can conclude $y \leq^*_\gamma x$ from $y \leq^*_\alpha x$ using (IIA-Input).

(*7) is proved similarly: Suppose $x <^\alpha y$ and $y <^*_\alpha x$. If γ does not agree with α on x, y then, since obviously $y <^*_\alpha x$ implies $y \leq^*_\alpha x$ we may apply the above proved property to conclude the required $y <^*_\gamma x$. If γ agrees with α on x, y then from (IIA-Input) we have $x \leq^*_\gamma y$ iff $x \leq^*_\alpha y$ and $y \leq^*_\gamma x$ iff $y \leq^*_\alpha x$. Hence we can conclude $y <^*_\gamma x$ from $y <^*_\alpha x$.

(*6) + (*7) \Rightarrow (IIA-Input). Suppose α and γ agree on x, y . To show (IIA-Input) it suffices by symmetry to show $x \leq^*_\alpha y$ implies $x \leq^*_\gamma y$. First suppose

$x <^\alpha y$, $x <^\gamma y$ and $x \leq_\alpha^* y$. If it were not the case that $x \leq_\gamma^* y$ then we would have $y <_\gamma^* x$ by (*1). Using this with $x <^\gamma y$ and (*7) would then give $y <_\alpha^* x$, contradicting $x \leq_\alpha^* y$. Hence we must have $x \leq_\gamma^* y$ as required. Now look at the case in which both $x \sim^\alpha y$ and $x \sim^\gamma y$. In this case using these with (*3) or (*4) we get $x \leq_\alpha^* y$ iff $x \leq y$ iff $x \leq_\gamma^* y$. Hence $x \leq_\alpha^* y$ implies $x \leq_\gamma^* y$ (and conversely) as required. Finally we consider the case in which both $y <^\alpha x$ and $y <^\gamma x$. This time we get $x \leq_\alpha^* y$ implies $x \leq_\gamma^* y$ (and conversely) using (*6). \square

Thus, given the basic properties (*1)–(*5) for tpo-revision, requiring $*$ to satisfy the two supplementary properties (*6) and (*7) amounts to enforcing (IIA-Input). Note this equivalence does not require the presence of the syntax-irrelevance property (*2). In fact, since sentences which are logically equivalent agree on *all* worlds x and y , we see that (*2) actually follows from (IIA-Input). Consequently, we have established that in the list (*1)–(*7), property (*2) is redundant.

For more discussion on social choice-like conditions and their relevance to tpo-revision we refer the reader to the work of Glaister [19].

5 On the Sentential Level

So far all our properties of tpo-revision operators have been expressed on the semantic level, directly in terms of worlds. But there is also a sentential level on which we can recast our properties. For any tpo \leq' over W and any $\beta \in L$ we let $\min(\beta, \leq')$ denote the set of \leq' -minimal elements of $[\beta]$, i.e., $\min(\beta, \leq') = \{x \in [\beta] \mid \nexists y \in [\beta] \text{ s.t. } y <' x\}$. Then we define:

$$\leq' \circ \beta = Th(\min(\beta, \leq')).$$

$\leq' \circ \beta$ represents what is believed in \leq' on the *supposition* that β is the case. If $\lambda \in \leq' \circ \beta$ then we might also say $\beta \mapsto \lambda$ is a *conditional belief* in \leq' . Note that we do not necessarily assume this is the same thing as saying λ would be believed after receiving β explicitly *as evidence*. This is because we want to support non-prioritised revision, so in particular β itself might not necessarily be believed after receiving it as evidence (it might be simply too far-fetched). Nevertheless, new evidence will have some impact on the set of conditional beliefs. Note that this notation enables us to denote the belief set associated to \leq' by $\leq' \circ \top$.

We can give all the properties (*2)–(*7) an equivalent formulation in terms of \circ , thus giving a set of sound and complete properties for our family of revision operators which has a different flavour:

- (o2) If $\alpha \equiv \gamma$ then $\leq_\alpha^* \circ \beta = \leq_\gamma^* \circ \beta$
- (o3) If $\beta \vdash \alpha$ then $\leq_\alpha^* \circ \beta = \leq \circ \beta$
- (o4) If $\beta \vdash \neg\alpha$ then $\leq_\alpha^* \circ \beta = \leq \circ \beta$
- (o5) If $\neg\alpha \notin \leq \circ \beta$ then $\alpha \in \leq_\alpha^* \circ \beta$

- (o6) If $\alpha \notin \leq_{\alpha}^* \circ \beta$ then $\alpha \notin \leq_{\gamma}^* \circ \beta$
- (o7) If $\neg\alpha \in \leq_{\alpha}^* \circ \beta$ then $\neg\alpha \in \leq_{\gamma}^* \circ \beta$

(o2) just says revising by logically equivalent sentences yields the same set of conditional beliefs. (o3) and (o4) are essentially the postulates (C1) and (C2) in Section 2, while (o5) corresponds to rule (P) of Booth and Meyer [7], also referred to as *Independence* by Jin and Thielscher [23]. The correspondences between these last three rules and their counterparts in the previous section were proved in those papers. (Although these papers all assume the prioritised setting for belief revision in which revision inputs are always believed after revision.) The last two rules are neatly explained with the help of the following terminology:

Definition 5 Given any revision operator $*$ for \leq and given $\alpha, \beta \in L$, we shall say β *overrules* α (relative to $*$) iff either β is inconsistent or $\alpha \notin \leq_{\alpha}^* \circ \beta$. We shall say β *strictly overrules* α (relative to $*$) iff $\neg\alpha \in \leq_{\alpha}^* \circ \beta$.

The inclusion of the clause “ β is inconsistent” in the definition of “overrules” allows for a smoother exposition. This way we get the intuitively expected chain of implications: $\beta \vdash \neg\alpha$ implies that β strictly overrules α , which implies that β overrules α . If $*$ satisfies (o5) then this in turn implies $\neg\alpha \in \leq \circ \beta$. Now suppose that evidence γ is received and we then make a further supposition β . (o6) says that if β overrules α and β is consistent then α will not be believed, while (o7) says that if β strictly overrules α then α will be *rejected*.

Proposition 4 Let $*$ be a revision operator for \leq which satisfies (*1). Then for each $i = 2, \dots, 7$, $*$ satisfies (*i) iff $*$ satisfies (oi).

Proof Suppose $*$ satisfies (*1), i.e., \leq_{α}^* is a tpo given any α . Note that for any tpo \leq' (in particular \leq_{α}^*) and $x, y \in W$, we have

$$x \leq' y \text{ iff } x \in \min(x \vee y, \leq') \tag{1}$$

where, recall, in the expression $x \vee y$, x and y stand for any sentences whose only model is x , respectively y , and so $[x \vee y] = \{x, y\}$.

(*2) \Leftrightarrow (o2) The left-to-right direction is obvious. For the converse direction suppose $\alpha \equiv \gamma$. Then using Eq. 1 we know, given any x, y , $x \leq_{\alpha}^* y$ iff $x \in \min(x \vee y, \leq_{\alpha}^*)$ and $x \leq_{\gamma}^* y$ iff $x \in \min(x \vee y, \leq_{\gamma}^*)$. But by (o2) $\min(x \vee y, \leq_{\alpha}^*) = \min(x \vee y, \leq_{\gamma}^*)$. Hence $x \leq_{\alpha}^* y$ iff $x \leq_{\gamma}^* y$ for all x, y , i.e., $\leq_{\alpha}^* = \leq_{\gamma}^*$ as required.

(*3) \Leftrightarrow (o3) and (*4) \Leftrightarrow (o4) Proofs given by Darwiche and Pearl [12] (Theorem 13).

(*5) \Leftrightarrow (o5) Proof given by Booth and Meyer [7] (Proposition 2) and Jin and Thielscher [23] (Theorem 5).

(*6) \Leftrightarrow (o6) For the left-to-right direction, suppose (*6) holds and suppose $\alpha \notin \leq_{\alpha}^* \circ \beta$. Then there exists $y \in [-\alpha] \cap \min(\beta, \leq_{\alpha}^*)$. Assume for contradiction $\alpha \in \leq_{\gamma}^* \circ \beta$. Then $y \notin \min(\beta, \leq_{\gamma}^*)$ so there exists $x \in \min(\beta, \leq_{\gamma}^*)$ such that $x <_{\gamma}^* y$. Since we assume $\alpha \in \leq_{\gamma}^* \circ \beta$ we must have $x \in [\alpha]$. Hence we may apply a contrapositive version of (*6) to obtain (with a little help from (*1)) $x <_{\alpha}^* y$. But this contradicts $y \in \min(\beta, \leq_{\alpha}^*)$. Hence it must be the case that $\alpha \notin \leq_{\gamma}^* \circ \beta$ as required.

For the converse direction, suppose (o6) holds and let $x \in [\alpha]$, $y \in [-\alpha]$ be such that $y \leq_{\alpha}^* x$. Then from Eq. 1, $y \in \min(x \vee y, \leq_{\alpha}^*)$. Hence, since $y \in [-\alpha]$, $\alpha \notin \leq_{\alpha}^* \circ (x \vee y)$. Using (o6) we infer $\alpha \notin \leq_{\gamma}^* \circ (x \vee y)$. Since necessarily $\min(x \vee y, \leq_{\gamma}^*) \subseteq \{x, y\}$, the only way we can have $\alpha \notin \leq_{\gamma}^* \circ (x \vee y)$ is if $y \in \min(x \vee y, \leq_{\gamma}^*)$, i.e., $y \leq_{\gamma}^* x$ as required to show (*6).

(*7) \Leftrightarrow (o7) For the left-to-right direction, suppose (*7) holds and suppose $\neg\alpha \in \leq_{\alpha}^* \circ \beta$. Suppose for contradiction $\neg\alpha \notin \leq_{\gamma}^* \circ \beta$. Then there exists $x \in [\alpha] \cap \min(\beta, \leq_{\alpha}^*)$. Since $\neg\alpha \in \leq_{\alpha}^* \circ \beta$ we know $x \notin \min(\beta, \leq_{\alpha}^*)$ so there exists $y \in \min(\beta, \leq_{\alpha}^*)$ such that $y <_{\alpha}^* x$. We know $y \in [-\alpha]$ since $\neg\alpha \in \leq_{\alpha}^* \circ \beta$, hence we may apply (*7) to deduce $y <_{\gamma}^* x$ – contradicting $x \in \min(\beta, \leq_{\alpha}^*)$. Hence $\neg\alpha \in \leq_{\gamma}^* \circ \beta$ as required.

For the converse direction, suppose (o7) holds and let $x \in [\alpha]$, $y \in [-\alpha]$ such that $y <_{\alpha}^* x$. Then $\min(x \vee y, \leq_{\alpha}^*) = \{y\}$ and so, since $y \in [-\alpha]$, $\neg\alpha \in \leq_{\alpha}^* \circ (x \vee y)$. Applying (o7) to this yields $\neg\alpha \in \leq_{\gamma}^* \circ (x \vee y)$ and so, since $x \in [\alpha]$, $x \notin \min(x \vee y, \leq_{\gamma}^*)$, i.e., $y <_{\gamma}^* x$ as required to show (*7). \square

Corollary 1 *Let $*$ be a revision operator for \leq . Then $*$ is generated from some \leq -faithful tpo \leq over W^{\pm} iff $*$ satisfies (*1) and (o2)–(o7).*

This sentential reformulation is useful since there are some interesting properties which can be formulated in sentential terms, but for which obvious semantic counterparts do not exist. For example:

$$\text{(Disj1)} \quad (\leq_{\alpha}^* \circ \beta) \cap (\leq_{\gamma}^* \circ \beta) \subseteq (\leq_{\alpha \vee \gamma}^* \circ \beta)$$

$$\text{(Disj2)} \quad (\leq_{\alpha \vee \gamma}^* \circ \beta) \subseteq (\leq_{\alpha}^* \circ \beta) \cup (\leq_{\gamma}^* \circ \beta)$$

These two properties were essentially first proposed by Schlechta et al. [35], and seem to be natural properties to have. The first one says if a conditional belief is held both after receiving evidence α and after receiving evidence γ , then it is also held after receiving their disjunction as evidence. The second one says a conditional belief is not held after receiving a disjunction as evidence, *without* being held after receiving just one of the disjuncts in isolation.

Proposition 5 *Every revision operator $*$ generated from some \leq -faithful tpo \leq over W^{\pm} satisfies (Disj1) and (Disj2).*

We prove this result by considering an arbitrary \leq -faithful tpo \leq , rather than trying to derive these rules syntactically from (*1) and (o2)–(o7). A key property used in the proof is that, for any $\alpha, \gamma \in L$ and $x \in W$, $r_{\alpha\vee\gamma}(x) = \min\{r_\alpha(x), r_\gamma(x)\}$.

Proof Let \leq be a given \leq -faithful tpo over W^\pm .

(Disj1) It suffices to show $\min(\beta, \leq_{\alpha\vee\gamma}^*) \subseteq \min(\beta, \leq_\alpha^*) \cup \min(\beta, \leq_\gamma^*)$. So let $x \in \min(\beta, \leq_{\alpha\vee\gamma}^*)$ and suppose for contradiction both $x \notin \min(\beta, \leq_\alpha^*)$ and $x \notin \min(\beta, \leq_\gamma^*)$. From these latter two we know there exist $y_1 \in \min(\beta, \leq_\alpha^*)$ and $y_2 \in \min(\beta, \leq_\gamma^*)$ such that $y_1 <_\alpha^* x$ and $y_2 <_\gamma^* x$. Equivalently $r_\alpha(y_1) < r_\alpha(x)$ and $r_\gamma(y_2) < r_\gamma(x)$. But since $x \in \min(\beta, \leq_{\alpha\vee\gamma}^*)$ we know $x \leq_{\alpha\vee\gamma}^* y_i$, equivalently $r_{\alpha\vee\gamma}(x) \leq r_{\alpha\vee\gamma}(y_i)$, for $i = 1, 2$. Since $r_{\alpha\vee\gamma}(y_i) = \min\{r_\alpha(y_i), r_\gamma(y_i)\}$ this means we have both $r_{\alpha\vee\gamma}(x) \leq r_\alpha(y_1)$ and $r_{\alpha\vee\gamma}(x) \leq r_\gamma(y_2)$. But since $r_{\alpha\vee\gamma}(x) = \min\{r_\alpha(x), r_\gamma(x)\}$ we know $r_{\alpha\vee\gamma}(x)$ is equal to either $r_\alpha(x)$ or $r_\gamma(x)$. In the first case we get $r_\alpha(x) \leq r_\alpha(y_1)$, contradicting $r_\alpha(y_1) < r_\alpha(x)$. In the second case we obtain $r_\gamma(x) \leq r_\gamma(y_2)$, contradicting $r_\gamma(y_2) < r_\gamma(x)$. Thus in either case we arrive at the required contradiction.

(Disj2) We first claim the following: Given any pair of worlds y_1, y_2 such that $y_1 \in \min(\beta, \leq_\alpha^*)$ and $y_2 \in \min(\beta, \leq_\gamma^*)$, at least one of these worlds must be in $\min(\beta, \leq_{\alpha\vee\gamma}^*)$. For suppose neither is an element of this set. Then there must exist $z \in \min(\beta, \leq_{\alpha\vee\gamma}^*)$ such that $z <_{\alpha\vee\gamma}^* y_i$, equivalently $r_{\alpha\vee\gamma}(z) < r_{\alpha\vee\gamma}(y_i)$, for $i = 1, 2$. Since $r_{\alpha\vee\gamma}(y_i) = \min\{r_\alpha(y_i), r_\gamma(y_i)\}$ we obtain from this both $r_{\alpha\vee\gamma}(z) < r_\alpha(y_1)$ and $r_{\alpha\vee\gamma}(z) < r_\gamma(y_2)$. Then since $r_{\alpha\vee\gamma}(z) = \min\{r_\alpha(z), r_\gamma(z)\}$ we get from these either $r_\alpha(z) < r_\alpha(y_1)$ or $r_\gamma(z) < r_\gamma(y_2)$. But in the former case we have $z <_\alpha^* y_1$, contradicting $y_1 \in \min(\beta, \leq_\alpha^*)$, while similarly in the latter case $z <_\gamma^* y_2$, which contradicts $y_2 \in \min(\beta, \leq_\gamma^*)$. Hence no such z can exist and so the claim must be true. This then allows us to show (Disj2), for suppose both $\lambda \notin \leq_\alpha^* \circ \beta$ and $\lambda \notin \leq_\gamma^* \circ \beta$. Then there must exist $y_1 \in \min(\beta, \leq_\alpha^*)$ and $y_2 \in \min(\beta, \leq_\gamma^*)$ such that $y_i \in [-\lambda]$ for $i = 1, 2$. From the above claim we know $y_i \in \min(\beta, \leq_{\alpha\vee\gamma}^*)$ for either $i = 1$ or $i = 2$. Either way we end up with some $y \in \min(\beta, \leq_{\alpha\vee\gamma}^*)$ such that $y \in [-\lambda]$, which is enough to prove $\lambda \notin \leq_{\alpha\vee\gamma}^* \circ \beta$. \square

The next result shows that $\leq_\alpha^* \circ \beta$ falls neatly into one of three categories. Note that we don't need (o6) and (o7), nor do we need (o2) for this.

Proposition 6 *Let $*$ be any revision operator for \leq satisfying (*1) and (o3)–(o5), and let the overrules relations be given relative to $*$. Then for all $\alpha, \beta \in L$,*

$$\leq_\alpha^* \circ \beta = \begin{cases} \leq(\alpha \wedge \beta) & \text{if } \beta \text{ doesn't overrule } \alpha \\ (\leq \circ(\alpha \wedge \beta)) \cap (\leq \circ \beta) & \text{if } \beta \text{ overrules } \alpha, \text{ but not strictly} \\ \leq \circ \beta & \text{if } \beta \text{ strictly overrules } \alpha \end{cases}$$

Proof We make use of the following standard properties, which hold for *any* tpo \leq' over W (note the assumption (*1) is satisfied permits us to apply these properties to \leq'_α):

- (i) If $\alpha \in \leq' \circ \beta$ then $\leq' \circ \beta = \leq' \circ (\alpha \wedge \beta)$ (Cumulativity).
- (ii) If $\neg\alpha \notin \leq' \circ \beta$ then $\leq' \circ \beta \subseteq \leq' \circ (\alpha \wedge \beta)$ (Rational Monotony).
- (iii) $(\leq' \circ \beta_1) \cap (\leq' \circ \beta_2) \subseteq (\leq' \circ (\beta_1 \vee \beta_2))$ (Or).

Suppose β does not overrule α . We must show $\leq^*_\alpha \circ \beta = \leq \circ (\alpha \wedge \beta)$. But if β does not overrule α then $\alpha \in \leq^*_\alpha \circ \beta$ so, using property (i) above, $\leq^*_\alpha \circ \beta = \leq^*_\alpha \circ (\alpha \wedge \beta)$. Using (o3) we conclude $\leq^*_\alpha \circ \beta = \leq \circ (\alpha \wedge \beta)$ as required.

Now suppose β strictly overrules α . We must show in this case $\leq^*_\alpha \circ \beta = \leq \circ \beta$. Firstly, if β is inconsistent then both these sets are equal to the entire set of sentences L and so the result clearly holds. So we assume β is consistent. We will in fact show both $\leq^*_\alpha \circ \beta$ and $\leq \circ \beta$ are equal to $\leq \circ (\neg\alpha \wedge \beta)$. For the former, if β strictly overrules α then $\neg\alpha \in \leq^*_\alpha \circ \beta$ so, again using (i) above, $\leq^*_\alpha \circ \beta = \leq^*_\alpha \circ (\neg\alpha \wedge \beta)$. Using (o4) we then obtain $\leq^*_\alpha \circ \beta = \leq \circ (\neg\alpha \wedge \beta)$ as required. Meanwhile from $\neg\alpha \in \leq^*_\alpha \circ \beta$ and the assumption β is consistent we can infer $\alpha \notin \leq^*_\alpha \circ \beta$. From this and (o5) we get $\neg\alpha \in \leq \circ \beta$ and so, from (i) once more, also $\leq \circ \beta = \leq \circ (\neg\alpha \wedge \beta)$ as required.

Finally we check the intermediate case where β overrules α , but not strictly, which means $\alpha, \neg\alpha \notin \leq^*_\alpha \circ \beta$. We must show $\leq^*_\alpha \circ \beta = (\leq \circ (\alpha \wedge \beta)) \cap (\leq \circ \beta)$. Looking back at the last sentence of the previous paragraph, we see we showed there that if $\alpha \notin \leq^*_\alpha \circ \beta$ then $\leq \circ \beta = \leq \circ (\neg\alpha \wedge \beta)$. Hence we may equivalently formulate our target identity as $\leq^*_\alpha \circ \beta = (\leq \circ (\alpha \wedge \beta)) \cap (\leq \circ (\neg\alpha \wedge \beta))$. Applying (o3) and (o4), this in turn is the same as requiring $\leq^*_\alpha \circ \beta = (\leq^*_\alpha \circ (\alpha \wedge \beta)) \cap (\leq^*_\alpha \circ (\neg\alpha \wedge \beta))$. We can prove the right-to-left inclusion here by noting by property (iii) above that $(\leq^*_\alpha \circ (\alpha \wedge \beta)) \cap (\leq^*_\alpha \circ (\neg\alpha \wedge \beta)) \subseteq \leq^*_\alpha \circ ((\alpha \wedge \beta) \vee (\neg\alpha \wedge \beta)) = \leq^*_\alpha \circ \beta$. For the converse direction note that from property (ii) we have $\alpha \notin \leq^*_\alpha \circ \beta$ implies $\leq^*_\alpha \circ \beta \subseteq \leq^*_\alpha \circ (\neg\alpha \wedge \beta)$ and $\neg\alpha \notin \leq^*_\alpha \circ \beta$ implies $\leq^*_\alpha \circ \beta \subseteq \leq^*_\alpha \circ (\alpha \wedge \beta)$. Thus the result holds. □

Thus if β doesn't overrule α then making the supposition β after receiving α as evidence is the same as supposing α and β together in the initial tpo \leq . If β strictly overrules α then evidence α is just ignored when making the further supposition β . In the intermediate case where β overrules α , but not strictly, supposing β following evidence α results in a mixture of these two.

In particular note what happens when $\beta \equiv \top$. We see that $\leq^*_\alpha \circ \top$ equals either (i) $\leq \circ \alpha$, or (ii) $(\leq \circ \alpha) \cap (\leq \circ \top)$, or (iii) $\leq \circ \top$. Thus either the evidence is fully incorporated into the belief set using the AGM revision operator corresponding to \leq [24] (case (i)), or the belief set remains unchanged (case (iii)), or there is an intermediate possibility (case (ii)), which amounts to removing $\neg\alpha$ from the initial belief set using the AGM contraction operator corresponding to \leq . That is, we don't commit to believing the evidence, but we leave open the possibility that it might hold. We will have more to say on these notions of overruling in the next section.

6 Notions of Strict Preference

In this section we shall assume a fixed \leq -faithful tpo \preceq over W^\pm . From a single \preceq we can extract *three* different notions of *strict preference* over W . First we have the simple one given by

$$x < y \text{ iff } x^+ < y^+$$

(equivalently $x < y$ iff $x^- < y^-$), i.e., $<$ is just the strict part of the tpo over W associated to \preceq . In terms of our graphical representation, $x < y$ iff the stick corresponding to x lies to the left of that associated to y , but *possibly with some overlap*. For example in Fig. 2 we have $x_1 < x_3$.

A second, stronger notion of strict preference can be expressed by:

$$x \lll y \text{ iff } x^- < y^+.$$

In other words, $x \lll y$ iff x , *even when represented negatively*, is preferred to y or, in terms of the picture, iff the stick associated to x lies *completely to the left* of that associated to y , and furthermore there is “daylight” between them. E.g., in Fig. 2 we see $x_2 \lll x_5$.

Finally a third case, intermediate between \lll and $<$, can be expressed by:


$$x \ll y \text{ iff } x^- \preceq y^+.$$

In other words $x \ll y$ iff x being represented negatively is *at least as* preferred to y . This third case captures a “hesitation” [32] between strong strict preference \lll and mere ordinary strict preference $<$. We will have $x \ll y$ and $x \lll y$ precisely when the right endpoint x^- of the x -stick and the left endpoint y^+ of the y -stick are vertically aligned with each other. E.g., in Fig. 2 we have $x_1 \ll x_4$ but $x_1 \lll x_4$. The next proposition collects some properties of these orderings.

Proposition 7

- (i) $\lll \subseteq \ll \subseteq <$ (where recall $<$ is the strict part of the initial tpo \preceq).
- (ii) \lll and \ll are both strict partial orders (i.e., irreflexive and transitive).
- (iii) \lll and \ll both satisfy the filtered condition [16], i.e., for all $x, y \in W$ and $\beta \in L$, if $x, y \in [\beta] \setminus \min(\beta, <')$ then there exists $z \in [\beta]$ such that $z <' x$ and $z <' y$.

Fig. 2 Example of abstract interval ordering



(Recall for a strict partial order $<'$, $\min(\beta, <') = \{x \in [\beta] \mid \nexists y \in [\beta] \text{ s.t. } y <' x\}$.)

Proof

- (i) The inclusion $\lll \subseteq \ll$ is immediate. The inclusion $\ll \subseteq <$ follows from $(\leq 4')$.
- (ii) The irreflexivity of \ll follows directly from the inclusion in (i). Since $\lll \subseteq \ll$ this means \lll must be irreflexive as well. To show transitivity of the two relations, we actually show something stronger holds, namely

$$\text{If } x \ll y \text{ and } y \ll z \text{ then } x \lll z. \tag{2}$$

This is true since if $x \ll y$ and $y \ll z$ then $x^- \preceq y^+$ and $y^- \preceq z^+$. Since $y^+ < y^-$ by (≤ 4) we obtain $x^- \preceq y^+ < y^- \preceq z^+$, thus $x^- < z^+$, i.e., $x \lll z$ as claimed. Equation 2 yields the transitivity of both \lll and \ll using the fact $\lll \subseteq \ll$.

- (iii) To show \lll satisfies the filtered condition let $x, y \in [\beta] \setminus \min(\beta, \lll)$. Since x, y are not minimal there exist $z_1, z_2 \in [\beta]$ such that $z_1 \lll x$ and $z_2 \lll y$, i.e., $z_1^- < x^+$ and $z_2^- < y^+$. Since \preceq is connected (since it is a tpo by (≤ 1)) we know either $z_1^- \preceq z_2^-$ or $z_2^- \preceq z_1^-$. In the first case we obtain $z_1^- < y^+$ from $z_2^- < y^+$ and so there exists some $z \in [\beta]$ (namely z_1) such that both $z \lll x$ and $z \lll y$ as required. In the second case we obtain $z_2^- < x^+$ from $z_1^- < x^+$ and so again we find a z (this time $z = z_2$) with the required properties. Hence \lll satisfies the filtered condition. The case for \ll is analogous. □

By (i) we see $<$, \ll and \lll form progressively more stringent notions of strict preference. If we let $* = *_{\preceq}$ then we see $x \lll y$ implies $r_{\gamma}(x) < r_{\gamma}(y)$ for all $\gamma \in L$, and so $x <_{\gamma}^* y$ for any γ . Thus \lll can also be viewed as a set of *core*, or *protected* strict preferences in $<$ which are always preserved in any revision. Meanwhile we have $x \ll y$ implies $x \leq_{\gamma}^* y$ for any γ . Thus \ll may be viewed as a set of *weakly protected* strict preferences, in the sense that if $x \ll y$ then no evidence will ever cause this preference to be reversed.

It turns out that these relations \ll and \lll are closely related to the notions of overruling and strict overruling from Definition 5.

Proposition 8 *Let the overrules relations be given relative to $*_{\preceq}$. Then (i) β overrules α iff $\min(\beta, \ll) \subseteq [-\alpha]$. (ii) β strictly overrules α iff $\min(\beta, \lll) \subseteq [-\alpha]$.*

Proof

- (i) We must show that $\min(\beta, \ll) \subseteq [-\alpha]$ iff either β is inconsistent or $\alpha \notin \leq_{\alpha}^* \circ \beta$.
 \Rightarrow Suppose $\min(\beta, \ll) \subseteq [-\alpha]$. If β is inconsistent we are done, so assume β is consistent. We must show $\alpha \notin \leq_{\alpha}^* \circ \beta$, i.e., $\min(\beta, \leq_{\alpha}^*) \cap [-\alpha] \neq \emptyset$. Suppose $\min(\beta, \ll) = \{y_1, \dots, y_k\}$. Since $\min(\beta, \ll) \subseteq$

$[\neg\alpha]$ we know $y_i \in [\neg\alpha]$ for all $i = 1, \dots, k$. We will show that at least one of these elements of $\min(\beta, \ll)$ must also be an element of $\min(\beta, \leq_\alpha^*)$, which will suffice. Suppose for contradiction $y_i \notin \min(\beta, \leq_\alpha^*)$ for all i . Then (since \leq_α^* is a tpo) there must be at least one element $z \in [\beta]$ such that $z <_\alpha^* y_i$, equivalently $r_\alpha(z) < r_\alpha(y_i)$, for all i . Since $r_\alpha(y_i) = y_i^-$ for all i , this gives $r_\alpha(z) < y_i^-$ for all i . Clearly it cannot be the case that $z = y_j$ for some j (since then we would have $r_\alpha(y_j) < r_\alpha(y_j)$, which is impossible), hence $z \notin \min(\beta, \ll)$. Hence it must be the case $y_j \ll z$, i.e., $y_j^- \leq z^+$ for some j . But this implies $y_j^- \leq r_\alpha(z)$, contradicting $r_\alpha(z) < r_\alpha(y_i)$, for all i . Hence there must exist some j such that $y_j \in \min(\beta, \leq_\alpha^*)$ as required.

\Leftarrow If β is inconsistent then $\min(\beta, \ll) = \emptyset$ and so the required conclusion $\min(\beta, \ll) \subseteq [\neg\alpha]$ holds true. So suppose β is consistent and $\alpha \notin \leq_\alpha^* \circ \beta$. Then there exists some $y \in \min(\beta, \leq_\alpha^*) \cap [\neg\alpha]$. Suppose for contradiction $\min(\beta, \ll) \not\subseteq [\neg\alpha]$, so there exists $x \in \min(\beta, \ll) \cap [\alpha]$. Using the minimality of y we get $y \leq_\alpha^* x$, i.e., $r_\alpha(y) \leq r_\alpha(x)$. Since $y \in [\neg\alpha]$ and $x \in [\alpha]$ this translates into $y^- \leq x^+$, i.e., $y \ll x$. But this contradicts $x \in \min(\beta, \ll)$. Hence $\min(\beta, \ll) \subseteq [\neg\alpha]$ as required.

(ii) We must show $\min(\beta, \ll\ll) \subseteq [\neg\alpha]$ iff $\neg\alpha \in \leq_\alpha^* \circ \beta$.

\Rightarrow As mentioned above, just after the proof of Proposition 7, we have $\ll\ll \subseteq <_\alpha^*$. This implies $\min(\beta, \leq_\alpha^*) \subseteq \min(\beta, \ll\ll)$. Hence if $\min(\beta, \ll\ll) \subseteq [\neg\alpha]$ then also $\min(\beta, \leq_\alpha^*) \subseteq [\neg\alpha]$, i.e., $\neg\alpha \in \leq_\alpha^* \circ \beta$ as required.

\Leftarrow Suppose $\neg\alpha \in \leq_\alpha^* \circ \beta$, i.e., $\min(\beta, \leq_\alpha^*) \subseteq [\neg\alpha]$, and suppose for contradiction $\min(\beta, \ll\ll) \not\subseteq [\neg\alpha]$. Then there exists $x \in \min(\beta, \ll\ll) \cap [\alpha]$. Since $x \in [\alpha]$ and $\min(\beta, \leq_\alpha^*) \subseteq [\neg\alpha]$ this implies that $x \notin \min(\beta, \leq_\alpha^*)$, so there exists $y \in \min(\beta, \leq_\alpha^*)$ such that $y <_\alpha^* x$, i.e., $r_\alpha(y) < r_\alpha(x)$. Since $y \in \min(\beta, \leq_\alpha^*)$ and $\min(\beta, \leq_\alpha^*) \subseteq [\neg\alpha]$ we know $y \in [\neg\alpha]$ and so $r_\alpha(y) = y^-$. Meanwhile since $x \in [\alpha]$ we know $r_\alpha(x) = x^+$. Hence $r_\alpha(y) < r_\alpha(x)$ translates into $y^- < x^+$, i.e., $y \ll\ll x$, which contradicts $x \in \min(\beta, \ll\ll)$. Hence $\min(\beta, \ll\ll) \subseteq [\neg\alpha]$ as required. \square

For each of the two overrules relations we may consider an interdefinable inference relation. We define:

$$\beta \Rightarrow \alpha \text{ iff } \beta \text{ overrules } \neg\alpha$$

$$\beta \Rightarrow\Rightarrow \alpha \text{ iff } \beta \text{ strictly overrules } \neg\alpha.$$

Using fundamental results by Freund [16] and Kraus et al. [28], classifying various families of nonmonotonic inference relations, Proposition 8 together with the properties of \ll and $\ll\ll$ now allows us to deduce many properties of \Rightarrow and $\Rightarrow\Rightarrow$, and thereby of the overrules relations:

Corollary 2 *The binary relations \Rightarrow and $\Rightarrow\Rightarrow$ are both (consistency-preserving) preferential inference relations, in the sense of Kraus et al. [28]. Furthermore*

they both satisfy the rule of Disjunctive Rationality, i.e., if $\beta \vee \gamma \Rightarrow \alpha$ then either $\beta \Rightarrow \alpha$ or $\gamma \Rightarrow \alpha$.

The first part is a consequence of the fact that \ll and \lll are strict partial orders [28]. In particular it implies that \Rightarrow and \Rightarrow both satisfy the following rules (among others):

$$\frac{\beta \Rightarrow \alpha, \alpha \vdash \gamma}{\beta \Rightarrow \gamma} \text{ (Right Weakening)}$$

$$\frac{\beta \Rightarrow \alpha, \beta \Rightarrow \gamma}{\beta \Rightarrow \alpha \wedge \gamma} \text{ (And)}$$


$$\frac{\beta \Rightarrow \alpha, \beta \Rightarrow \gamma}{\beta \wedge \gamma \Rightarrow \alpha} \text{ (Cautious Monotony)}$$

Switching things around in terms of the corresponding overrules relations, Right Weakening implies if β (strictly) overrules α then β (strictly) overrules every sentence logically *stronger* than α . The And-rule tells us that if β (strictly) overrules both α and γ separately, then it (strictly) overrules their *disjunction*. While Cautious Monotony translates into the rule that if β (strictly) overrules α , then so does $\beta \wedge \neg\gamma$, *provided* β (strictly) overrules γ .

The second part of Corollary 2 follows from results by Freund [16] and Proposition 7(iii). It implies a disjunction $\beta \vee \gamma$ cannot (strictly) overrule α without at least one of its disjuncts doing so. However it's possible for neither \Rightarrow nor \Rightarrow to satisfy the well-known rule Rational Monotony [28] (and thus also Monotony). I.e., if $\beta \Rightarrow \alpha$ and $\beta \not\Rightarrow \neg\gamma$ then $\beta \wedge \gamma \Rightarrow \alpha$. This is because it can be shown that the relations \ll and \lll are not in general *modular*, i.e., they do not verify the property $x <' y$ implies either $x <' z$ or $z <' y$. In fact the following condition, which is easily seen to be weaker than the both the modularity of \ll and of \lll , fails to hold in general:

$$\text{If } x \lll y \text{ then either } x \ll z \text{ or } z \ll y.$$

For the counterexample, consider the following picture:




Then clearly we have $x \lll y$, but neither $x \ll z$ nor $z \ll y$.

7 Limiting Cases

In this section we investigate some special limiting cases of our family of revision operators. Firstly, suppose we insist on the following strengthening of property (≤ 4):

$$(\leq L) \quad x^+ < y^-.$$

In other words, given a choice between a positive representation of any world and a negative representation of any world, we choose the world with a positive representation every time. This is equivalent to the limiting case where $\ll = \emptyset$ (thus also $\lll = \emptyset$). Hence this condition can be thought of as expressing minimal confidence behind the initial tpo \leq . Note that adding this rule to (≤ 2) and (≤ 3) is enough to specify a unique tpo over W^\pm , thus causing (≤ 1) to become redundant. Indeed we are left with the tpo defined by, for all $x, y \in W$ and $\delta, \epsilon \in \{+, -\}$, $x^\delta \leq y^\epsilon$ iff either $(\delta = + \text{ and } \epsilon = -)$ or $(\delta = \epsilon \text{ and } x \leq y)$. In terms of the graphical representation of \leq , this corresponds to the case where every right-point of a stick appears strictly to the right of the left end-points of every stick.



The revision operator $*_L$ defined by this \leq then reduces to:

$$x \leq_\alpha^* y \text{ iff either } x <^\alpha y \text{ or } (x \sim^\alpha y \text{ and } x \leq y)$$

This is the well-known *lexicographic* revision operator studied and axiomatised in the context of iterated belief revision [19, 30, 36]. It amounts to \leq^α being refined by \leq . We can characterise $*_L$ within our family in the following way:

Proposition 9 *If $*$ is generated from some \leq -faithful tpo over W^\pm satisfying $(\leq L)$ then $*$ satisfies:*

$$(*L) \text{ If } x \in [\alpha] \text{ and } y \in [-\alpha] \text{ then } x <_\alpha^* y.$$

Furthermore if $$ is any revision operator for \leq which satisfies $(*L)$ then the \leq -faithful tpo \leq_* defined in the completeness proof of Theorem 1 satisfies $(\leq L)$.*

Proof Suppose \leq satisfies $(\leq L)$ and let $* = *_\leq$. Let $x \in [\alpha]$ and $y \in [-\alpha]$. Then $r_\alpha(x) = x^+$ and $r_\alpha(y) = y^-$. By $(\leq L)$ $r_\alpha(x) < r_\alpha(y)$, i.e., $x <_\alpha^* y$ as required to show $(*L)$.

Conversely suppose $*$ is a revision operator for \leq satisfying $(*L)$ and let \leq_* be as defined in the completeness proof of Theorem 1. We must show $x^+ <_* y^-$ for all x, y . If $x = y$ then $x^+ <_* x^-$ directly by construction. So suppose $x \neq y$. We need to show $x^+ \leq_* y^-$ and $y^- \not\leq_* x^+$. By construction these are equivalent to $x \leq_x^* y$ and $y \not\leq_x^* x$ respectively, i.e., $x <_x^* y$. But by $(*L)$ $x <_x^* z$ for all $z \neq x$. Hence $x <_x^* y$ as required. \square

From this result we see that $*_L$ is axiomatically characterised by $(*1)$ – $(*7)$ plus $(*L)$. However it is easy to see that $(*L)$ implies $(*5)$ – $(*7)$. $(*1)$ also becomes redundant, since $(*3)$, $(*4)$ and $(*L)$ are enough to force the

unique tpo \leq_α^* , and we already established after Proposition 3 that (*2) can be removed. Hence (*3), (*4) and (*L) form a sound and complete axiomatisation for $*_L$. The sentential counterpart of (*L) is the rule *Recalcitrance* of Nayak et al. [30], i.e.,


$$(\circ L) \quad \text{If } \beta \not\vdash \neg\alpha \text{ then } \alpha \in \leq_\alpha^* \circ \beta.$$

Note also that new evidence is always believed after lexicographic revision. A characterisation of $*_L$ in terms of social choice-like conditions was given by Glaister [19], who referred to it as *J-revision*.

At the other extreme, suppose instead we insist on

$$(\leq P) \quad x < y \text{ implies } x^- < y^+.$$

This rule is equivalent to saying $\lll = <$. (Thus also $\ll = <$.) This property expresses maximal confidence behind the initial tpo \leq , or skepticism towards new evidence. Adding this rule to (≤ 2) – (≤ 4) is again enough to specify \leq completely. It is not difficult to show this time we are left with $x^\delta \leq y^\epsilon$ iff either $x < y$ or $[x \sim y \text{ and } (\delta = + \text{ or } \epsilon = -)]$:



The associated revision operator $*_P$ is then given by

$$x \leq_\alpha^* y \text{ iff either } x < y \text{ or } (x \sim y \text{ and } x \leq^\alpha y).$$

This is a so-called reverse lexicographic method, studied in the context of iterated belief revision [33]. This time it corresponds to \leq being refined by \leq^α . In this case new evidence is not always believed.

Proposition 10 *If $*$ is generated from some \leq -faithful tpo over W^\pm satisfying $(\leq P)$ then $*$ satisfies*

$$(*P) \quad \text{If } x \in [-\alpha], y \in [\alpha] \text{ and } x < y \text{ then } x <_\alpha^* y.$$

Furthermore if $$ is any revision operator for \leq which satisfies $(*P)$ then the \leq -faithful tpo \leq_* defined in the completeness proof of Theorem 1 satisfies $(\leq P)$.*

Proof Suppose $* = *_\leq$ for some \leq satisfying $(\leq P)$. To show $*$ satisfies $(*P)$ suppose $x \in [-\alpha], y \in [\alpha]$ and $x < y$. Then $r_\alpha(x) = x^-$ and $r_\alpha(y) = y^+$. Since $x < y$ we may apply $(\leq P)$ to deduce $r_\alpha(x) < r_\alpha(y)$, i.e., $x <_\alpha^* y$ as required.

For the second part let $*$ be a revision operator which satisfies $(*P)$ and suppose $x < y$. We want to show $x^- <_* y^+$, i.e., both $x^- \leq_* y^+$ and $y^+ \not\leq_* x^-$. If $x < y$ then clearly $x \neq y$, hence by construction this is equivalent to showing $x \leq_y^* y$ and $y \not\leq_y^* x$, i.e., $x <_y^* y$. But from $(*P)$ we know $z <_y^* y$ for all $z \neq y$ such that $z < y$. Hence $x <_y^* y$ as required. \square

This result implies that $*_P$ may be characterised axiomatically by (*1)–(*7) plus (*P). However we may significantly simplify this list by observing the following:

Proposition 11 *Let $*$ be any revision operator for \leq satisfying (*3), (*4) and (*5). Then $*$ together satisfies (*6), (*7) and (*P) iff $*$ satisfies:*

$$(*p) \quad < \subseteq <^*_\alpha.$$

Proof

(*6), (*7), (*P) \Rightarrow (*p) In fact we show that, in the presence of the other rules, (*P) is enough to prove (*p) on its own. Suppose $x < y$. To show (*p) we must show $x <^*_\alpha y$. We look at each of the cases $y <^\alpha x$, $x <^\alpha y$ and $x \sim^\alpha y$. If $y <^\alpha x$ then the required conclusion follows immediately from (*P). If $x <^\alpha y$ then the conclusion follows from (*5). Finally if $x \sim^\alpha y$ then the conclusion follows from (*3) or (*4).

(*p) \Rightarrow (*6), (*7), (*P) (*p) \Rightarrow (*P) is immediate. To show (*p) implies the other two rules we show in fact (*p) implies the following property, which is easily seen to be stronger than both (*6) and (*7):

$$\text{If } x <^\alpha y \text{ and } y \leq^*_\alpha x \text{ then } y <^*_\alpha x.$$

This property holds since if $x <^\alpha y$ and $y \leq^*_\alpha x$ then $y < x$ by (*5). Hence $y <^*_\alpha x$ follows by (*p). □

Again (*1) becomes redundant, and so we arrive at the following characterisation of $*_P$.

Proposition 12 *$*_P$ is the unique revision operator for \leq which satisfies (*3)–(*5) plus (*p).*

It is easy to see that the sentential counterpart of (*p) is the following rule:

$$(\circ p) \quad \leq \circ \beta \subseteq \leq^*_\alpha \circ \beta.$$

($\circ p$) states that *all* conditional beliefs in \leq are preserved after revision.

As the following example shows (partly based on one by Darwiche and Pearl [12]), rigid use of either of these limiting cases $*_L$ and $*_P$ can lead to counter-intuitive results.

Example 2 Suppose we have a murder trial with two main suspects, John and Mary. Let p represent “John is the murderer” and q represent “Mary is the

murderer”. Furthermore let r represent “The victim is an alien from outer-space”.²

Initially we believe the murder was committed by one person, either John or Mary. However we *wouldn't be surprised* to discover that either both or neither were involved in the crime. What *would* be surprising—indeed highly shocking—would be if we found out the victim was an alien. However we are still capable of imagining a hypothetical situation in which this turns out to be the case, and we think this would not alter our belief that either John or Mary acted alone. If we were to represent all this using a tpo \leq , it seems the following is the best candidate:

100	○			
010	○			
110		○		
000		○		
101			○	
011			○	
111				○
001				○

Now during the trial we receive testimony that John is the murderer, leading us to revise \leq by p . Supposing we then receive testimony that Mary is the murderer, the most reasonable conclusion would be that both John and Mary were involved in the murder. But using the operator $*_p$ gives

$$\leq_p^{*p} \circ q = Cn(\neg p \wedge q \wedge \neg r)$$

We are forced to drop our belief that John is the murderer.

Now consider the situation where we receive testimony that John is the murderer, followed by the supposition that if John is the murderer, then the victim is an alien. In this case it seems the reasonable thing to do is drop the acquired belief that John is the murderer. However, using the operator $*_L$ gives

$$\leq_p^{*L} \circ (p \rightarrow r) = Cn(p \wedge \neg q \wedge r)$$

That is, we end up believing John murdered an alien!

The move to our more general family of tpo-revision operators enables a correct treatment of both these scenarios simultaneously. Consider the \leq -faithful tpo \leq represented by Fig. 3.

In the first case where we receive evidence pointing towards John's guilt followed by the supposition Mary did it, we have

$$\leq_p^* \circ q = Cn(p \wedge q \wedge \neg r)$$

²In the diagrams below we represent propositional worlds as triples of 0s and 1s corresponding to the propositional variables p , q and r respectively, with 1 denoting truth and 0 denoting falsity. For example, 101 represents the world in which p is true, q is false, and r is true.




Fig. 3 Example of abstract interval ordering

which is the intuitive result. In the case where we receive evidence for John being the murderer, followed by supposing that if John is the murderer then the victim is an alien, we have

$$\leq_p^* \circ (p \rightarrow r) = Cn(\neg p \wedge q \wedge \neg r)$$

which is what we would expect .

8 Another Subclass

Close inspection reveals that both the limiting cases mentioned above share something in common—in both cases we have $\lll = \llr$. Writing out this condition in full, the unique \leq defined in each case satisfies:

$$(\leq 5) \quad x^- \leq y^+ \text{ iff } x^- < y^+.$$

This condition states that no x^- appears in the same \leq -rank as a y^+ . In this section we take a look at the subclass of our family of revision operators defined by enforcing this condition.

One thing to notice is that if $\lll = \llr$ then the distinction between the overrules relation and the strictly overrules relation relative to \leq^* disappears—they collapse into the same binary relation. As for an axiomatic characterisation of this subfamily, the next result points the way:

Proposition 13 *If \ast is generated from some \leq -faithful tpo over W^\pm satisfying (≤ 5) then \ast satisfies*

$$(\ast 8) \quad \text{For } x \in [\alpha] \text{ and } y \in [-\alpha], \text{ either } x <_\alpha^* y \text{ or } y <_\alpha^* x.$$

Furthermore if \ast is any revision operator for \leq which satisfies $(\ast 8)$ then the \leq -faithful tpo \leq_\ast defined in the completeness proof of Theorem 1 satisfies (≤ 5) .

Proof For the first part let $* = *_{\leq}$ for some \leq -faithful tpo satisfying (≤ 5). Let $x \in [\alpha]$ and $y \in [-\alpha]$. Then to show the consequent of ($*8$) we need to show that either $x^+ < y^-$ or $y^- < x^+$. By (≤ 5) we can replace the second disjunct here by $y^- \leq x^+$. But since \leq is a tpo (by (≤ 1)) we always have either $x^+ < y^-$ or $y^- \leq x^+$. Hence the consequent of ($*8$) holds.

For the second part let $*$ be a revision operator satisfying ($*8$). We want to show $x^- \leq_* y^+$ iff $x^- <_* y^+$. If $x = y$ we know $x^+ <_* x^-$ so neither of these conditions can hold, making the biconditional true in this case. So suppose $x \neq y$. In this case the first condition is equivalent to $x \leq_y^* y$ while the second is equivalent to $x <_y^* y$. But from ($*8$) (since $x \neq y$) we know either $x <_y^* y$ or $y <_y^* x$, i.e., $x \not\prec_y^* y$. This means $x \leq_y^* y$ can hold iff $x <_y^* y$, as required. \square

Condition ($*8$) means that after revising by α , there is a separation between α -worlds and $-\alpha$ -worlds, in the sense that each \leq_{α}^* -rank contains *either* only α -worlds or only $-\alpha$ -worlds. This property is called (UR) by Booth and Meyer [7], where it is shown that its sentential counterpart is:

$$(o8) \quad \text{If } \neg\alpha \notin \leq_{\alpha}^* \circ\beta \text{ then } \alpha \in \leq_{\alpha}^* \circ\beta.$$

Postulate (o8) says that after receiving α as evidence and then making the supposition β , α should be believed as long as it is consistent to do so.

($*8$), alias (o8), is quite a strong rule, and adding it to the list ($*1$)–($*7$) causes some redundancies. Since ($*8$) implies the equivalence of $x \leq_{\alpha}^* y$ with $x <_{\alpha}^* y$ for $x \not\prec_{\alpha}^* y$, we see ($*6$) now follows from ($*7$). Meanwhile ($*5$) becomes equivalent to “if $x <^{\alpha} y$ and $x \leq y$ then $x \leq_{\alpha}^* y$ ” (i.e., (CR4) proposed by Darwiche and Pearl [12]. But using the fact that $\leq = \leq_{\top}^*$ (which follows from ($*3$)), this is seen as just the instance of ($*7$) in which $\gamma = \top$. Hence ($*5$) also disappears. Thus the class of tpo-revision operators generated by those \leq -faithful tpos over W^{\pm} satisfying (≤ 5) may be characterised as follows:

Theorem 2 *Let $*$ be a revision operator for \leq . Then $*$ is generated from some \leq -faithful tpo over W^{\pm} satisfying (≤ 5) iff $*$ satisfies ($*1$), ($*3$), ($*4$), ($*7$) and ($*8$).*

Of course we can if we wish replace the last four rules above with their sentential equivalents.

9 Improvement Operators

The problem of defining tpo-revision operators has also been studied recently by Konieczny et al. [25, 27]. Their purpose is to study *iterable* tpo-revision operators in which repeated revision by α eventually leads to acceptance of α into the tpo’s associated belief set. This “weak success” property rules out, for example, just blindly using $*_{\text{P}}$ to always revise the current tpo at every turn. (We remark that this property is fully formalised by Konieczny et al. We just provide an intuitive description here.)

Konieczny et al. study and axiomatise a series of classes of such operators. The general class of *improvement operators* satisfies, in addition to the above weak success requirement, the rules (*1)–(*5). Next comes the class of *soft improvement operators*, which is obtained by adding the following postulate (referred to as (S4) in [25]):

(*soft) If $x \in [\alpha]$ and $y \in [\neg\alpha]$ and $y < x$ then $y \leq^*_\alpha x$

This rule limits the mobility of α -worlds when revising by α . It says an α -world is not allowed to *overtake* a $\neg\alpha$ -world which was initially considered strictly more preferred. This obviously excludes $*_L$. However, the following result shows that soft improvement operators do satisfy another of our postulates. (Recall that (Pareto) is a consequence of (*1)–(*5)—see the discussion just before Proposition 2.)

Proposition 14 *Let $*$ be any tpo-revision operator satisfying (Pareto) and (*soft). Then $*$ satisfies (*6).*

Proof Let $x \in [\alpha]$, $y \in [\neg\alpha]$, i.e., $x <^\alpha y$, and suppose $y \leq^*_\alpha x$. We must show $y \leq^*_\gamma x$ for any γ . From $x <^\alpha y$, $y \leq^*_\alpha x$ and (Pareto) we know $y < x$. If $y \leq^\gamma x$ then $y <^*_\gamma x$ from (Pareto) again and so we obtain the conclusion. If $x <^\gamma y$ then we obtain the required $y \leq^*_\gamma x$ by (*soft). \square

We will see below that soft improvement operators do *not* generally satisfy (*7).

Konieczny et al. go on to describe three distinguished members of the family of soft improvement operators, which we describe informally below (we refer the reader to [25, 27] for the formal details). In the following we use $x \triangleleft y$ to denote the fact that $x < y$ and there is no $z \in W$ such that $x < z < y$:

One-improvement The one-improvement operator, which we denote here by $*_o$ satisfies the following property (called (S5) in [27]):

(*o) If $x \in [\alpha]$ and $y \in [\neg\alpha]$ and $y \triangleleft x$ then $x \leq^*_\alpha y$

Combined with (*soft), this means that if $y <^\alpha x$ and $y < x$, but y and x were neighbours in \leq , then, after revision by α , x ‘moves in’ with y , in the sense that they now share the same rank in \leq^*_α . To illustrate, look at the example in Fig. 4. The leftmost box depicts an initial tpo \leq over $W = \{x_1, x_2, x_3, x_4\}$. The box to the right shows the result of using $*_o$ to revise \leq by some sentence α whose models are $\{x_2, x_4\}$. After revision, x_2 joins the rank of the the immediately \leq -preceding $\neg\alpha$ -world x_1 , and x_4 joins the rank of the the immediately \leq -preceding $\neg\alpha$ -world x_3 .

Half-improvement The half-improvement operator $*_h$ is just like one-improvement, except that x ‘moves in’ with y only if there were no $\neg\alpha$ -worlds sharing the same rank as x in the initial tpo \leq . Thus, in Fig. 4 (third box from




Fig. 4 An initial tpo \leq and the result of revising \leq by α using $*_o$, $*_h$ and $*_b$ resp., where $[\alpha] = \{x_2, x_4\}$

the left) we see that x_2 increases its plausibility with respect to x_3 (in keeping with $(*5)$), but remains strictly less preferred than x_1 , due to the presence of the $\neg\alpha$ -world x_3 in its \leq -rank. For x_4 however, since there is no $\neg\alpha$ -world in the same \leq -rank, x_4 moves into the same rank as x_3 , as in the case of one-improvement.

Best-improvement The best-improvement operator $*_b$ behaves as half-improvement, except that now x moves in with y only if there are no ranks *at all* in \leq which contain both an α -world and a $\neg\alpha$ -world. In such a case α is said to be *separated* in \leq . Effectively best-improvement behaves like $*_p$, unless α is separated in \leq , in which case it behaves like one-improvement. For example, in Fig. 4 α is not separated in \leq (since x_2 and x_3 share the same rank), leading to the result of revision shown in the last box.


Of the three specific soft improvement operators mentioned above, only one-improvement falls within the general family of tpo-revision which we described in the previous sections:

Proposition 15 $*_o$ satisfies $(*7)$, but $*_h$ and $*_b$ generally do not.

Proof To show $*_o$ satisfies $(*7)$ we will show how $(*7)$ may be derived from $(*1)$ – $(*5)$, $(*soft)$ and $(*o)$. Suppose $x <^\alpha y$ and $y <^*_\alpha x$. We must show $y <^*_\gamma x$. Firstly, from $x <^\alpha y$ and $y <^*_\alpha x$ we know $y < x$ from (Pareto). If it were the case that $y \leq^\gamma x$ then we would get the desired conclusion from (Pareto), so suppose $x <^\gamma y$. Now, if it were the case that $y < x$ then we would obtain $x \leq^*_\alpha y$ from this and $x <^\alpha y$ using $(*o)$, thus yielding a contradiction. Hence we have shown that $y < x$, but that it is *not* the case that $y < x$, and so there is some z such that $y < z < x$. We split into two cases, according to whether $z \in [\gamma]$ or not. If $z \in [\gamma]$ then $y \leq^*_\gamma z$ by $(*soft)$ and $z <^*_\gamma x$ by $(*3)$, giving the required $y <^*_\gamma x$ by $(*1)$. If $z \in [\neg\gamma]$ then $y <^*_\gamma z$ by $(*4)$ and $z \leq^*_\gamma x$ by $(*soft)$, again giving the required $y <^*_\gamma x$ by $(*1)$.

To see that $*_h$ and $*_b$ fail to satisfy $(*7)$ in general, consider again the example of Fig. 4. We have $x_1 \in [\neg\alpha]$, $x_2 \in [\alpha]$ and both $x_1 <^{*h}_\alpha x_2$ and $x_1 <^{*b}_\alpha x_2$. If $*_h$ and $*_b$ satisfied $(*7)$ then we would also expect $x_1 <^{*h}_\gamma x_2$ and $x_1 <^{*b}_\gamma x_2$,

where $[\gamma] = \{x_2, x_3, x_4\}$. However both $*_h$ and $*_b$ yield the following tpo when revising by γ .




So both $x_1 <_{\gamma}^{*h} x_2$ and $x_1 <_{\gamma}^{*b} x_2$, contradicting $(*7)$. □

Since one-improvement satisfies $(*1)$ – $(*7)$ we know it may be generated from some \leq -faithful tpo over W^\pm . In fact it is generated by the unique tpo \preceq satisfying $(\preceq 1)$ – $(\preceq 4)$ together with the property

$$(\preceq 0) \quad x < y \text{ implies } x^- \sim y^+$$

In other words, if $x < y$ and x, y are in adjacent ranks in \leq , then x^- and y^+ are in the same rank in \preceq :



10 Strict Preference Hierarchies (or Interval Orderings)

In this section we introduce a way of re-packaging a given ordering \leq over W^\pm satisfying $(\preceq 1)$ – $(\preceq 4)$. We show that this alternative representation is equivalent to using the class of orderings \preceq . This representation in terms of *strict preference hierarchies* will be used in Section 11 to describe desirable properties for the revision of, not just tpos, but *the strict preference hierarchies themselves*. This is equivalent to the revision of the class of orderings \preceq , and therefore goes beyond the revision of just tpos, to provide the first steps in the description of an approach for revising epistemic states containing the enriched structure to be found in the class of orderings \preceq .

As observed in Section 6, from a single \preceq we can extract *three* different notions of *strict preference* over W .

1. $x < y$ iff $x^+ < y^+$
2. $x \ll y$ iff $x^- \preceq y^+$
3. $x \lll y$ iff $x^- < y^+$

We are now in a position to define our alternative representation of \preceq .

Definition 6 The triple $\$ = (\lll, \ll, <)$ of binary relations over W is a *strict preference hierarchy (over W)* (SPH for short) iff there is some relation \preceq over W^\pm satisfying $(\preceq 1)$ – $(\preceq 4)$ such that \lll, \ll and $<$ can all be defined from \preceq as above. We shall sometimes say that $\$$ is *relative to* $<$.

Such “interval orderings” like the above have already been studied in the context of temporal reasoning [2], as well as in preference modelling [32]. Indeed, concerning the former case, the relations $\lll, \ll, <$ could all be defined in terms of the relations *before*, *meets* and *overlaps* between temporal intervals studied by Allen [2].

What are the properties of the three relations $(\lll, \ll, <)$? A couple were already mentioned in Section 6. For example we already know from there that \lll and \ll are strict partial orders (i.e., irreflexive and transitive). But what else do they satisfy? In particular how do they *interrelate* with each other? Furthermore, given any *arbitrary* triple $\$ = (\lll, \ll, <)$ of binary relations over W , under what conditions on $\$$ can we be sure that $\$$ forms an SPH, i.e., under what conditions can we be sure there is *some* \preceq satisfying $(\preceq 1)$ – $(\preceq 4)$ such that $\$$ can be derived from \preceq in the above manner. These questions are answered by the following representation result for SPHs. We point out that part (iii) of the “only if” part (but not the “if” part) was essentially already proved, in the temporal reasoning context, by Allen [2].

Theorem 3 Let \lll, \ll and $<$ be three binary relations over W . Then $\$ = (\lll, \ll, <)$ is an SPH iff the following conditions hold (where $x \preceq y$ iff $y \neq x$):

- (i) \preceq is a total preorder.
- (ii) $\lll \subseteq \ll \subseteq <$.
- (iii) The following are satisfied, for all $x, y, z \in W$:
 - (SPH1) $z \preceq x$ and $x \lll y$ implies $z \lll y$
 - (SPH2) $x \lll y$ and $y \preceq z$ implies $x \lll z$
 - (SPH3) $z \preceq x$ and $x \ll y$ implies $z \ll y$
 - (SPH4) $x \ll y$ and $y \preceq z$ implies $x \ll z$
 - (SPH5) $z < x$ and $x \ll y$ implies $z \lll y$
 - (SPH6) $x \ll y$ and $y < z$ implies $x \lll z$

Proof The “only if” direction is straightforward, and in fact easy to visualise given our graphical representation of \preceq . For the “if” direction suppose $\$ = (\lll, \ll, <)$ satisfies (i)–(iii). We must find some relation \preceq_S over W^\pm such that (a) \preceq_S satisfies $(\preceq 1)$ – $(\preceq 4)$, and (b) the relations \lll, \ll and $<$ may be defined from \preceq_S as above. We define \preceq_S as follows. For each $x, y \in W$ and $\epsilon, \delta \in \{+, -\}$ we must specify the conditions under which $x^\epsilon \preceq_S y^\delta$ or not. First, in the case $\epsilon = \delta$ we define

$$x^\epsilon \preceq_S y^\epsilon \text{ iff } x \leq y.$$

This clearly ensures \leq_S satisfies conditions (≤ 2) and (≤ 3) . If $\epsilon \neq \delta$ but $x = y$ then we declare

$$x^+ \prec_S x^-.$$

This ensures (≤ 4) is satisfied. Finally if $\epsilon \neq \delta$ and $x \neq y$ then we set

$$x^+ \leq_S y^- \text{ iff } y \lll x \quad x^- \leq_S y^+ \text{ iff } x \lll y.$$

We still need to show \leq_S satisfies (≤ 1) , i.e., \leq_S is connected and transitive.

Connectedness We need to show for any $x, y \in W$ and $\epsilon, \delta \in \{+, -\}$ either $x^\epsilon \leq_S y^\delta$ or $y^\delta \leq_S x^\epsilon$. If $\epsilon = \delta$ then, by construction of \leq_S , this reduces to showing that either $x \leq y$ or $y \leq x$, which obviously holds since \leq is connected. So suppose $\epsilon \neq \delta$. In this case if furthermore $x = y$ then we know by construction that precisely *one* of $x^\epsilon \leq_S x^\delta$ and $x^\delta \leq_S x^\epsilon$ holds, namely $x^+ \leq_S x^-$. So suppose both $\epsilon \neq \delta$ and $x \neq y$. Assume $\epsilon = +$ and $\delta = -$ (the reverse case is symmetrical). Then by construction we need to show either $y \lll x$ or $y \ll x$. But this follows since $\lll \subseteq \ll$ as required.

Transitivity We need to show, for any $x, y, z \in W$ and $\epsilon, \delta, \nu \in \{+, -\}$,

$$x^\epsilon \leq_S y^\delta \text{ and } y^\delta \leq_S z^\nu \text{ implies } x^\epsilon \leq_S z^\nu.$$

If $\nu = \delta = \epsilon$ then this follows from the transitivity of \leq . Now we consider the other possible combinations of ϵ, δ and ν .

$\epsilon = \delta \neq \nu$ In this case, since $x^\epsilon \leq_S y^\delta$ reduces to $x \leq y$, we must show

$$x \leq y \text{ and } y^\delta \leq_S z^\nu \text{ implies } x^\delta \leq_S z^\nu.$$

So suppose the antecedent holds. Suppose $\delta = +$ and $\nu = -$. If $z = x$ then the consequent holds by construction of \leq_S so suppose $z \neq x$. Then the consequent reduces to $z \lll x$. If $y = z$ then from $x \leq y$ we get $x \leq z$ and so, since $\lll \subseteq <$, we obtain the required consequent. So suppose also $y \neq z$. Then from $y^\delta \leq_S z^\nu$ we obtain $z \lll y$. Then this together with $x \leq y$ gives us the required $z \lll x$ using (SPH2).

Suppose instead $\delta = -$ and $\nu = +$. Then $y \neq z$ (since otherwise $y^\delta \leq_S z^\nu$ becomes $y^- \leq_S y^+$, contrary to the construction of \leq_S) so $y^\delta \leq_S z^\nu$ becomes $y \ll z$. If it were the case $z = x$ then this would give $y \ll x$ and so, since $\ll \subseteq <$, $y < x$ – contradiction. Hence $z \neq x$ which means to show the consequent holds we need $x \ll z$. But this follows from $y \ll z$ and $x \leq y$ using (SPH3) as required.

$\nu = \epsilon \neq \delta$ In this case the consequent reduces to $x \leq z$, and so we must show

$$x^\nu \leq_S y^\delta \text{ and } y^\delta \leq_S z^\nu \text{ implies } x \leq z.$$

Suppose the antecedent holds. First suppose $\nu = +$ and $\delta = -$. Then we know $z \neq y$ (since otherwise $y^\delta \leq_S z^\nu$ would be $z^- \leq_S z^+$ which is not possible), hence from $y^\delta \leq_S z^\nu$ we know $y \ll z$. If $y = x$ then the consequent becomes

$y \leq z$, which then follows from $y \ll z$ using the fact that $\ll \subseteq <$. So suppose $y \neq x$. Then $x^\nu \leq_S y^\delta$ becomes $y \lll x$, and this together with $y \ll z$ gives the required $x \leq z$ using (SPH6).

Suppose instead $\nu = -$ and $\delta = +$. Then we know $y \neq x$ (otherwise $x^\nu \leq_S y^\delta$ would lead to $y^- \leq_S y^+$) so $x^\nu \leq_S y^\delta$ reduces to $x \ll y$. If $z = y$ then this in turn gives $x \ll z$ which implies the required $x \leq z$ using the fact that $\ll \subseteq <$. So suppose $z \neq y$. Then from $y^\delta \leq_S z^\nu$ we get $z \lll y$ which, together with $x \ll y$ gives the required $x \leq z$ using (SPH5).

$\nu = \delta \neq \epsilon$ In this case we must show

$$x^\epsilon \leq_S y^\delta \text{ and } y \leq z \text{ implies } x^\epsilon \leq_S z^\delta.$$

First suppose $\delta = +$ and $\epsilon = -$. Then $x \neq y$ since otherwise $x^\epsilon \leq_S y^\delta$ would become $x^- \leq_S x^+$ which is impossible. Hence $x^\epsilon \leq_S y^\delta$ becomes $x \ll y$. From this we know $x \neq z$ since otherwise $z \ll y$, which contradicts $y \leq z$ (since $\ll \subseteq <$). Hence the consequent is $x \ll z$. But this follows from $z \ll y$ and $y \leq z$ using (SPH4).

Now suppose $\delta = -$ and $\epsilon = +$. If $x = z$ then the consequent becomes $x^+ \leq_S x^-$, which holds automatically by construction of \leq_S . So assume $x \neq z$, which means the consequent is equivalent to $z \lll x$. Now if $x = y$ then this would be just $z \lll y$, which then follows from $y \leq z$ using the fact $\lll \subseteq <$. So assume also $x \neq y$. Then $x^\epsilon \leq_S y^\delta$ becomes $y \lll x$. But this together with $y \leq z$ still implies the desired consequent $z \lll x$ using (SPH1).

Thus we have proved that \leq_S satisfies all three rules (≤ 1)–(≤ 4). It remains to show \lll , \ll and $<$ may all be recaptured from \leq_S . For \lll we need to show for any $x, y \in W$, $x \lll y$ iff $x^- <_S y^+$, in other words $x \lll y$ iff both $x^- \leq_S y^+$ and $y^+ \not\leq_S x^-$. If $x = y$ then the left-hand-side will be false (since \lll is clearly irreflexive) so in this case we must show $x^- \not\leq_S x^+$, i.e., $x^+ \leq_S x^-$. But this holds by construction of \leq_S . If $x \neq y$ the construction tells us we must show $x \lll y$ iff both $x \ll y$ and $x \lll y$. But this holds since $\lll \subseteq \ll$. Hence \lll may indeed be defined from \leq_S . For \ll we need to show $x \ll y$ iff $x^- \leq_S y^+$. The case $x = y$ holds as in the case above for \lll , while the case $x \neq y$ follows immediately by construction. Finally for $<$ we need $x < y$ iff $x^+ <_S y^+$. Again this is immediate from the construction. □

The rules (SPH1)–(SPH6) each represent some sort of transitivity condition across the relations of the SPH. Note it follows easily from these conditions that \lll and \ll are strict partial orders.

Two special limiting cases of SPHs were already mentioned in Section 7: Given any tpo \leq over W with strict part $<$, the triples $(\emptyset, \emptyset, <)$ and $(<, <, <)$ each *always* forms an SPH, as can easily be seen by checking conditions (i)–(iii) of the theorem. In fact these are the SPH forms of the well-known lexicographic tpo-revision operator [30] and Papini’s [33] “reverse” lexicographic tpo-revision operator respectively.

SPHs seem quite closely related to the notion of “PQI interval order” studied by Öztürk et al. [32]. Indeed several representation results in the same spirit as Theorem 3 can be found in their work. The main difference with ours is that PQI interval orders make use of an explicit numerical scale, so the endpoints of the intervals are ordinary real numbers, whereas our intervals are “abstract”, having endpoints only in some totally preordered set (but see Section 11.1 of this paper). Also, with PQI interval orders, different possibilities (i.e., possible worlds for us) may be assigned intervals of different length. It is even possible for the interval assigned to one possibility to be completely *enclosed* in the interval assigned to another. This is something we do not allow. We are currently in the process of examining in more detail the relationship between SPHs and PQI interval orders.

To summarise the findings of this section, we now see we have two different, but equivalent ways of describing the structure required to revise a tpo \preceq :

1. As a \preceq -faithful tpo \preceq over W^\pm satisfying (≤ 1) – (≤ 4) .
2. As a triple $(\lll, \ll, <)$ of binary relations over W satisfying conditions (i)–(iii) from Theorem 3 (with $<$ being the strict part of \preceq).

Recall that the revision operator $*$ for \preceq derived from a \preceq -faithful tpo \preceq over W^\pm is defined by setting $x \preceq_\alpha^* y$ iff $r_\alpha(x) \preceq r_\alpha(y)$. The next result shows how we can describe $*$ purely in terms of the SPH corresponding to \preceq .

Proposition 16 *Let \preceq be a tpo over W and let \preceq be a given \preceq -faithful tpo over W^\pm . Let $\mathcal{S} = (\lll, \ll, <)$ be the SPH corresponding to \preceq and let $*$ be the revision operator for \preceq derived from \preceq . Then, for all $x, y \in W$,*

$$x \preceq_\alpha^* y \text{ iff } \begin{cases} x \sim^\alpha y \text{ and } x \leq y \\ \text{or } x <^\alpha y \text{ and } y \lll x \\ \text{or } y <^\alpha x \text{ and } x \ll y. \end{cases}$$

Proof Given $x, y \in W$, we can clearly split into three mutually exhaustive and exclusive cases $x \sim^\alpha y$, $x <^\alpha y$ and $y <^\alpha x$. In the first case we know already $x \preceq_\alpha^* y$ iff $x \leq y$ by $(*3)$ and $(*4)$. This takes care of the first clause in the above identity for \preceq_α^* . If $x <^\alpha y$ then $r_\alpha(x) = x^+$ and $r_\alpha(y) = y^-$ so by definition of \preceq_α^* we have $x \preceq_\alpha^* y$ iff $x^+ \preceq y^-$. This is equivalent to $y^- \not\prec x^+$, i.e., $y \lll x$. This takes care of the middle clause in the above identity. Finally if $y <^\alpha x$ then $r_\alpha(x) = x^-$ and $r_\alpha(y) = y^+$, so now $x \preceq_\alpha^* y$ iff $x^- \preceq y^+$, i.e., $x \ll y$ as required to show the last clause in the above identity. \square

Since the class of orderings \preceq and the class of SPHs are equivalent, any way of revising one of these two types of structure will automatically give us a way of revising the other. We are free to use whichever one seems more appropriate at the time. For the purpose of expressing *desirable properties* of revising \preceq , it is easier to express such properties in terms of SPHs than \preceq .

11 Properties of SPH Revision

Given an SPH $\$$ and a sentence α , we want to determine the new SPH $\$ \circledast \alpha$ which is the result of revising the entire SPH $\$$ by α . Assume $\$ = (\lll, \ll, <)$ and let's denote $\$ \circledast \alpha$ by $(\lll', \ll', <')$. Firstly, we have the following three fundamental properties:


- (\circledast 1) $\$ \circledast \alpha$ is an SPH
- (\circledast 2) $<' = <^*_\alpha$
- (\circledast 3) If $\alpha \equiv \gamma$ then $\$ \circledast \alpha = \$ \circledast \gamma$

In (\circledast 2), $<^*_\alpha$ is the strict version of the tpo \leq^*_α determined using \leq, \ll and \lll as in Proposition 16. In other words, $\$ \circledast \alpha$ should be an SPH relative to $<^*_\alpha$. (\circledast 3) is a syntax-irrelevance property.


With $<'$ settled, it remains to specify \lll' and \ll' . An initial suggestion for the new strong strict preferences \lll' might be to keep it unchanged. That is, to set \lll' equal to \lll . This can be seen as a pure application of minimal change to \lll . In addition, it is easy to see that $\lll \subseteq <'$ and so such a choice is not at odds with part (ii) of Theorem 3. However, the following example shows this can't be done in general. For $\$ \circledast \alpha$ to be an SPH it is necessary to satisfy

(SPH1) $z \leq^*_\alpha x$ and $x \lll' y$ implies $z \lll' y$

But if we set $\lll = \lll'$ this might not hold in general. For suppose we are given a portion of the \leq corresponding to $\$$ as follows:



So $x \lll y$ and $z \not\lll y$. Now suppose we revise by a sentence α such that $z \in [\alpha]$ and $x, y \in [-\alpha]$.



Then $z <^*_\alpha x$, thus giving the required counterexample. Note, incidentally, that it is still a counterexample if we assume $y \in [\alpha]$. Thus there are times when the set of strong strict preferences *must* change. In the above counterexample, when we move from \lll to \lll' we must *either* lose $x \lll y$, *or* gain $z \lll y$. How do we decide which? A useful approach is to distinguish between the case $y \in [-\alpha]$, as indicated in the counterexample above, and the case $y \in [\alpha]$. In the former case intuition dictates that $x \lll y$ ought to be retained since α does not discriminate between x and y : they are both in $[-\alpha]$. Moreover, it is justifiable to gain $z \lll y$ since we have a positive representation of z ($z \in [\alpha]$) and a negative representation of y ($y \in [-\alpha]$). On the other hand, in the case where $y \in [\alpha]$ it can be argued that the strong preference $x \lll y$ can be lost since

we don't have such a strong case to prefer x over y anymore when $x \in [-\alpha]$ and $y \in [\alpha]$. Also, note that in this case it seems reasonable to require that the relative ordering of z and y with respect to $<$, \ll and \lll ought to remain unchanged since α does not distinguish between z and y : they are both in $[\alpha]$. This brings us, in fact, to what can be regarded as the basic postulates for SPH revision, once $(\otimes 1)$ – $(\otimes 3)$ are included as well:

- $(\otimes 4a)$ If $x \sim^\alpha y$ then $x \ll y$ iff $x \ll' y$
- $(\otimes 4b)$ If $x \sim^\alpha y$ then $x \lll y$ iff $x \lll' y$
- $(\otimes 5a)$ If $x <^\alpha y$ then $x \leq y$ implies $x \ll' y$
- $(\otimes 5b)$ If $x <^\alpha y$ then $x < y$ implies $x \lll' y$

Definition 7 The SPH-revision operator \otimes is *admissible* iff it satisfies $(\otimes 1)$ – $(\otimes 3)$, $(\otimes 4a)$, $(\otimes 4b)$, $(\otimes 5a)$ and $(\otimes 5b)$.

We refer to this as admissible SPH revision since it corresponds closely to admissible revision as defined by [7]. $(\otimes 4a)$ and $(\otimes 4b)$ are versions of Darwiche and Pearl's (CR1) and (CR2) [12], or rules $(\ast 3)$ and $(\ast 4)$ defined earlier. They require that the ordering of two elements x and y be unchanged, wrt to \ll and \lll , provided that the circumstances for x and y are the same (i.e. either both are in $[\alpha]$ or both are in $[-\alpha]$). This can be seen as an application of minimal change to \ll and \lll . The postulates $(\otimes 5a)$ and $(\otimes 5b)$ are versions of rule $(\ast 5)$ defined earlier. In fact, in the presence of the fundamental rules $(\otimes 1)$ and $(\otimes 2)$, $(\otimes 5a)$ is a *strengthening* of $(\ast 5)$. They ensure that a “widening of the gap” between x and y occurs when we have a positive representation of x and a negative representation of y . This can be viewed as making sure that the evidence α is taken seriously. A world x in $[\alpha]$ will be more preferred with respect to a world y in $[-\alpha]$, provided that y was not preferred to x to start with. So, informally, admissible SPH revision effects a “slide to the right” of those worlds in $[-\alpha]$ in a manner similar to that described by Booth and Meyer [7]. The difference here is that, with the aid of \ll and \lll , we can specify more precisely how such a slide is allowed to take place.

We now turn to some additional properties which, on the face of it, seem to be desirable, and then investigate how they square up against admissible SPH revision. The first one we consider is

$$(\otimes 6) \quad S \otimes T = S$$

which states that everything remains unchanged if we revise by a tautology. And indeed, $(\otimes 6)$ follows immediately from $(\otimes 2)$, $(\otimes 4a)$ and $(\otimes 4b)$.

Next we consider the pair of properties

- $(\otimes 7a)$ If $x \ll y$ and $x \lll' y$ then $y <^\alpha x$
- $(\otimes 7b)$ If $x \lll y$ and $x \lll' y$ then $y <^\alpha x$

which state that losing a \ll -preference or a \lll -preference of x over y must be the result of y being represented positively ($y \in [\alpha]$) and x being represented

negatively ($x \in [-\alpha]$). It's easy to verify that $(\otimes 7a)$ follows from $(\otimes 4a)$ and $(\otimes 5a)$, while $(\otimes 7b)$ follows from $(\otimes 4b)$ and $(\otimes 5b)$.

Next is the pair of properties

$$(\otimes 8a) \text{ If } x \ll y \text{ and } x \ll' y \text{ then } x <^\alpha y$$

$$(\otimes 8b) \text{ If } x \ll\ll y \text{ and } x \ll\ll' y \text{ then } x <^\alpha y$$

which state that *gaining* a \ll -preference or an $\ll\ll$ -preference of x over y must be the result of x being represented positively ($x \in [\alpha]$) and y being represented negatively ($y \in [-\alpha]$). It turns out that $(\otimes 8a)$ follows from $(\otimes 1)$, $(\otimes 2)$ and $(\otimes 4a)$, while $(\otimes 8b)$ follows from $(\otimes 1)$, $(\otimes 2)$ and $(\otimes 4b)$.

Next we mention a property *not* compatible with admissible SPH revision:

$$(\otimes 9) \text{ If } (\ll\ll, \ll \cap <_\alpha^*, <_\alpha^*) \text{ is an SPH then } \$ \otimes \alpha = (\ll\ll, \ll \cap <_\alpha^*, <_\alpha^*)$$

Property $(\otimes 9)$ is an attempt to enforce the principle of minimal change with respect to both \ll and $\ll\ll$. To see that it is incompatible with admissible revision, suppose $\$$ is of the form $(\emptyset, \emptyset, <)$, i.e., $\ll\ll = \ll = \emptyset$. Assume furthermore that $x < y$ and suppose we then revise by α such that $x <^\alpha y$. Then $(\ll\ll, \ll \cap <_\alpha^*, <_\alpha^*) = (\emptyset, \emptyset, <_\alpha^*)$ is an SPH and so $(\otimes 9)$ dictates that $\$ \otimes \alpha = (\emptyset, \emptyset, <_\alpha^*)$. But observe that admissible SPH revision, and more specifically $(\otimes 5b)$, requires that $x \ll\ll' y$, which contradicts $\ll\ll' = \emptyset$.

The difference between the approach advocated by $(\otimes 9)$ and admissible SPH revision is that $(\otimes 9)$ requires all three orderings to change as little as possible, while with $(\otimes 5a)$ and $(\otimes 5b)$ we are advocating that the new evidence α overrides the principle of minimal change.

Finally we mention a couple of plausible properties which go *beyond* those of admissible revision, in that they relate the results of revising by *different* sentences. Recall (Definition 4) that we say α and γ *agree* on x, y iff they both “say the same thing” regarding the relative plausibility of x, y . The next two rules express that whether or not $x \ll' y$ and $x \ll\ll' y$ should depend only on $\$$ and on what the input sentence says about the relative plausibility between x, y . They express a principle of “Independence of Irrelevant Alternatives in the Input”. Here we are writing $\$ \otimes \alpha = (\ll\ll_\alpha^*, \ll_\alpha^*, <_\alpha^*)$ and $\$ \otimes \gamma = (\ll\ll_\gamma^*, \ll_\gamma^*, <_\gamma^*)$.

$$(\otimes 10a) \text{ If } \alpha \text{ and } \gamma \text{ agree on } x, y \text{ then } x \ll_\alpha^* y \text{ iff } x \ll_\gamma^* y$$

$$(\otimes 10b) \text{ If } \alpha \text{ and } \gamma \text{ agree on } x, y \text{ then } x \ll\ll_\alpha^* y \text{ iff } x \ll\ll_\gamma^* y$$

We omit the case for $<_\alpha^*, <_\gamma^*$, since these were already proved in Proposition 3 to follow from $(*)1$ – $(*)7$. It is thus already handled by $(\otimes 2)$. It can be shown that adding these two rules to those for admissible revision leads to the redundancy of $(\otimes 3)$ and allows $(\otimes 4a)$ and $(\otimes 4b)$ to be replaced by the simple rule $(\otimes 6)$.

11.1 A Concrete Revision Operator

In the previous section we proposed that any reasonable SPH-revision operator should at the very least be admissible according to Definition 7. In

this section we demonstrate that such operators exist by defining a concrete operator for SPH revision which is admissible. This operator employs yet more structure which goes beyond SPHs and their corresponding orderings \preceq over W^\pm , and which takes us a step closer to the PQI interval orders of Öztürk et al. [32] and also to semi-quantitative representations of epistemic states such as that of Spohn [36]. But we expect there will be other, interesting, admissible revision operators which can still be defined in a purely qualitative fashion. This is a topic for further research.

To describe our operator it will be useful to switch back to the \preceq -representation of our tpo-revising structure rather than work directly with SPHs. The basic idea is to enrich the \preceq -representation with numerical information. More precisely, to each element $x^\epsilon \in W^\pm$ we assign a real number $p(x^\epsilon)$ such that for all $x \in W$,

$$p(x^-) - p(x^+) = a > 0,$$

where a is some given fixed real number. The idea is that the smaller the number $p(x^\epsilon)$, the more preferred x^ϵ is. To each such assignment p we may associate an ordering \preceq_p over W^\pm given by

$$x^\epsilon \preceq_p y^\delta \text{ iff } p(x^\epsilon) \leq p(y^\delta).$$

Essentially we replace our abstract intervals (x^+, x^-) with the real intervals $(p(x^+), p(x^-))$, all of length a . It is obvious that \preceq_p satisfies $(\preceq 1)$ – $(\preceq 4)$. (Again, we point out it is not *absolutely* necessary for all the intervals to be of the *same* length a in order for \preceq_p to satisfy $(\preceq 2)$ and $(\preceq 3)$.)

To revise a given SPH $\$$ by sentence α we will use the following procedure:

1. Convert $\$$ to its corresponding tpo \preceq over W^\pm
2. Choose some p such that $\preceq = \preceq_p$
3. Revise p to get a new assignment $p * \alpha$
4. Take $\$ \circledast \alpha$ to be the SPH corresponding to $\preceq_{p*\alpha}$

Clearly the crucial step here is step 3. How should we determine $p * \alpha$? We propose a very simple method here. We define $p * \alpha$ by setting, for each $x^\epsilon \in W^\pm$,

$$(p * \alpha)(x^\epsilon) = \begin{cases} p(x^\epsilon) & \text{if } x \in [\alpha] \\ p(x^\epsilon) + a & \text{if } x \in [-\alpha] \end{cases}$$

In other words, the interval $(p(x^+), p(x^-))$ associated to x remains unchanged if x satisfies α , but is “moved back” by amount a to $(p(x^-), p(x^-) + a)$ if x satisfies $\neg\alpha$. Essentially this boils down to nothing more than an operation familiar from the context of Spohn-type rankings known as *L-conditionalisation* [18].

The following result reveals what $\$ \circledast \alpha$ will look like.

Proposition 17 *Assume $\$ = (\lll, \ll, <)$ and let $\$ \circledast \alpha = (\lll', \ll', <')$ be as defined in the above procedure, for suitable p in step 2. Then, for any $x, y \in W$,*

(i) $<' = <^*_\alpha$, where $*$ is the revision operator corresponding to $\$$ as in Proposition 16.

(ii)

$$x \ll' y \text{ iff } \begin{cases} x \sim^\alpha y \text{ and } x \ll y \\ \text{or } x <^\alpha y \text{ and } x \leq y \\ \text{or } y <^\alpha x \text{ and } p(x^-) + a \leq p(y^+). \end{cases}$$

(iii)

$$x \lll' y \text{ iff } \begin{cases} x \sim^\alpha y \text{ and } x \lll y \\ \text{or } x <^\alpha y \text{ and } x < y \\ \text{or } y <^\alpha x \text{ and } p(x^-) + a < p(y^+). \end{cases}$$

Proof We assume \leq is the tpo over W^\pm corresponding to $\$$ and that p is chosen such that $\leq = \leq_p$.

(i) We need to show $x <' y$ iff $x <^*_\alpha y$. By construction of $<'$, the left hand side here is equivalent to $x^+ <_{p*\alpha} y^+$, i.e.,

$$(p * \alpha)(x^+) < (p * \alpha)(y^+).$$

Meanwhile, using the identity in Prop. 16 we may reformulate the right hand side, as follows:

$$x <^*_\alpha y \text{ iff } \begin{cases} x \sim^\alpha y \text{ and } x < y \\ \text{or } y <^\alpha x \text{ and } x \lll y \\ \text{or } x <^\alpha y \text{ and } y \ll x. \end{cases}$$

We now check for identity in each of the cases $x \sim^\alpha y$, $y <^\alpha x$ and $x <^\alpha y$. First suppose $x \sim^\alpha y$. Then we must show $(p * \alpha)(x^+) < (p * \alpha)(y^+)$ iff $x < y$. But if $x \sim^\alpha y$ then either both worlds satisfy α or both do not. In the former case $(p * \alpha)(x^+) = p(x^+)$ and $(p * \alpha)(y^+) = p(y^+)$, while in the latter case $(p * \alpha)(x^+) = p(x^+) + a$ and $(p * \alpha)(y^+) = p(y^+) + a$. In both cases we are left with $(p * \alpha)(x^+) < (p * \alpha)(y^+)$ iff $p(x^+) < p(y^+)$. But this is the same as $x^+ < y^+$, i.e., $x < y$ as required.

If $y <^\alpha x$ then we must show $(p * \alpha)(x^+) < (p * \alpha)(y^+)$ iff $x \lll y$. But in this case we get $(p * \alpha)(x^+) = p(x^+) + a = p(x^-)$ and $(p * \alpha)(y^+) = p(y^+)$, so $(p * \alpha)(x^+) < (p * \alpha)(y^+)$ iff $p(x^-) < p(y^+)$. But this is the same as $x^- < y^+$, i.e., $x \lll y$ as required.

Finally if $x <^\alpha y$ then we must show $(p * \alpha)(x^+) < (p * \alpha)(y^+)$ iff $y \ll x$. But in this case $(p * \alpha)(x^+) = p(x^+)$ and $(p * \alpha)(y^+) = p(y^+) + a = p(y^-)$, so $(p * \alpha)(x^+) < (p * \alpha)(y^+)$ iff $p(x^+) < p(y^-)$ iff $p(y^-) \leq p(x^+)$. But this is the same as $y^- \leq x^+$, i.e., $y \ll x$ as required.

(ii) We have $x \lll' y$ iff $x^- \leq_{p*\alpha} y^+$ iff $(p * \alpha)(x^-) \leq (p * \alpha)(y^+)$. Again we check for identity in each of the cases $x \sim^\alpha y$, $x <^\alpha y$ and $y <^\alpha x$. First if $x \sim^\alpha y$ then we need to show $(p * \alpha)(x^-) \leq (p * \alpha)(y^+)$ iff $x \lll y$. But as in part (i) above we have $(p * \alpha)(x^-) \leq (p * \alpha)(y^+)$ iff $p(x^-) \leq p(y^+)$ iff $x^- \leq y^+$, i.e., $x \lll y$ as required.

If $x <^\alpha y$ then we need $(p * \alpha)(x^-) \leq (p * \alpha)(y^+)$ iff $x \leq y$. But in this case $(p * \alpha)(x^-) = p(x^-)$ and $(p * \alpha)(y^+) = p(y^+) + a = p(y^-)$, so $(p * \alpha)(x^-) \leq (p * \alpha)(y^+)$ becomes $p(x^-) \leq p(y^-)$, i.e., $x^- \preceq y^-$, which is the same as $x \leq y$ as required.

If $y <^\alpha x$ then we need $(p * \alpha)(x^-) \leq (p * \alpha)(y^+)$ iff $p(x^-) + a \leq p(y^+)$. But this holds since $(p * \alpha)(x^-) = p(x^-) + a$ and $(p * \alpha)(y^+) = p(y^+)$ directly by definition of $p * \alpha$.

- (iii) Proved along exactly similar lines to (ii), but with strict inequalities $<$ replacing the weak ones \leq . \square

From this result we can see that \otimes satisfies $(\otimes 2)$, $(\otimes 4a)$, $(\otimes 4b)$, $(\otimes 5a)$ and $(\otimes 5b)$. We can also see from this that the result of revision depends on $[\alpha]$ rather than α , thus $(\otimes 3)$ is also satisfied. Meanwhile rule $(\otimes 1)$ obviously holds. Thus:

Corollary 3 *The SPH-revision operator \otimes defined via the above procedure from a given assignment p is admissible. Furthermore $(\otimes 10a)$ and $(\otimes 10b)$ also hold.*

12 Conclusion

We have introduced a new family of operators for revising tpos by sentences based on the simple intuitive idea that when we compare possibilities, we are often able to imagine these possibilities with regard to best case and worst case scenarios. We then extended this framework to revise not only tpos, but also the *structure* required to guide the revision of the tpo. We showed that this structure may be described in terms of strict preference hierarchies (SPHs), and proved the equivalence of this representation with the class of orderings \preceq . We provided some properties which any reasonable SPH-revision operator ought to satisfy, and proved their consistency by giving a concrete example of an SPH-revision operator which satisfies them. We placed our work firmly in the context of the problem of iterated belief revision, and showed that our results significantly extend current work on this topic.

In this paper we have proposed a type of structure which can be placed on top of the usual tpo representation \leq of an agent's epistemic state and whose role is in fact to calculate the agent's new tpo in its revised epistemic state. This structure takes a specific form in the guise of a \preceq -faithful tpo over the set W^\pm or, equivalently, a SPH relative to $<$. The question naturally arises as to whether any other kinds of structure are conceivable which equally go beyond the basic tpo representation and which can be deployed to obtain a new tpo following receipt of new propositional information. In fact several different representations can be found in the literature. Bochman [5] represents an epistemic state as an ordered set of states, each state being labelled with a logically closed set of sentences. Lehmann [29] uses a sequence of sentences representing the revision history of the agent (see also [26]). More generally one can use a sequence of *sets* of sentences [13]. Revision is enacted by placing

the new sentence in an appropriate position in the sequence (e.g., at the right-most end), though the precise position may be determined by extra means which may depend on the context of the revision episode. Another enrichment was considered by Booth et al. in [9], with the $\text{tpo} \leq$ over W being augmented by a second ordering, also over W . However, there the extra structure was deployed to calculate the result of a *single-step belief contraction* rather than iterated *revision*. Finally there are the more quantitative (or at least semi-quantitative) accounts of epistemic states, particularly Spohn rankings [36] and possibility functions [14]. Deeper connections between these representations and the one of the present paper remain to be worked out, as well as the question of which representation can be said to be the *best* way to represent an agent's epistemic state. However such an investigation would probably require a paper of its own.

On the level of belief sets, our operators for revising tpos fall within the realm of non-prioritised revision, in that revision inputs are not necessarily elements of the belief set associated to the epistemic state. This is in contrast to most work on iterated belief change, which are usually given in the prioritised setting (with the works of Booth [6] and Konieczny et al. [25, 27] being exceptions). We envisage prioritised revision by α as a two-stage process, with the first stage being carried out by one of the operators in this paper, and then the second stage consisting of an application of Boutilier's *natural revision* [11] of the resulting tpo by α , i.e., the most preferred α -worlds are simply brought if necessary to the front of the new tpo . For the special case of the operator $*_P$, this was already done by Booth and Meyer [7] (Section 5), leading to the *restrained revision* operator. For future work we plan to apply this to the more general family.

Another direction for future research is the investigation of larger families of revision operators, such as those obtained by weakening one, or both, of (≤ 2) and (≤ 3). Observe that this is equivalent to weakening ($*3$) and ($*4$), or ($\circ 3$) and ($\circ 4$). The weakening of ($\circ 4$) will be of particular interest, since it is essentially equivalent to the much-criticised postulate (C2) proposed by Darwiche and Pearl [12] and reproduced in Section 2.

Conversely, it would be interesting to consider special subclasses of our general family. We considered some in Sections 7 and 8. Another example could be the family obtained by taking \ll or \lll to be modular orderings. Finally note that our operators do not conform to the principle of *categorical matching* – from an initial $\text{tpo} \leq$ together with a \leq -faithful $\text{tpo} \leq$ over W^\pm they return a new $\text{tpo} \leq_\alpha^*$, but give no help on defining a new \leq_α^* -faithful tpo over W^\pm which can then be used to further revise \leq_α^* . One way of rectifying this might be to preserve as much of \lll and \ll as possible.

For future work on SPH-revision we plan to investigate more desirable properties, and to examine useful equivalent ways to reformulate the ones we already have. In this paper all our properties are formulated as rules for single-step revision of SPHs. But since an SPH encodes the structure required to revise its associated tpo , these properties correspond to properties for *double-step* revision of tpos . To give an example, property ($\otimes 5a$) corresponds to the

following rule governing revision of a tpo \leq by α followed by β , which we denote for now by $\leq_{\alpha,\beta}^*$:

$$\text{If } x <^\alpha y \text{ and } x \leq y \text{ then } x \leq_{\alpha,\beta}^* y.$$

As mentioned above we intend to come up with other concrete SPH-revision operators, which perhaps can be described in purely qualitative terms rather than requiring extra numerical information like the operator described in this paper.

On a more fundamental level, as noted previously, the framework presented here can be viewed as a special case of *preference aggregation* or *social choice theory* [3]. We intend to pursue this link by an investigation into extending the positive and negative representations of worlds to a finer-grained version in which representations of worlds cover a larger spectrum. A comparison of such an extended framework with existing work in preference aggregation and social choice theory may well prove to be illuminating for both disciplines.

In a similar vein, there seems to be a close connection between our work and the work on preference modelling using interval orderings by Öztürk et al. [32]. The possible relationships between iterated belief revision and works such as these have, as far as we are aware, not been previously explored. We plan to look more closely at this. More generally, the question of how our work fits into the more general use of interval orderings [15] is also well worth exploring.

Finally, it is worth noting that our work can be seen as refuting the conjecture by Spohn [36] that an adequate treatment of prioritised iterated revision has to be quantitative in nature. The conjecture is based on the assumption that the only reasonable candidates for qualitative iterated revision are Boutilier's natural revision [11] and Nayak's lexicographic revision [31], but that both are flawed. The conjecture formed part of his motivation for developing a theory of *ordinal conditional functions*. While we agree with the claim that both natural and lexicographic revision can be problematic, we have shown in this paper that a qualitative setting leaves room for much more than just these two iterated revision operators.

Acknowledgements Thanks are due to Alexander Nittka for some useful comments, and to Ka-Shu Wong for contributing to an earlier version of this paper which appeared in the proceedings of KR 2006, the 10th International Conference on Principles of Knowledge Representation and Reasoning [8]. We are also grateful for insightful comments on the true contribution of this paper by two anonymous referees. Finally, this material is based upon work supported by the National Research Foundation under Grant number 65152.

References

1. Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2), 510–530.
2. Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11), 832–843.

3. Arrow, K. (1963). *Social choice and individual values*. New York: Wiley.
4. Benferhat, S., Konieczny, S., Papini, O., & Pérez, R. P. (2000). Iterated revision by epistemic states: Axioms, semantics and syntax. In *ECAI* (pp. 13–17).
5. Bochman, A. (2001). *A logical theory of nonmonotonic inference and belief change*. Berlin: Springer.
6. Booth, R. (2005). On the logic of iterated non-prioritised revision. In *Conditionals, information and inference—Selected papers from the workshop on conditionals, information and inference, 2002* (pp. 86–107). LNAI 3301. New York: Springer.
7. Booth, R., & Meyer, T. (2006). Admissible and restrained revision. *Journal of Artificial Intelligence Research (JAIR)*, 26, 127–151.
8. Booth, R., Meyer, T., & Wong, K. S. (2006). A bad day surfing is better than a good day working: How to revise a total preorder. In *KR* (pp. 230–238). AAAI Press.
9. Booth, R., Chopra, S., Ghose, A., & Meyer, T. (2010). Double preference relations for generalised belief change. *Artificial Intelligence*, 174(16–17), 1339–1368.
10. Booth, R., & Meyer, T. A. (2007). On the dynamics of total preorders: Revising abstract interval orders. In *ECSQARU* (pp. 42–53).
11. Boutilier, C. (1996). Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25(3), 263–305.
12. Darwiche, A., & Pearl, J. (1997). On the logic of iterated belief revision. *Artificial Intelligence*, 89, 1–29.
13. Delgrande, J., Dubois, D., & Lang, J. (2006). Iterated revision as prioritized merging. In *KR* (pp. 210–220).
14. Dubois, D., Lang, J., & Prade, H. (1994). *Possibilistic logic*. Clarendon.
15. Fishburn, P. C. (1985). *Interval orders and interval graphs*. New York: Wiley.
16. Freund, M. (1993). Injective models and disjunctive relations. *Journal of Logic and Computation*, 3(3), 231.
17. Freund, M., & Lehmann, D. (1994). *Belief revision and rational inference*. Technical Report TR 94-16, The Leibniz Centre for Research in Computer Science, Institute of Computer Science, Hebrew University of Jerusalem.
18. Geffner, H., & Pearl, J. (1992). Conditional entailment: Bridging two approaches to default entailment. *Artificial Intelligence*, 53, 209–244.
19. Glaister, S. M. (1998). Symmetry and belief revision. *Erkenntnis*, 49(1), 21–56.
20. Grove, A. (1988). Two modelings for theory change. *Journal of Philosophical Logic*, 17, 157–170.
21. Hansson, S. O., Fermé, E., Cantwell, J., & Falappa, M. (2001). Credibility-limited revision. *Journal of Symbolic Logic*, 66(4), 1581–1596.
22. Hansson, S. O. (1999). A survey of non-prioritized belief revision. *Erkenntnis*, 50(2), 413–427.
23. Jin, Y., & Thielscher, M. (2007). Iterated belief revision, revised. *Artificial Intelligence*, 171(1), 1–18.
24. Katsuno, H., & Mendelzon, A. O. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3), 263–294.
25. Konieczny, S., Medina Grespan, M., & Pino Pérez, R. (2010). Taxonomy of improvement operators and the problem of minimal change. In *KR*.
26. Konieczny, S., & Pérez, R. P. (2000). A framework for iterated revision. *Journal of Applied Non-Classical Logics*, 10(3–4), 339–367.
27. Konieczny, S., & Pérez, R. P. (2008). Improvement operators. In *KR* (pp. 177–186).
28. Kraus, S., Lehmann, D., & Magidor, M. (1991). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44, 167–207.
29. Lehmann, D. (1995). Belief revision, revised. In *IJCAI* (pp. 1534–1540).
30. Nayak, A., Pagnucco, M., & Peppas, P. (2003). Dynamic belief revision operators. *Artificial Intelligence*, 146, 193–228.
31. Nayak, A. C. (1994). Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41(3), 353–390.
32. Öztürk, M., Tsoukiàs, A., & Vincke, P. (2005). Preference modelling. In *Multiple criteria decision analysis: State of the art surveys* (Vol. 78, pp. 27–71). New York: Springer.
33. Papini, O. (2001). Iterated revision operations stemming from the history of an agent's observations. In *Frontiers of belief revision* (pp. 279–301).

34. Rott, H. (2001). *Change, choice and inference: A study of belief revision and nonmonotonic reasoning*. London, UK: Oxford University Press.
35. Schlechta, K., Lehmann, D., & Magidor, M. (1996). Distance semantics for belief revision. In *TARK* (pp. 137–145).
36. Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. *Causation in Decision, Belief Change, and Statistics*, 2, 105–134.