

On Strengthening the Logic of Iterated Belief Revision: Proper Ordinal Interval Operators

Richard Booth
Cardiff University
Cardiff, UK

Jake Chandler
La Trobe University
Melbourne, Australia

Abstract

Darwiche and Pearl’s seminal 1997 article outlined a number of baseline principles for a logic of iterated belief revision. These principles, the DP postulates, have been supplemented in a number of alternative ways. Most suggestions have resulted in a form of ‘reductionism’ that identifies belief states with orderings of worlds. However, this position has recently been criticised as being unacceptably strong. Other proposals, such as the popular principle (P), aka ‘Independence’, characteristic of ‘admissible’ operators, remain commendably more modest. In this paper, we supplement the DP postulates and (P) with a number of novel conditions. While the DP postulates constrain the relation between a prior and a posterior conditional belief set, our new principles notably govern the relation between two posterior conditional belief sets obtained from a common prior by different revisions. We show that operators from the resulting family, which subsumes both lexicographic and restrained revision, can be represented as relating belief states associated with a ‘proper ordinal interval’ (POI) assignment, a structure more fine-grained than a simple ordering of worlds. We close the paper by noting that these operators satisfy iterated versions of many AGM era postulates, including Superexpansion, that are not sound for admissible operators in general.

1 Introduction

Darwiche & Pearl’s (1997) seminal paper put forward a number of now popular baseline principles of iterated belief revision. These principles, the DP postulates, have been strengthened in various manners. Most proposals for doing so—such as natural (Boutilier 1996), lexicographic (Nayak, Pagnucco, and Peppas 2003), and restrained (Booth and Meyer 2006) revision (see (Peppas 2014) for an overview)—have yielded sets of principles strong enough to entail the following strong ‘reductionist’ principle: the set of beliefs held by an agent after a sequence of two revisions is fully determined by the agent’s single-step revision dispositions. This thesis can alternatively be cashed out in terms of an identification of belief states, the relata of the revision function, with total preorders (TPO’s) over possible worlds. Booth & Chandler (2017) have however recently provided considerations that suggest this reductionist posi-

tion to be too strong. Other supplements to the DP postulates, however, have fallen short of having such a consequence. This is true of the popular principle termed ‘(P)’ by Booth & Meyer (2006) and ‘Independence’ by Jin & Thielscher (2007), which, together with the DP postulates, characterises the family of ‘admissible’ revision operators that includes both lexicographic and restrained operators but excludes Boutilier’s natural ones.

In this paper, we supplement both the DP postulates and (P) with a number of novel conditions. While the DP postulates constrain the relation between a prior and a posterior conditional belief set, our new principles notably govern the relation between two posterior conditional belief sets obtained from a common prior by different revisions. We take as our foil two postulates of this variety considered by Booth & Meyer (2011). These characterised a family of *non-prioritised* revision operators, for which they offered a representation in terms of what we shall call ‘proper ordinal interval assignments’. Here, we show that these two postulates become implausible in the context of *prioritised* revision, which is the focus of the present paper. First of all, they turn out to characterise lexicographic revision when one supplements the remaining postulates of Booth & Meyer, i.e. (P) and the DP postulates, with the AGM postulate of Success. Secondly, they fall prey to an intuitive class of counterexample. After noting this, we then consider two, more plausible, weaker counterparts that have not yet been discussed in the literature. We show that these can be obtained from Booth & Meyer’s construction by adding a ‘naturalisation’ step. This is essentially an application of Boutilier’s natural revision operation to the posterior TPO obtained by Booth & Meyer’s method of non-prioritised revision. We call the resulting family of iterated revision operators, which subsumes both lexicographic and restrained revision operators, the family of ‘Proper Ordinal Interval (POI)’ operators and offer semantic and syntactic characterisations thereof. We close the paper by noting that POI revision operators satisfy iterated versions of a large number of AGM era postulates, including Superexpansion.

The plan of the remainder of the article is as follows. In Section 2, we first introduce some basic terminology and definitions. Section 3 recapitulates Booth & Meyer’s framework and introduces its two key postulates. These postulates are then critically discussed in Section 4. Section 5 outlines

our construction of the POI family of operators. In Section 6, we discuss the weakenings of Booth & Meyer’s postulates that are satisfied by the members of our new family. In Section 7, the family is then characterised semantically and syntactically, in two different manners. We wrap up the paper with a fairly substantial discussion, in Section 8, of the extent to which the members of the POI family satisfy extensions of various strong AGM era postulates to the iterated case. We then conclude in Section 9.

Due to space limitations, proofs or proof sketches are only provided for a small subset of the propositions and theorems stated. The appendix contains a detailed sketch of the proof of what is arguably our main result (Theorem 2). The remainder of the proofs are available in an extended version of the paper, which can be accessed at <http://arxiv.org/abs/1807.09942>.

2 Preliminaries

The beliefs of an agent are represented by a *belief state* Ψ . Ψ determines a *belief set* $[\Psi]$, a deductively closed set of sentences, drawn from a finitely generated propositional, truth-functional language L . Logical equivalence is denoted by \equiv and the set of logical consequences of $\Gamma \subseteq L$ by $\text{Cn}(\Gamma)$. The set of propositional worlds is denoted by W , and the set of models of a given sentence A is denoted by $\llbracket A \rrbracket$. We shall occasionally use x to denote, not the world x , but an arbitrary sentence whose set of models is $\{x\}$.

In terms of belief dynamics, our principal focus is on iterated revision—rather than contraction—operators, which return, for any prior belief state Ψ and consistent sentence A , the posterior belief state $\Psi * A$ that results from an adjustment of Ψ to accommodate the inclusion of A in $[\Psi]$.

The function $*$ is assumed to satisfy the AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985; Darwiche and Pearl 1997)—henceforth ‘AGM’, for short—which notably include the postulate of Success ($A \in [\Psi * A]$). This ensures the following convenient representability of single-shot revision: each Ψ has associated with it a total preorder \preceq_Ψ over W such that $\llbracket [\Psi * A] \rrbracket = \min(\preceq_\Psi, \llbracket A \rrbracket)$ (Katsuno and Mendelzon 1991; Grove 1988). This ordering is sometimes interpreted in terms of relative ‘(im)plausibility’, so that $x \preceq_\Psi y$ iff x is considered more ‘plausible’ than y in state Ψ . In this context, Success corresponds to the requirement that $\min(\preceq_{\Psi * A}, W) \subseteq \llbracket A \rrbracket$.

The single-shot revision dispositions associated with Ψ can also be represented by a ‘*conditional belief set*’ $[\Psi]_c$. This set extends the belief set $[\Psi]$ by further including various ‘conditional beliefs’, of the form $A \Rightarrow B$, where \Rightarrow is a non-truth-functional conditional connective. This is achieved by means of the so-called Ramsey Test, according to which $A \Rightarrow B \in [\Psi]_c$ iff $B \in [\Psi * A]$.

Following convention, we shall call principles couched in terms of belief sets ‘syntactic’, and principles couched in terms of TPOs ‘semantic’. The principles that we will discuss will be given in both types of format, with the distinction reflected in the nomenclature by the use of a subscript ‘ \preceq ’ to denote semantic principles.

We shall also be touching on a broader class of non-prioritised iterated ‘revision’ operators, for which Success

does not necessarily hold. These will be denoted by the symbol \circ . To avoid ambiguity, we will follow a convention of superscripting every principle governing a belief change operator with the relevant operator symbol (here: $*$ or \circ).

Finally, $*$ will be assumed to satisfy a principle of irrelevance of syntax that we shall call ‘Equivalence’:

$$(Eq^*) \quad \text{If } A \equiv B \text{ and } C \equiv D, \text{ then} \\ \llbracket (\Psi * A) * C \rrbracket = \llbracket (\Psi * B) * D \rrbracket$$

or semantically

$$(Eq^*_\preceq) \quad \text{If } A \equiv B, \text{ then } \preceq_{\Psi * A} = \preceq_{\Psi * B}$$

as well as the DP postulates, which constrain the belief set resulting from two successive revisions, or, equivalently, the conditional belief set resulting from a single revision:

$$(C1^*) \quad \text{If } A \in \text{Cn}(B), \text{ then } \llbracket (\Psi * A) * B \rrbracket = \llbracket \Psi * B \rrbracket$$

$$(C2^*) \quad \text{If } \neg A \in \text{Cn}(B), \text{ then } \llbracket (\Psi * A) * B \rrbracket = \llbracket \Psi * B \rrbracket$$

$$(C3^*) \quad \text{If } A \in [\Psi * B], \text{ then } A \in \llbracket (\Psi * A) * B \rrbracket$$

$$(C4^*) \quad \text{If } \neg A \notin [\Psi * B], \text{ then } \neg A \notin \llbracket (\Psi * A) * B \rrbracket$$

whose semantic counterparts are given by:

$$(C1^*_\preceq) \quad \text{If } x, y \in \llbracket A \rrbracket, \text{ then } x \preceq_{\Psi * A} y \text{ iff } x \preceq_\Psi y$$

$$(C2^*_\preceq) \quad \text{If } x, y \in \llbracket \neg A \rrbracket, \text{ then } x \preceq_{\Psi * A} y \text{ iff } x \preceq_\Psi y$$

$$(C3^*_\preceq) \quad \text{If } x \in \llbracket A \rrbracket, y \in \llbracket \neg A \rrbracket \text{ and } x \prec_\Psi y, \text{ then } x \prec_{\Psi * A} y$$

$$(C4^*_\preceq) \quad \text{If } x \in \llbracket A \rrbracket, y \in \llbracket \neg A \rrbracket \text{ and } x \preceq_\Psi y, \text{ then } x \preceq_{\Psi * A} y$$

In fact we shall further assume that $*$ satisfies the principle (P*), which strengthens both (C3*) and (C4*):

$$(P^*) \quad \text{If } \neg A \notin [\Psi * B], \text{ then } A \in \llbracket (\Psi * A) * B \rrbracket$$

Its semantic counterpart is given by:

$$(P^*_\preceq) \quad \text{If } x \in \llbracket A \rrbracket, y \in \llbracket \neg A \rrbracket \text{ and } x \preceq_\Psi y, \text{ then } x \prec_{\Psi * A} y$$

Satisfaction of AGM, (Eq*), (C1*), (C2*) and (P*) means that $*$ is an ‘*admissible*’ revision operator, in the sense of (Booth and Meyer 2006).

The constraints considered so far are notably satisfied by two well-known kinds of revision operators: restrained operators and lexicographic operators.¹ In semantic terms, these both promote the minimal A -worlds in the prior TPO to become minimal worlds in the posterior TPO. Regarding the rest of the ordering, restrained revision operators preserve the strict ordering \prec_Ψ while additionally making every A -world x strictly lower ranked than every $\neg A$ -world y for which $x \sim_\Psi y$ (where \sim_Ψ is the symmetric closure of \preceq_Ψ), so that $x \preceq_{\Psi * A} y$ iff: (i) $x \in \min(\preceq_\Psi, \llbracket A \rrbracket)$, or (ii) $x, y \notin \min(\preceq_\Psi, \llbracket A \rrbracket)$ and either (a) $x \prec_\Psi y$ or (b) $x \sim_\Psi y$ and ($x \in \llbracket A \rrbracket$ or $y \in \llbracket \neg A \rrbracket$). Lexicographic revision operators make every A -world lower ranked than every $\neg A$ -world, while preserving the ordering within each of $\llbracket A \rrbracket$ and $\llbracket \neg A \rrbracket$, so that $x \preceq_{\Psi * A} y$ iff: (i) $x \in \llbracket A \rrbracket$ and $y \in \llbracket \neg A \rrbracket$, or (ii) ($x \in \llbracket A \rrbracket$ iff $y \in \llbracket A \rrbracket$) and $x \preceq_\Psi y$. Natural revision operators, however, fail to satisfy (P*) and are thus not members of the family of admissible revision

¹Note the use of the plural here: we speak of restrained/lexicographic operators. It is of course customary, in the literature, to refer to *the* restrained/lexicographic operator. However, this way of speaking is only appropriate to the extent that belief states are identifiable with TPOs.

operators. These operators simply promote the minimal A -worlds to be $\preceq_{\Psi * A}$ -minimal, while leaving everything else unchanged, so that $x \preceq_{\Psi * A} y$ iff: (i) $x \in \min(\preceq_{\Psi}, \llbracket A \rrbracket)$, or (ii) $x, y \notin \min(\preceq_{\Psi}, \llbracket A \rrbracket)$ and $x \preceq_{\Psi} y$.

3 Two principles of non-prioritised revision

The DP postulates, as well as (P*), constrain the relation between a prior conditional belief set on the one hand, and a posterior one on the other. But one might wonder what kinds of constraints govern the relation between *two posterior conditional belief sets* obtained from a common prior by different revisions.

To the best of our knowledge, the only two articles to consider principles of this nature are (Booth and Meyer 2011) and, more briefly, (Schlechta, Lehmann, and Magidor 1996). In the former, a slightly more general form of the following pair of syntactic principles is discussed:

- ($\beta 1+^*$) If $A \notin [(\Psi * A) * B]$, then $A \notin [(\Psi * C) * B]$
($\beta 2+^*$) If $\neg A \in [(\Psi * A) * B]$, then $\neg A \in [(\Psi * C) * B]$

whose semantic counterparts are given by:

- ($\beta 1+_{\preceq}^*$) If $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$ and $y \preceq_{\Psi * A} x$, then
 $y \preceq_{\Psi * C} x$
($\beta 2+_{\preceq}^*$) If $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$ and $y \prec_{\Psi * A} x$, then
 $y \prec_{\Psi * C} x$

On the relative plausibility interpretation of \preceq , the latter can be informally glossed as follows: if (i) there exists some potential evidence, consistent with a world x but not with a world y , such that x would be considered no more plausible than (respectively: strictly less plausible than) y after receiving it, then (ii) there is no potential evidence *whatsoever* that would lead x to be considered more plausible than (respectively: at least as plausible as) y .

It is easy to see that, on the assumption that $\preceq_{\Psi * \top} = \preceq_{\Psi}$ (which follows from (C1 $_{\preceq}^*$)), these respectively generalise (C3 $_{\preceq}^*$) and (C4 $_{\preceq}^*$), which correspond to the special cases in which C is a tautology.

These postulates can be interpreted in a number of ways. One way is in terms of the binary relations (over consistent sentences in L) of *overrules* and *strictly overrules* (Booth and Meyer 2011). We say B overrules A (in Ψ) iff $A \notin [(\Psi * A) * B]$, while B strictly overrules A (in Ψ) iff $\neg A \in [(\Psi * A) * B]$.² Then ($\beta 1+^*$) says that, if B overrules A in Ψ , then A will not be believed following any sequence of two revisions starting in Ψ ending with B , while ($\beta 2+^*$) says that, if B strictly overrules A in Ψ , then A will be rejected following any such sequence of two revisions.

We noted above that it was a more *general* form of ($\beta 1+^*$), ($\beta 2+^*$) and their semantic counterparts that interested Booth & Meyer. The reason for this is that their topic of interest was not in fact $*$, but rather a more *general* kind of operator: a *non-prioritised* “revision” operator \circ , which does

²Incidentally, the first relation also corresponds to the condition under which Chandler (2017) proposed that one takes B to provide a reason to not believe A . The second relation is related to the condition under which he claimed one takes B to provide a reason to believe $\neg A$ (Chandler 2013).

not necessarily satisfy the Success postulate. They showed that these operators could be represented as relating belief states to which a certain type of structure is associated. We provide in what follows a brief overview of their framework. First, some key definitions:

Definition 1. \leq is a proper ordinal interval (POI) assignment to W iff it is a relation over $W^{\pm} = \{w^i \mid w \in W \text{ and } i \in \{-, +\}\}$ such that:

- (≤ 1) \leq is a TPO
(≤ 2) $x^+ < x^-$
(≤ 3) $x^+ \leq y^+$ iff $x^- \leq y^-$.

Definition 2. Where \preceq is a TPO over W and \leq is a POI assignment to W , we say that \leq is faithful to \preceq iff it satisfies:

- (≤ 4) $x^+ \leq y^+$ iff $x \preceq y$.

Booth & Meyer then assumed that each belief state Ψ is associated, not only with a TPO \preceq_{Ψ} , but with a POI assignment \leq_{Ψ} that is faithful to it (they remained agnostic as to whether states are to be identified with POI assignments; we will follow suit). This assignment was then taken to determine the agent’s posterior TPO upon revision by A , i.e. $\preceq_{\Psi \circ A}$, in the following manner:

Definition 3. \circ is a non-prioritised POI revision operator iff \circ is a function from state-sentence pairs to states, such that for every state Ψ there is a POI assignment \leq_{Ψ} such that, for any sentence A , $x \preceq_{\Psi \circ A} y$ iff $r_A(x) \leq_{\Psi} r_A(y)$, where

$$r_A(x) = \begin{cases} x^+ & \text{if } x \in \llbracket A \rrbracket \\ x^- & \text{if } x \in \llbracket \neg A \rrbracket. \end{cases}$$

General forms of our principles ($\beta 1+_{\preceq}^*$) and ($\beta 2+_{\preceq}^*$) turn out to play a key role in this model. Indeed, Booth & Meyer (2011, Theorem 1) show that \circ is a non-prioritised POI revision operator if and only if it satisfies (C1 $_{\preceq}^{\circ}$), (C2 $_{\preceq}^{\circ}$), (P $_{\preceq}^{\circ}$), ($\beta 1+_{\preceq}^{\circ}$) and ($\beta 2+_{\preceq}^{\circ}$), where these principles are obtained from their counterparts for (prioritised) revision in the obvious manner, by substituting the \circ symbol for $*$.

Non-prioritised POI revision operators can helpfully be understood diagrammatically. Figure 1 represents a proper ordinal interval assignment that is faithful to $x \prec_{\Psi} y \prec_{\Psi} z$. The left and right interval endpoints respectively represent the positive $(\cdot)^+$ and negative $(\cdot)^-$ counterparts of each world. Figure 2 represents, by means of the filled circles, the TPO resulting from the corresponding non-prioritised revision by $y \vee z$, i.e., $x \sim_{\Psi \circ y \vee z} y \prec_{\Psi \circ y \vee z} z$. It also illustrates failure of Success, since $x \in \min(\preceq_{\Psi \circ y \vee z}, W)$.




Figure 1: proper ordinal interval assignment

Figure 2: posterior TPO after non-prioritised POI revision by $y \vee z$

We note that lexicographic revision operators are special cases of this family in which $x^+ <_{\Psi} y^-$ for all $x, y \in W$.

4 The principles in a prioritised setting

In spite of their arguable appropriateness in a non-prioritised setting, $(\beta1+_{\succeq}^*)$ and $(\beta2+_{\succeq}^*)$ prove to be problematically strong when one imposes Success.

For one, it turns out that, in such a context the only kind of operators satisfying $(\beta1+_{\succeq}^*)$ are lexicographic revision operators, and hence that $(\beta1+_{\succeq}^*)$ imposes the reductionist assumption that we have suggested is objectionable. Indeed:

Theorem 1. *Let $*$ be a revision operator satisfying AGM and $(\beta1+_{\succeq}^*)$. Then it also satisfies the Recalcitrance property (Nayak, Pagnucco, and Peppas 2003):*

(Rec *) *If $A \wedge B$ is consistent, then $A \in [(\Psi * A) * B]$.*

Proof: If $A \wedge B$ is consistent, then $A \in [(\Psi * A \wedge B) * B]$ from AGM. Then $A \in [(\Psi * A) * B]$ by $(\beta1+_{\succeq}^*)$.³ \square

Since we know (see, e.g., (Booth and Meyer 2011; Nayak, Pagnucco, and Peppas 2003)), that lexicographic revision operators are the only admissible operators satisfying (Rec *), we obtain the following corollary, which also gives us an alternative characterisation of lexicographic revision:

Corollary 1. *The only operators satisfying AGM, (C1 *), (C2 *) and $(\beta1+_{\succeq}^*)$ are lexicographic revision operators.⁴*

These principles also face a class of direct counterexamples that match the following general pattern: A provides a defeasible reason to believe $\neg B$ (for example, let $B =$ ‘She missed the target at 5 yards’ and $A =$ ‘She is a pro archer’) and C is equivalent to the conjunction of A and a defeater for A ’s support for $\neg B$ (for example, let $C =$ ‘She is a pro archer but isn’t wearing her glasses’). Under these conditions, it can plausibly be the case that $\neg A \in [(\Psi * A) * B]$ but $A \in [(\Psi * C) * B]$, contradicting both principles.

This negative result raises the following question: Is there any way to weaken $(\beta1+_{\succeq}^*)$ and $(\beta2+_{\succeq}^*)$ to allow a wider, but intuitively plausible family of iterated prioritised revision operators? The answer, as we will now show, is ‘yes’.

5 Success via naturalisation

The guiding idea in what follows is to take the family of operators discussed in the section before last and ensure satisfaction of Success, not by adding the principle to the list of characteristic postulates but rather by minimally transforming the TPO associated with the posterior belief state by means of an operation analogous to natural revision.

More precisely, the proposal is to define $*$ as the *composition* of a non-prioritised POI revision operator \circ and a natural revision operator \boxplus :

$$\succeq_{\Psi * A} = \succeq_{(\Psi \circ A) \boxplus A}$$

Recalling the definition of natural revision in Section 2, we can equivalently say:

³Assuming consistency of inputs, the implication would also run the other way, yielding equivalence of (Rec *) and $(\beta1+_{\succeq}^*)$.

⁴What about $(\beta2+_{\succeq}^*)$? Lexicographic revision satisfies it trivially, since it satisfies: If $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$, then $x \prec_{\Psi * A} y$. We can analogously show that it implies, given AGM, the following weakening of (Rec *): If $A \wedge B$ is consistent, then $\neg A \notin [(\Psi * A) * B]$. But this is too weak to allow us to recover (Rec *) and indeed, $(\beta2+_{\succeq}^*)$ is not uniquely satisfied by lexicographic revision.

Definition 4. *$*$ is a naturalisation of \circ iff:*

- $x \preceq_{\Psi * A} y$ iff either
- (i) $x \in \min(\preceq_{\Psi \circ A}, \llbracket A \rrbracket)$, or
 - (ii) $x, y \notin \min(\preceq_{\Psi \circ A}, \llbracket A \rrbracket)$ and $x \preceq_{\Psi \circ A} y$.

We use $\mathbb{N}(*, \circ)$ to denote the fact that this relation obtains between the two functions.

Definition 5. *$*$ is a proper interval order (POI) revision operator iff $\mathbb{N}(*, \circ)$ for some non-prioritised POI revision operator \circ .*

This kind of suggestion generalises one that was made in (Booth and Meyer 2006), in which restrained revision operators were shown to be naturalisations of a particular class of non-prioritised revision operators due to Papini (2001). Indeed, the latter satisfy: $x \preceq_{\Psi \circ A} y$ iff (a) $x \prec_{\Psi} y$ or (b) $x \sim_{\Psi} y$ and $[x \in \llbracket A \rrbracket$ or $y \in \llbracket \neg A \rrbracket]$. These conditions, of course, simply correspond to (ii)(a) and (ii)(b) in the definition of restrained revision operators given in Section 2. The proposal is also somewhat reminiscent of the manner in which the Levi Identity (Levi 1977) treats non-iterated revision as the composition of a contraction and an expansion ($[\Psi * A] = [\Psi \div \neg A] + A$), with our natural revision step \boxplus playing the role of the expansion step $+$.




Figure 3: functional dependencies in POI revision

Figure 3 provides a general overview of the model, with the various arrows denoting functional determination. From bottom to top, each belief state Ψ is mapped onto a POI \leq_{Ψ} . This POI determines a TPO \preceq_{Ψ} , such that $x \preceq_{\Psi} y$ iff $x^+ \leq_{\Psi} y^+$. Finally the TPO in turn determines a belief set $[\Psi]$, such that $A \in [\Psi]$ iff $\min(\preceq, W) \subseteq \llbracket A \rrbracket$. These mappings are potentially many-to-one, so that we obtain increasingly coarse descriptions of an agent’s beliefs as one moves upwards. From left to right, the function \circ maps the prior belief state Ψ onto an ‘intermediate’ state $\Psi \circ A$, before the function \boxplus maps the latter onto the posterior state $\Psi * A = (\Psi \circ A) \boxplus A$.

We have used dashed arrows to denote some further functional dependencies. The constraints of (Booth and Meyer 2011) ensure that the prior POI assignment \leq_{Ψ} determines the ‘intermediate’ TPO $\preceq_{\Psi \circ A}$. Finally, the constraints oper-

ating on the function \boxplus ensure that *this* in turn determines the posterior TPO $\preceq_{\Psi * A}$.

This last step is achieved by moving the $\preceq_{\Psi \circ A}$ -minimal A -worlds to the front of the ordering. Figure 4 represents the result of naturalising the posterior TPO depicted in Figure 2, with y being moved into the leftmost position.




Figure 4: posterior TPO after naturalisation

The naturalisation step ensures that we have $B \in [\Psi * A]$ iff $\min(\preceq_{\Psi}, \llbracket A \rrbracket) \subseteq \llbracket B \rrbracket$, so AGM will now clearly be satisfied, including Success. Furthermore, the following general fact about naturalisation establishes that the set of POI revision operators is a subset of the set of admissible operators:

Proposition 1. *For any iterated revision operators \circ and $*$, such that $\mathbb{N}(*, \circ)$, if \circ satisfies (Eq_{\preceq}°) , $(C1_{\preceq}^{\circ})$, $(C2_{\preceq}^{\circ})$ and (P_{\preceq}°) , then $*$ will satisfy (Eq_{\preceq}^*) , $(C1_{\preceq}^*)$, $(C2_{\preceq}^*)$ and (P_{\preceq}^*) .*

Indeed, we have already noted that non-prioritised POI revision operators satisfy $(C1_{\preceq}^{\circ})$, $(C2_{\preceq}^{\circ})$ and (P_{\preceq}°) . Furthermore, Booth & Meyer show that they also satisfy (Eq_{\preceq}°) .

The family of POI revision operators includes some familiar figures:

Proposition 2. *Both lexicographic and restrained revision operators are POI revision operators.*

Indeed, we have pointed out, at the end of Section 3, that lexicographic revision operators are themselves non-prioritised POI revision operators. Furthermore, since they satisfy Success, they will be identical with their own naturalisations. Regarding restrained revision operators, the result was established in Proposition 14 of (Booth and Meyer 2006): they are, as we noted above, naturalisations of Papini’s ‘reverse’ lexicographic revision operators, which are non-prioritised POI revision operators.

6 Two weaker principles

It is easy to see that neither $(\beta 1_{\preceq}^*)$, nor $(\beta 2_{\preceq}^*)$ are generally satisfied by POI revision operators. Indeed, let $W = \{x, y, z\}$ and \preceq_{Ψ} be given as follows: $z^+ < y^+ < z^- < y^- < x^+ < x^-$. Then $y \prec_{\Psi * x \vee z} x$, but $x \prec_{\Psi * x} y$. However, as we shall see from Proposition 6 in the next section, we *do* nevertheless obtain the following weakened versions of these principles, which incorporate into their antecedents the further requirement that $x \notin \min(\preceq, \llbracket C \rrbracket)$:

$(\beta 1_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$, and $y \preceq_{\Psi * A} x$, then $y \preceq_{\Psi * C} x$

$(\beta 2_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$, and $y \prec_{\Psi * A} x$, then $y \prec_{\Psi * C} x$

Regarding the syntactic counterparts of these principles:

Proposition 3. (a) *Given AGM, $(\beta 1_{\preceq}^*)$ is equivalent to:*

$(\beta 1^*)$ If $A \notin [(\Psi * A) * B]$ and $B \rightarrow \neg A \in [\Psi * C]$, then $A \notin [(\Psi * C) * B]$

(b) *Given AGM, $(\beta 2_{\preceq}^*)$ is equivalent to:*

$(\beta 2^*)$ If $\neg A \in [(\Psi * A) * B]$ and $B \rightarrow \neg A \in [\Psi * C]$, then $\neg A \in [(\Psi * C) * B]$

These principles are particularly interesting insofar as they avoid the kind of counterexample to $(\beta 1_{\preceq}^*)$ and $(\beta 2_{\preceq}^*)$ that we raised earlier. Indeed, in the scenarios in question, we also intuitively have $B \rightarrow \neg A \notin [\Psi * C]$, rendering them perfectly consistent with the weaker $(\beta 1^*)$ and $(\beta 2^*)$.

It will turn out to be useful, in the final sections of the paper, to have noted the following equivalent formulations of $(\beta 1_{\preceq}^*)$ and $(\beta 2_{\preceq}^*)$:

Proposition 4. (a) *Given $(C2_{\preceq}^*)$ and $(C4_{\preceq}^*)$, $(\beta 1_{\preceq}^*)$ is equivalent to the conjunction of the following two principles:*

$(\gamma 1_{\preceq}^*)$ If $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$ and $y \preceq_{\Psi * A} x$, then $y \preceq_{\Psi * A \vee C} x$

$(\gamma 3_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \vee C \rrbracket$, $y \in \llbracket \neg(A \vee C) \rrbracket$, and $y \preceq_{\Psi * A \vee C} x$, then $y \preceq_{\Psi * C} x$.

(b) *Given $(C1_{\preceq}^*)$ and $(C3_{\preceq}^*)$, $(\beta 2_{\preceq}^*)$ is equivalent to the conjunction of the following two principles:*

$(\gamma 2_{\preceq}^*)$ If $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$ and $y \prec_{\Psi * A} x$, then $y \prec_{\Psi * A \vee C} x$

$(\gamma 4_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \vee C \rrbracket$, $y \in \llbracket \neg(A \vee C) \rrbracket$ and $y \prec_{\Psi * A \vee C} x$, then $y \prec_{\Psi * C} x$.

Note that, given the assumption that $\preceq_{\Psi * \top} = \preceq_{\Psi}$, which follows from $(C1_{\preceq}^*)$, $(\gamma 1_{\preceq}^*)$ and $(\gamma 2_{\preceq}^*)$ respectively entail $(C3_{\preceq}^*)$ and $(C4_{\preceq}^*)$ (let $C = \neg A$). However, none of these four new principles, and hence neither of $(\beta 1_{\preceq}^*)$ and $(\beta 2_{\preceq}^*)$, are generally sound for admissible operators:

Proposition 5. *None of $(\gamma 1_{\preceq}^*)$ to $(\gamma 4_{\preceq}^*)$ follows from AGM, $(C1_{\preceq}^*)$, $(C2_{\preceq}^*)$ and (P_{\preceq}^*) alone.*

7 Characterisations of POI operators

We have now identified a number of sound principles for the class of POI revision operators, which, we would like to remind the reader, subsumes both restrained and lexicographic operators. Next, we would like to *characterise* it.

7.1 Semantic characterisation

For our semantic characterisation, we need to introduce three more postulates, the first two of which are respective strengthenings of $(\beta 1_{\preceq}^*)$ and $(\beta 2_{\preceq}^*)$, which can be recovered by setting $z = y$:

$(\alpha 1_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$, $z \preceq_{\Psi} y$ and $y \preceq_{\Psi * A} x$, then $z \preceq_{\Psi * C} x$

$(\alpha 2_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$, $z \preceq_{\Psi} y$ and $y \prec_{\Psi * A} x$, then $z \prec_{\Psi * C} x$

$(\alpha 3_{\preceq}^*)$ If $x \notin \min(\preceq, \llbracket C \rrbracket)$, $x \in \llbracket A \rrbracket$, $y \in \llbracket \neg A \rrbracket$, $z \prec_{\Psi} y$ and $y \preceq_{\Psi * A} x$, then $z \prec_{\Psi * C} x$

Proposition 6. *$(\alpha 1_{\preceq}^*)$, $(\alpha 2_{\preceq}^*)$ and $(\alpha 3_{\preceq}^*)$ are satisfied by POI revision operators.*

These principles can perhaps be viewed as qualified pseudo-‘transitivity’ principles, if one ignores the subscripts. $(\alpha1_{\underline{z}}^*)$ and $(\alpha2_{\underline{z}}^*)$ amount to the conjunctions of $(\beta1_{\underline{z}}^*)$ and $(\beta2_{\underline{z}}^*)$ with the following semantic postulates, respectively:

$(\beta3_{\underline{z}}^*)$ If $z \neq y$, $x \notin \min(\preceq, [C])$, $x \in [A]$, $y \in [\neg A]$,
 $z \preceq_{\Psi} y$, and $y \preceq_{\Psi^*A} x$, then $z \preceq_{\Psi^*C} x$

$(\beta4_{\underline{z}}^*)$ If $z \neq y$, $x \notin \min(\preceq, [C])$, $x \in [A]$, $y \in [\neg A]$,
 $z \preceq_{\Psi} y$, and $y \prec_{\Psi^*A} x$, then $z \prec_{\Psi^*C} x$

So Proposition 6 shows that $(\beta1_{\underline{z}}^*)$ and $(\beta2_{\underline{z}}^*)$ —and hence, in view of Proposition 4, $(\gamma1_{\underline{z}}^*)$ to $(\gamma4_{\underline{z}}^*)$ —are sound for POI revision operators. We can now present our main result, which is a semantic characterisation of the family:

Theorem 2. ** is a POI revision operator iff it satisfies AGM, $(Eq_{\underline{z}}^*)$, $(C1_{\underline{z}}^*)$, $(C2_{\underline{z}}^*)$, $(P_{\underline{z}}^*)$, $(\alpha1_{\underline{z}}^*)$, $(\alpha2_{\underline{z}}^*)$, and $(\alpha3_{\underline{z}}^*)$.*

In conjunction with the results of Booth & Meyer regarding non-prioritised POI revision operators, Propositions 1 and 6 establish the left-to-right direction of the above claim. For the other direction we need to show that, if * satisfies the relevant semantic properties, then there exists a non-prioritised POI revision operator \circ such that $N(*, \circ)$. The construction works as follows: From *, define \circ by setting, for all $x, y \in W$, $x \preceq_{\Psi \circ A} y$ iff $x \preceq_{\Psi^*A \vee \neg(x \vee y)} y$. Given $(Eq_{\underline{z}}^*)$, $(C1_{\underline{z}}^*)$ and $(C2_{\underline{z}}^*)$ this is equivalent to:

$$x \preceq_{\Psi \circ A} y \text{ iff } \begin{cases} x \preceq_{\Psi} y & \text{if } x \sim^A y \\ x \preceq_{\Psi^* \neg y} y & \text{if } x \triangleleft^A y \\ x \preceq_{\Psi^* \neg x} y & \text{if } y \triangleleft^A x \end{cases}$$

where (i) $x \triangleleft^A y$ iff $x \in [A]$ or $y \in [\neg A]$, (ii) $x \sim^A y$ iff $x \triangleleft^A y$ and $y \triangleleft^A x$, (iii) $x \triangleleft^A y$ iff $x \triangleleft^A y$ but not $y \triangleleft^A x$. A detailed outline of the proof is given in the appendix.

7.2 Two syntactic characterisations

In this part we offer two different syntactic characterisations of the family of POI revision operators. The first involves the following postulates:

$(\Omega1^*)$ If $\neg A \notin [\Psi^*A \vee B]$ and $A \notin [(\Psi^*A)^*B]$,
then $B \notin [(\Psi^*B)^*A]$

$(\Omega2^*)$ If $\neg A \notin [\Psi^*A \vee B]$ and $\neg A \in [(\Psi^*A)^*B]$,
then $\neg B \in [(\Psi^*B)^*A]$

$(\Omega3^*)$ If $\neg B \in [\Psi^*A \vee B]$ and $A \notin [(\Psi^*A)^*B]$,
then $\neg B \in [(\Psi^*B)^*A]$

These principles admit an interpretation in terms of the notions of overruling and strictly overruling that we introduced in connection with $(\beta1+^*)$ and $(\beta2+^*)$. Indeed, $(\Omega1^*)$ and $(\Omega2^*)$ stipulate conditions under which the obtaining of these relations entail that of their converses, while $(\Omega3^*)$ offers a condition that is sufficient for B ’s overruling A to entail A ’s strictly overruling B .

We remark that, for both lexicographic and restrained operators, it can be shown that the overrules and strictly overrules relations collapse into the same relation. Furthermore, for lexicographic revision, we have that B overrules A iff $A \wedge B$ is inconsistent (cf. the postulate (Rec^*) in Theorem 1), while, for restrained revision, B overrules A iff both $\neg A \in [\Psi^*B]$ and $\neg B \in [\Psi^*A]$ (i.e., iff A and B counteract, to use the terminology from (Booth and Meyer 2006)).

Clearly, in both cases, the overrules relation is symmetric, and so unrestricted versions of $(\Omega1^*)$ – $(\Omega3^*)$ hold for these two sets of operators.

With Proposition 2 in mind, we now offer:

Proposition 7. *Given AGM, the following are equivalent: (a) $(Eq_{\underline{z}}^*)$, $(C1_{\underline{z}}^*)$, $(C2_{\underline{z}}^*)$, $(P_{\underline{z}}^*)$, $(\alpha1_{\underline{z}}^*)$ – $(\alpha3_{\underline{z}}^*)$, and (b) (Eq^*) , $(C1^*)$, $(C2^*)$, $(\beta1^*)$, $(\beta2^*)$, $(\Omega1^*)$ – $(\Omega3^*)$.*

While it employs some fairly accessible principles, this first characterisation ‘bundles’ the contribution of $(P_{\underline{z}}^*)$ into the principles $(\Omega1^*)$ – $(\Omega3^*)$. For this reason, we offer a second characterisation that separates out the contributions and maps each characteristic semantic principle onto a corresponding syntactic counterpart. Indeed, it turns out that the exact syntactic counterparts of $(\beta3_{\underline{z}}^*)$, $(\beta4_{\underline{z}}^*)$, and $(\alpha3_{\underline{z}}^*)$ are given as follows, where \vee denotes exclusive OR:

Proposition 8. (a) *Given AGM, $(\beta3_{\underline{z}}^*)$ is equivalent to*

$(\beta3^*)$ If $B_2 \notin [\Psi^*B_1]$, $B_1 \rightarrow A \notin [(\Psi^*A)^*B_2]$,
and $B_2 \rightarrow \neg A \in [\Psi^*C]$,
then $B_2 \wedge A \notin [(\Psi^*C)^*B_1 \vee B_2]$.

(b) *Given AGM and $(C4_{\underline{z}}^*)$, $(\beta4_{\underline{z}}^*)$ is equivalent to:*

$(\beta4^*)$ If $B_2 \notin [\Psi^*B_1]$, $B_1 \wedge \neg A \in [(\Psi^*A)^*B_2]$,
and $B_2 \rightarrow \neg A \in [\Psi^*C]$,
then $B_2 \rightarrow \neg A \in [(\Psi^*C)^*B_1 \vee B_2]$.

(c) *Given AGM and $(C3_{\underline{z}}^*)$, $(\alpha3_{\underline{z}}^*)$ is equivalent to:*

$(\alpha3^*)$ If $\neg B_2 \in [\Psi^*B_1]$, $B_1 \rightarrow A \notin [(\Psi^*A)^*B_2]$,
and $B_2 \rightarrow \neg A \in [\Psi^*C]$,
then $B_2 \rightarrow \neg A \in [(\Psi^*C)^*B_1 \vee B_2]$.

Since we already have syntactic counterparts for $(\beta1_{\underline{z}}^*)$ and $(\beta2_{\underline{z}}^*)$, as well as $(P_{\underline{z}}^*)$, the above result completes a second syntactic characterisation of the POI family. This one-to-one correspondence between semantic and syntactic principles, however, comes at a cost, since we note that $(\alpha1^*)$ – $(\alpha3^*)$ are clearly much harder to interpret than $(\Omega1^*)$ – $(\Omega3^*)$.

8 Iterated versions of AGM era postulates

In this final section of the paper, we investigate various further properties of POI revision operators, discussing in the process an interesting issue that has somewhat been neglected in the literature: the extension, to the iterated case, of the various AGM era postulates for revision.

8.1 Some postulates that are sound

In Section 6, we briefly noted that $(\beta1_{\underline{z}}^*)$ and $(\beta2_{\underline{z}}^*)$ could each be reformulated as the conjunction of a pair of principles. We showed that these principles, which had not been discussed in the literature to date, are not generally satisfied by admissible revision operators. It turns out, furthermore, that they are particularly noteworthy, since we can show that, in various combinations, they enable us to recover iterated generalisations of the following strong AGM postulates for revision and related well-known principles:

$(K7^*)$ $[\Psi^*A \wedge C] \subseteq \text{Cn}([\Psi^*A] \cup \{C\})$

(DR^*) $[\Psi^*A \vee C] \subseteq [\Psi^*A] \cup [\Psi^*C]$

(DO^*) $[\Psi^*A] \cap [\Psi^*C] \subseteq [\Psi^*A \vee C]$

(DI^*) If $\neg A \notin [\Psi^*A \vee C]$, then $[\Psi^*A \vee C] \subseteq [\Psi^*A]$

(K7*) is one of the two ‘supplementary’ AGM postulates, and is also known as ‘Superexpansion’. ‘DR’, ‘DO’ and ‘DI’ respectively abbreviate ‘Disjunctive Rationality’, ‘Disjunctive Overlap’ and ‘Disjunctive Inclusion’. As is well known in the literature, given the other AGM postulates, (DR*) is a consequence of the second supplementary postulate (K8*), aka ‘Subexpansion’, while (DO*) is equivalent to (K7*) and (DI*) to (K8*).

The iterated generalisations that we recover are obtained by replacing all mentions of the belief states in the principles above by that of their corresponding revisions by a common sentence B and making some minor adjustments. In each case, assuming Success and $[\Psi * \top] = [\Psi]$, setting $B = \top$ enables us to recover the non-iterated counterpart. We have:

$$\begin{aligned} \text{(iK7*)} \quad & [(\Psi * A \wedge C) * B] \subseteq \\ & \text{Cn}([\Psi * A] * B \cup \{A \wedge C\}) \\ \text{(iDR*)} \quad & [(\Psi * A \vee C) * B] \subseteq \\ & [(\Psi * A) * B] \cup [(\Psi * C) * B] \\ \text{(iDO*)} \quad & [(\Psi * A) * B] \cap [(\Psi * C) * B] \subseteq \\ & [(\Psi * A \vee C) * B] \\ \text{(iDI*)} \quad & \text{If } \neg A \notin [(\Psi * A \vee C) * B], \\ & \text{then } [(\Psi * A \vee C) * B] \subseteq [(\Psi * A) * B] \end{aligned}$$

Although (iK7*) and (iDI*) are, to the best of our knowledge, new to the literature, we note that (iDR*) and (iDO*) were already discussed and endorsed by Schlechta *et al* (1996). Our results are the following.

Proposition 9. *In the presence of AGM, (C1_≤*) and (C2_≤*), (a) (γ1_≤*) and (γ4_≤*) jointly entail (iDO*) and (b) (γ2_≤*) and (γ3_≤*) jointly entail (iDR*).*

Proposition 10. *Given AGM and (C1_≤*), (a) (γ1_≤*) is equivalent to (iDI*) and (b) (γ2_≤*) is equivalent to (iK7*).*

Before moving on to the next subsection, a brief comment is in order regarding the well-known results, starting with the work of Gärdenfors (1986), which show that iterated versions of even the weak ‘basic’ AGM postulates lead to triviality. How do these results square with our claims to recover iterated versions of various postulates that are substantially stronger? The answer lies in a difference in the approach to generalising these principles. In the literature surrounding the triviality results, the generalisations were obtained by replacing all references to belief sets by references to corresponding conditional belief sets. But this way of proceeding yields a significantly different set of generalisations to ours. To illustrate, consider the following ‘basic’ AGM postulate, which is a weakening of (K7*):

$$\text{(K3*)} \quad [\Psi * A] \subseteq \text{Cn}([\Psi] \cup \{A\})$$

In triviality result literature, however, its iterated counterpart is: $[\Psi * A]_c \subseteq \text{Cn}([\Psi]_c \cup \{A\})$. Etlin (2009) shows this principle to play badly with a pair of principles of conditional logic that he argues to be plausible. On our method for generating generalisations, however, we obtain the following weakening of our (iK7*):

$$\text{(iK3*)} \quad [(\Psi * A) * B] \subseteq \text{Cn}([\Psi * B] \cup \{A\})$$

This principle is insufficiently strong to play the required role in the derivations of triviality and simply turns out to be a consequence of the DP postulates.⁵

8.2 Some postulates that are *not* sound

We have not recovered the iterated version of Subexpansion:

$$\text{(K8*)} \quad \text{If } \neg C \notin [\Psi * A], \text{ then } \text{Cn}([\Psi * A] \cup \{C\}) \subseteq [\Psi * A \wedge C]$$

which is given by:

$$\text{(iK8*)} \quad \text{If } \neg(A \wedge C) \notin [(\Psi * A) * B], \text{ then } \text{Cn}([\Psi * A] * B \cup \{A \wedge C\}) \subseteq [(\Psi * A \wedge C) * B]$$

For this, we consider the following rather strong principle:

$$\text{(P+_≤*)} \quad \text{If } x \in \llbracket A \rrbracket, y \in \llbracket \neg A \rrbracket \text{ and } x \preceq_{\Psi * A \vee C} y, \text{ then } x \prec_{\Psi * A} y$$

This principle strengthens the conjunction of (γ1_≤*) and (γ2_≤*) in much the same way that (P_≤*) strengthens the conjunction of (C3_≤*) and (C4_≤*) (which, recall, are respective weakenings of (γ1_≤*) and (γ2_≤*)). Taken contrapositively, the principle inherits the weak antecedent of (γ1_≤*) but the strong consequent of (γ2_≤*). Given $\preceq_{\Psi * \top} = \preceq_{\Psi}$, which follows from (C1_≤*), (P_≤*) is recovered as the special case of (P+_≤*) in which $C = \neg A$. We now note:

Proposition 11. *(iK8*) is equivalent, given AGM and (C1_≤*), to the conjunction of (γ1_≤*) and (P+_≤*).*

Where does our POI family stand with respect to this principle? Well we can establish the following:

Proposition 12. *(P+_≤*) is satisfied by both lexicographic and restrained revision operators.*

Since lexicographic and restrained revision operators satisfy (γ1+_≤*), this establishes, that they satisfy (iK8*). This is interesting, since it shows, not only that the principle is consistent with our previous constraints, but that adding it to these does not yield the kind of ‘reductionism’ that has been argued to be objectionable. However, it remains the case that

Proposition 13. *(P+_≤*) is not generally satisfied by POI revision operators.*

In fact, a weaker property than this one fails to hold across the family. Indeed, (P+_≤*) generalises the following *Separation* property, discussed by Booth & Meyer (2006) under the name of ‘UR’, the special case of (P+_≤*) in which $C = A$:

$$\text{(Sep_≤*)} \quad \text{If } x \in \llbracket A \rrbracket \text{ and } y \in \llbracket \neg A \rrbracket, \text{ then } x \prec_{\Psi * A} y \text{ or } y \prec_{\Psi * A} x$$

⁵More specifically, it follows from (C1_≤*) and (C4_≤*). Note that (C1_≤*) and (C3_≤*) enable us to also recover the following iterated version of ‘Preservation’, which, given the other AGM postulates, is equivalent, in its non-iterated form, to (K4*): (iPres*) If $\neg A \notin [\Psi * B]$, then $[\Psi * B] \subseteq [(\Psi * A) * B]$. If one adds (P_≤*) to these principles, one also recovers the corresponding iterated version of (K4*): (iK4*) If $\neg A \notin [\Psi * B]$, then $\text{Cn}([\Psi * B] \cup \{A\}) \subseteq [(\Psi * A) * B]$. We omit the proofs here, since these claims are not central to our discussion. (iK3*) and (iPres*) are respective weakenings of (iK7*) and (iDI*). (iK4*) is a weakening of the iterated version of Subexpansion, (iK8*), discussed in the next section.

This principle can be captured by a ‘*Non-Flush*’ constraint on the POI assignment, which states that it is never the case that two intervals line up flush, in the sense that $x^+ \sim y^-$. This condition is not satisfied in general by POI assignments. Indeed, consider the following assignment to $W = \{x, y, z\}$: $x^+ < y^+ < x^- < y^- \sim z^+ < z^-$. Non-Flush fails, with the result that so too does (Sep_{\leq}^*) fails, and hence so does $(\text{P}+_{\leq}^*)$, since $y \sim_{\Psi * x \vee z} z$. This establishes Proposition 13.

At this point, a natural question arises: Why has the narrower family of POI revision operators satisfying (Sep_{\leq}^*) , or indeed, $(\text{P}+_{\leq}^*)$, not made a more central appearance in the present paper? The answer to this is that (Sep_{\leq}^*) remains in our view an extremely strong property. This becomes most apparent when one considers its syntactic counterpart:

(Sep^*) Either $\neg A \in [(\Psi * A) * B]$ or $A \in [(\Psi * A) * B]$

This principle states that, once one has revised one’s beliefs by a certain sentence, one will remain opinionated as to whether or not that sentence is true upon any further single revision. But this seems too strong: let A be any sentence and B be the sentence ‘The Oracle says that it might not be the case that A ’. Plausibly $A, \neg A \notin [(\Psi * A) * B]$.

Also of interest is the iterated version of ‘Disjunctive Factoring’, which is equivalent to the conjunction of $(\text{K}7^*)$ and $(\text{K}8^*)$, in the presence of the other AGM postulates:⁶

- (DF^*)
- (i) If $\neg C \in [\Psi * A \vee C]$, then $[\Psi * A \vee C] = [\Psi * A]$
 - (ii) If $\neg A, \neg C \notin [\Psi * A \vee C]$, then $[\Psi * A \vee C] = [\Psi * A] \cap [\Psi * C]$
 - (iii) If $\neg A \in [\Psi * A \vee C]$, then $[\Psi * A \vee C] = [\Psi * C]$

The iterated version is given by:

- (iDF^*)
- (i) If $\neg C \in [(\Psi * A \vee C) * B]$, then $[(\Psi * A \vee C) * B] = [(\Psi * A) * B]$
 - (ii) If $\neg A, \neg C \notin [(\Psi * A \vee C) * B]$, then $[(\Psi * A \vee C) * B] = [(\Psi * A) * B] \cap [(\Psi * C) * B]$
 - (iii) If $\neg A \in [(\Psi * A \vee C) * B]$, then $[(\Psi * A \vee C) * B] = [(\Psi * C) * B]$

$(\text{iDF}^*)(\text{ii})$ is entailed by the combination of (iDO^*) , for the right-to-left direction, and (iDI^*) , for the left-to-right direction, both of which we have established to be sound for POI revision operators. Regarding $(\text{iDF}^*)(\text{i})$:

Proposition 14. *The semantic counterparts of the right-to-left and left-to-right directions of $(\text{iDF}^*)(\text{i})$ are respectively:*

- $(\gamma 5_{\leq}^*)$ If $x, y \in [\neg A]$ and $y \preceq_{\Psi * A \vee C} x$, then $y \preceq_{\Psi * C} x$
 $(\gamma 6_{\leq}^*)$ If $y \in [\neg A]$, and $y \prec_{\Psi * A \vee C} x$, then $y \prec_{\Psi * C} x$.

We note that the second of these two principles, in conjunction with $(\gamma 2_{\leq}^*)$, obviously gives us $(\beta 2+_{\leq}^*)$. However, due to the requirement that $x \in [\neg A]$ in the antecedent of the

⁶This condition is typically stated in weaker terms, as: $[\Psi * A \vee C]$ is equal to either $[\Psi * A]$, $[\Psi * C]$, or $[\Psi * A] \cap [\Psi * C]$. However, the equivalence that is proven is in fact with the stronger principle. See (Gärdenfors 1988, Proposition 3.16).

second principle, the latter does not give us $(\beta 1+_{\leq}^*)$, in conjunction with $(\gamma 1_{\leq}^*)$. Where do they stand in relation to our family of operators? The answer is the following:

Proposition 15. *Neither $(\gamma 5_{\leq}^*)$ nor $(\gamma 6_{\leq}^*)$ are generally satisfied by POI revision operators.*

However:

Proposition 16. *Both $(\gamma 5_{\leq}^*)$ and $(\gamma 6_{\leq}^*)$ are satisfied by lexicographic revision operators.*

This establishes that (iDF^*) is satisfied by lexicographic revision operators, since we have individually shown that they satisfy all the component principles.

9 Conclusions and further work

This paper has investigated a significant, yet comparatively restrained, strengthening of the seminal framework introduced two decades ago by Darwiche and Pearl. Unlike the majority of existing models of iterated revision, ours falls short of identifying belief states with simple total preorders over worlds. Indeed, it incorporates further structure into these, in the form of proper ordinal intervals.⁷ This is achieved by combining Booth & Meyer’s framework for non-prioritised revision with a ‘naturalisation’ step. The resulting family of POI revision operators, which is a subfamily of the so-called ‘admissible’ family, has been characterised both semantically and syntactically. It has also been shown that POI revision operators are distinctive, within the class of admissible ones, in satisfying iterated counterparts of many (albeit not all) classic AGM era postulates.

In future work, we first plan to consider the consequences of relaxing (P^*) . This condition fails for a more general family of ‘*basic ordinal interval (BOI)*’ revision operators. These operators, which include natural revision operators, are naturalisations of non-prioritised operators based on ordinal interval assignments that satisfy, not (≤ 2) , but the weaker requirement that $x^+ \leq x^-$. As it turns out, our proof of the soundness of $(\alpha 1_{\leq}^*)$ – $(\alpha 3_{\leq}^*)$ with respect to POI operators carries over here, leaving us in a strong position to provide a characterisation for this more general family. Secondly, as Figure 3 reminds us, the constraints that we have discussed impose few constraints on the result of more than two iterations of the revision operation. While the structure associated with belief states currently determines the posterior *TPO*, nothing has been said regarding the nature of the posterior *POI*.

⁷Others have also enriched belief states beyond mere TPOs or equivalent structures. Spohn, for instance, identifies states with ‘*ranking functions*’, aka ‘*OCFs*’ (Spohn 1988; 2012). Others, such as Konieczny & Pérez (2000) identify states with *histories of input sentences* (see also (Lehmann 1995) and (Delgrande, Dubois, and Lang 2006) for related work). However, the comparison with Spohn’s approach is not straightforward, since his model features a whole parameterised family of revision-like functions. Regarding Konieczny & Pérez’s model, if $(\text{C}2^*)$ is imposed, then their proposal reduces to lexicographic revision.

Acknowledgements

This research was partially supported by the Australian Government through an Australian Research Council Future Fellowship (project number FT160100092) awarded to Jake Chandler.

Appendix

Theorem 2. $*$ is a POI revision operator iff it satisfies AGM, (Eq_{\leq}^*) , $(C1_{\leq}^*)$, $(C2_{\leq}^*)$, (P_{\leq}^*) , $(\alpha1_{\leq}^*)$, $(\alpha2_{\leq}^*)$, and $(\alpha3_{\leq}^*)$.

Detailed proof sketch: As remarked earlier, soundness follows from the results of Booth & Meyer and Propositions 1 and 6. For the completeness part, let \circ be defined from $*$ as it was just after Theorem 2 in the main text. We establish the result by proving two main lemmas: first, we show that \circ is a non-prioritised POI revision operator (Lemma 1) and then we show that $\mathbb{N}(*, \circ)$ (Lemma 3).

Lemma 1. \circ is a non-prioritised POI revision operator

We show that \circ satisfies each of $(C1_{\leq}^{\circ})$, $(C2_{\leq}^{\circ})$, (P_{\leq}°) , $(\beta1_{\leq}^{\circ})$ and $(\beta2_{\leq}^{\circ})$, as well as the requirement that $\leq_{\Psi \circ A}$ is a TPO over W .

The main difficulty lies in establishing the soundness of this last property. To do this we will use the following useful auxiliary lemma:

Lemma 2. Let x, y, z be distinct worlds such that $y \leq_{\Psi} z$. Then the following are equivalent: (i) If $x \leq_{\Psi^* \neg y} y$, then $x \leq_{\Psi^* \neg z} z$, and (ii) If $x \leq_{\Psi^* x \vee z} y$, then $x \leq_{\Psi^* x \vee y} z$.

Totality of $\leq_{\Psi \circ A}$ follows from totality of $\leq_{\Psi^* A \vee \neg(x \vee y)}$ and (Eq_{\leq}^*) , since we have $x \leq_{\Psi \circ A} y \Leftrightarrow x \leq_{\Psi^* A \vee \neg(x \vee y)} y$. To show that $\leq_{\Psi \circ A}$ is transitive, we consider the eight cases according to whether each of x, y , and z is in $\llbracket A \rrbracket$ or not:

- (i) $x, y, z \in \llbracket A \rrbracket$ or $x, y, z \in \llbracket \neg A \rrbracket$: Follows from transitivity for $*$.
- (ii) $x, y \in \llbracket A \rrbracket, z \in \llbracket \neg A \rrbracket$: Then we must show that, if $x \leq_{\Psi} y$ and $y \leq_{\Psi^* \neg z} z$, then $x \leq_{\Psi^* \neg z} z$. Since $x, y \in \llbracket A \rrbracket$ and $z \in \llbracket \neg A \rrbracket$, we know that $x \neq z$ and $y \neq z$. Then from $x \leq_{\Psi} y$ and $(C1_{\leq}^*)$, we obtain $x \leq_{\Psi^* \neg z} y$. From the latter and $y \leq_{\Psi^* \neg z} z$, we then obtain $x \leq_{\Psi^* \neg z} z$ by transitivity for $*$.
- (iii) $x \in \llbracket A \rrbracket, y \in \llbracket \neg A \rrbracket, z \in \llbracket A \rrbracket$: Then we must show that, if $x \leq_{\Psi^* \neg y} y$ and $y \leq_{\Psi^* \neg y} z$, then $x \leq_{\Psi} z$. By transitivity for $*$, it follows, from $x \leq_{\Psi^* \neg y} y$ and $y \leq_{\Psi^* \neg y} z$, that $x \leq_{\Psi^* \neg y} z$. From $x, z \in \llbracket A \rrbracket$ and $y \in \llbracket \neg A \rrbracket$, we know $x \neq y$ and $z \neq y$. So from $x \leq_{\Psi^* \neg y} z$ and $(C1_{\leq}^*)$, we obtain $x \leq_{\Psi} z$.
- (iv) $x \in \llbracket A \rrbracket, y, z \in \llbracket \neg A \rrbracket$: Then we must show that, if $x \leq_{\Psi^* \neg y} y$ and $y \leq_{\Psi} z$, then $x \leq_{\Psi^* \neg z} z$. If $z = y$, then $x \leq_{\Psi^* \neg z} z$ follows immediately from $x \leq_{\Psi^* \neg y} y$. So we may assume $z \neq y$. By Lemma 2, what we must establish is then equivalent to: if $x \leq_{\Psi^* x \vee z} y$ and $y \leq_{\Psi} z$, then $x \leq_{\Psi^* x \vee y} z$. Or contrapositing: if $z \prec_{\Psi^* x \vee y} x$ and $y \leq_{\Psi} z$, then $y \prec_{\Psi^* x \vee z} x$. So assume $z \prec_{\Psi^* x \vee y} x$ and $y \leq_{\Psi} z$. Now, if $x \leq_{\Psi} z$, then $x \leq_{\Psi^* x \vee y} z$ by $(C3_{\leq}^*)$. So assume $z \prec_{\Psi} x$. We therefore have: $x \in \llbracket x \vee y \rrbracket, z \notin \llbracket x \vee y \rrbracket, z \prec_{\Psi^* x \vee y} x, y \leq_{\Psi} z$ and $x \notin \min(\leq, \llbracket x \vee z \rrbracket)$. From this, by $(\alpha2_{\leq}^*)$, we can then infer that $y \prec_{\Psi^* x \vee z} x$.

(v) $x \in \llbracket \neg A \rrbracket, y, z \in \llbracket A \rrbracket$: Then we must show that, if $x \leq_{\Psi^* \neg x} y$ and $y \leq_{\Psi} z$, then $x \leq_{\Psi^* \neg x} z$. Since $x \in \llbracket \neg A \rrbracket$ and $y, z \in \llbracket A \rrbracket$, we know that $x \neq y$ and $x \neq z$. Hence from $y \leq_{\Psi} z$, we know $y \leq_{\Psi^* \neg x} z$. The implication then follows from transitivity for $*$.

(vi) $x \in \llbracket \neg A \rrbracket, y \in \llbracket A \rrbracket, z \in \llbracket \neg A \rrbracket$: Then we must show that, if $x \leq_{\Psi^* \neg x} y$ and $y \leq_{\Psi^* \neg z} z$, then $x \leq_{\Psi} z$, or, equivalently, that, if $x \leq_{\Psi^* \neg x} y$ and $z \prec_{\Psi} x$, then $z \prec_{\Psi^* \neg z} y$. So suppose $x \leq_{\Psi^* \neg x} y$ and $z \prec_{\Psi} x$. If $y \leq_{\Psi} x$, then by (P_{\leq}^*) we would have $y \prec_{\Psi^* \neg x} x$: contradiction. Hence we may assume $x \prec_{\Psi} y$. From this and $z \prec_{\Psi} x$ we have, by transitivity, $z \prec_{\Psi} y$ and hence $y \notin \min(\leq, \llbracket y \vee z \rrbracket)$. From this and $x \leq_{\Psi^* \neg x} y$, using postulate $(\beta1_{\leq}^*)$, we can deduce $x \leq_{\Psi^* y \vee z} y$. We therefore have: $y \in \llbracket y \vee z \rrbracket, x \notin \llbracket y \vee z \rrbracket, x \leq_{\Psi^* y \vee z} y, z \prec_{\Psi} x$ and $y \notin \min(\leq_{\Psi}, \llbracket x \vee y \rrbracket)$. From this, by $(\alpha3_{\leq}^*)$, we can then infer that $z \prec_{\Psi^* x \vee y} x$, and so $z \prec_{\Psi^* \neg z} y$, by $(\gamma2_{\leq}^*)$, as required.

(vii) $x, y \in \llbracket \neg A \rrbracket, z \in \llbracket A \rrbracket$: Then we must show that, if $x \leq_{\Psi} y$ and $y \leq_{\Psi^* \neg y} z$, then $x \leq_{\Psi^* \neg x} z$. If $x = y$, then this holds immediately, so we may assume $x \neq y$. Now suppose $x \leq_{\Psi} y$ and $y \leq_{\Psi^* \neg y} z$. If $z \leq_{\Psi} y$, then $z \prec_{\Psi^* \neg y} y$ by (P_{\leq}^*) : contradiction. So we may assume $y \prec_{\Psi} z$. From this and $x \leq_{\Psi} y$, we know, by transitivity, that $x \prec_{\Psi} z$, so $z \notin \min(\leq_{\Psi}, \llbracket x \vee z \rrbracket)$. It then follows that $y \leq_{\Psi^* x \vee z} z$ by postulate $(\beta1_{\leq}^*)$. We therefore have: $z \in \llbracket x \vee z \rrbracket, y \notin \llbracket x \vee z \rrbracket, y \leq_{\Psi^* x \vee z} z, x \prec_{\Psi} y$ and $z \notin \min(\leq_{\Psi}, \llbracket y \vee z \rrbracket)$. From this, by $(\alpha1_{\leq}^*)$, we can then infer that $x \prec_{\Psi^* y \vee z} z$, and so $x \leq_{\Psi^* \neg x} z$, by $(\gamma1_{\leq}^*)$, as required.

We now prove our second main lemma:

Lemma 3. $\mathbb{N}(*, \circ)$

We require: $x \leq_{\Psi^* A} y$ iff (i) $x \in \min(\leq_{\Psi \circ A}, \llbracket A \rrbracket)$, or (ii) $x, y \notin \min(\leq_{\Psi \circ A}, \llbracket A \rrbracket)$ and $x \leq_{\Psi^* A \vee \neg(x \vee y)} y$. Since \circ satisfies $(C1_{\leq}^{\circ})$, we can however replace this with: $x \leq_{\Psi^* A} y$ iff (i) $x \in \min(\leq, \llbracket A \rrbracket)$, or (ii) $x, y \notin \min(\leq, \llbracket A \rrbracket)$ and $x \leq_{\Psi^* A \vee \neg(x \vee y)} y$. With this in mind:

(a) **Left-to-right direction:** Suppose that $x \leq_{\Psi^* A} y$ and $x \notin \min(\leq, \llbracket A \rrbracket)$. If $y \in \min(\leq, \llbracket A \rrbracket)$, then $y \prec_{\Psi^* A} x$, by Success: contradiction. Hence $y \notin \min(\leq, \llbracket A \rrbracket)$. It remains to be shown that $x \leq_{\Psi^* A \vee \neg(x \vee y)} y$. If $x \sim^A y$, then the conclusion follows by $(C1_{\leq}^*)$ – $(C2_{\leq}^*)$. If $y \triangleleft^A x$, then the conclusion follows from $x \leq_{\Psi^* A} y$ and $(\gamma1_{\leq}^*)$. Finally, if $x \triangleleft^A y$, then $x \in \llbracket A \vee \neg(x \vee y) \rrbracket$ and $y \in \llbracket \neg(A \vee \neg(x \vee y)) \rrbracket$. Together with $x \notin \min(\leq_{\Psi}, \llbracket A \rrbracket)$ and $x \leq_{\Psi^* A} y$, the desired conclusion then follows by postulate $(\beta2_{\leq}^*)$.

(b) **Right-to-left direction:** If $x \in \min(\leq, \llbracket A \rrbracket)$, then $x \leq_{\Psi^* A} y$ by Success. So suppose $x, y \notin \min(\leq, \llbracket A \rrbracket)$ and $x \leq_{\Psi^* A \vee \neg(x \vee y)} y$. We must show $x \leq_{\Psi^* A} y$. If $x \sim^A y$, then the conclusion follows by $(C1_{\leq}^*)$ – $(C2_{\leq}^*)$. If $x \triangleleft^A y$, then the conclusion follows by $(\gamma2_{\leq}^*)$. Finally, if $y \triangleleft^A x$, then the conclusion follows from postulate $(\beta1_{\leq}^*)$. \square

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2):510–530.
- Booth, R., and Chandler, J. 2017. The irreducibility of iterated to single revision. *Journal of Philosophical Logic* 46(4):405–418.
- Booth, R., and Meyer, T. 2006. Admissible and restrained revision. *Journal of Artificial Intelligence Research* 26:127–151.
- Booth, R., and Meyer, T. 2011. How to revise a total pre-order. *Journal of Philosophical Logic* 40(2):193–238.
- Boutilier, C. 1996. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic* 25(3):263–305.
- Chandler, J. 2013. Transmission failure, AGM-style. *Erkenntnis* 78(2):383–398.
- Chandler, J. 2017. Preservation, commutativity and modus ponens: Two recent triviality results. *Mind* 126(502):579–602.
- Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial intelligence* 89(1-2):1–29.
- Delgrande, J. P.; Dubois, D.; and Lang, J. 2006. Iterated revision as prioritized merging. In *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR 2006)*, 210–220.
- Etlin, D. 2009. The problem of noncounterfactual conditionals. *Philosophy of Science* 76(5):676–688.
- Gärdenfors, P. 1986. Belief revisions and the Ramsey Test for conditionals. *The Philosophical Review* 95(1):81–93.
- Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press Cambridge, Mass.
- Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17(2):157–170.
- Jin, Y., and Thielscher, M. 2007. Iterated belief revision, revised. *Artificial Intelligence* 171(1):1–18.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.
- Konieczny, S., and Pino Pérez, R. 2000. A framework for iterated revision. *Journal of Applied Non-Classical Logics* 10(3-4):339–367.
- Lehmann, D. 1995. Belief revision, revised. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI 1995)*, 1534–1540.
- Levi, I. 1977. Subjunctives, dispositions and chances. *Synthese* 34(4):423–455.
- Nayak, A. C.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *Artificial Intelligence* 146(2):193–228.
- Papini, O. 2001. Iterated revision operations stemming from the history of an agent's observations. In *Frontiers in Belief Revision*. Springer Netherlands. 279–301.
- Peppas, P. 2014. A panorama of iterated revision. In *David Makinson on Classical Methods for Non-classical Problems*. Springer. 71–94.
- Schlechta, K.; Lehmann, D.; and Magidor, M. 1996. Distance semantics for belief revision. In *Proceedings of the 6th conference on Theoretical Aspects of Rationality and Knowledge (TARK 1996)*, 137–145.
- Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In *Causation in Decision, Belief Change, and Statistics, vol. II*. Kluwer Academic Publishers.
- Spohn, W. 2012. *The Laws of Belief: Ranking Theory and its Philosophical Applications*. Oxford University Press.