Revision by Conditionals: From Hook to Arrow

Jake Chandler[1] , **Richard Booth**[2]

[1]La Trobe University
[2]Cardiff University
jacob.chandler@latrobe.edu.au, boothr2@cardiff.ac.uk

Abstract

The belief revision literature has largely focussed on the issue of how to revise one's beliefs in the light of information regarding matters of fact. Here we turn to an important but comparatively neglected issue: How might one extend a revision operator to handle *conditionals* as input? Our approach to this question of 'conditional revision' is distinctive insofar as it abstracts from the controversial details of how to revise by factual sentences. We introduce a 'plug and play' method for uniquely extending any iterated belief revision operator to the conditional case. The flexibility of our approach is achieved by having the result of a conditional revision by a Ramsey Test conditional ('arrow') determined by that of a plain revision by its corresponding material conditional ('hook'). It is shown to satisfy a number of new constraints that are of independent interest.

1 Introduction

The past three decades have witnessed the development of a substantial, if inconclusive, body of work devoted to the issue of *belief revision*, namely

(A) determining the impact of a local change in belief on both (i) the remainder of one's prior beliefs and (ii) one's prior *conditional beliefs* ('Ramsey Test conditionals').

Surprisingly, however, very little has been done to this date on the question of *conditional belief revision*, that is

(B) determining the impact of a local change in conditional beliefs on both (i) and (ii).

Furthermore, nearly all of the few proposals to tackle issue (B), namely (Hansson 1992), (Boutilier and Goldszmidt 1993), and (Nayak et al. 1996), have typically rested on somewhat contentious assumptions about how to approach (A). (A noteworthy exception to this (Kern-Isberner 1999), who introduced a number of plausible general postulates governing revision by conditionals whose impact on revision simpliciter remains fairly modest. More on these below.)

In this paper, we consider the prospects of providing a 'plug and play' solution to issue (B) that is independent of the details of how to address (A). Its remainder is organised as follows. First, in Section 2, we present some standard background on problem (A), introducing along the way the well-known notion of a Ramsey Test conditional or again

conditional belief. In Section 3, we outline and discuss our proposal regarding (B). Subsection 3.1 presents the basic idea, according to which computing the result of a revision by a Ramsey Test conditional can be derived by minimal modification, under constraints, of the outcome of a revision by its corresponding material conditional. Our key technical contribution is presented in Subsection 3.2, where we prove that this minimal change under constraints can be achieved by means of a simple and familiar transformation. Subsection 3.3 outlines some interesting general properties of the proposal. These strengthen, in a plausible manner, the aforementioned constraints presented in (Kern-Isberner 1999) and are of independent interest. Subsection 3.4 considers the upshot of pairing our proposal regarding (B) with some well-known suggestions regarding how to tackle (A). Finally, in Section 4, we compare the suggestion made with existing work on the topic noting some important shortcomings of the latter. We close the paper in Section 5 with a number of questions for future research.

Due to space limitations, only a couple of the more important proofs have been provided, in Section 6. A version of the paper containing all proofs can be accessed online at http://arxiv.org/abs/2006.15811.

2 Revision

The beliefs of an agent are represented by a *belief state*. Such states will be denoted by upper case Greek letters $\Psi, \Theta, \ldots$. We denote by $\mathbb{S}$ the set of all such states. Each state determines a *belief set*, a consistent and deductively closed set of sentences, drawn from a finitely generated propositional, truth-functional language L, equipped with the standard connectives $\supset$, $\wedge$, $\vee$, and $\neg$. We denote the belief set associated with state Ψ by $[\Psi]$. Logical equivalence is denoted by $\equiv$ and the set of classical logical consequences of $\Gamma \subseteq L$ by $\text{Cn}(\Gamma)$, with $\top$ denoting an arbitrary propositional tautology. The set of propositional worlds will be denoted by W and the set of models of a given sentence A by $[\![A]\!]$.

The operation of *revision* $*$ returns the posterior state $\Psi * A$ that results from an adjustment of Ψ to accommodate the inclusion of the consistent input A in its associated belief set, in such a way as to maintain consistency of the resulting belief set.

The beliefs resulting from single revisions are conveniently representable by a *conditional belief set* $[\Psi]_c$, which

can be viewed as encoding the agent's rules of inference over L in state Ψ. It is defined via the Ramsey Test:

(RT)   For all $A, B \in L$, $A \Rightarrow B \in [\Psi]_c$ iff $B \in [\Psi * A]$

We shall call L_c the minimal extension of L that additionally includes all sentences of the form $A \Rightarrow B$, with $A, B \in L$. We shall call sentences of the form $A \Rightarrow B$ 'conditionals' and sentences of the form $A \supset B$ 'material conditionals'. We shall say that a sentence of the form $A \Rightarrow B$ is consistent just in case $A \wedge B$ is consistent (later in the paper, we shall explicitly disallow revisions by inconsistent conditionals).

Conditional belief sets are constrained by the AGM postulates of (Alchourrón, Gärdenfors, and Makinson 1985; Darwiche and Pearl 1997) (henceforth 'AGM'). Given these, $[\Psi]_c$ corresponds to a consistency-preserving *rational consequence relation*, in the sense of (Lehmann and Magidor 1992). Equivalently, it is representable by a *total preorder (TPO)* $\preccurlyeq_\Psi$ of worlds, such that $A \Rightarrow B \in [\Psi]_c$ iff $\min(\preccurlyeq_\Psi, [\![A]\!]) \subseteq [\![B]\!]$ (Grove 1988; Katsuno and Mendelzon 1991). Note that $A \in [\Psi]$ iff $\top \Rightarrow A \in [\Psi]_c$ or equivalently iff $\min(\preccurlyeq_\Psi, W) \subseteq [\![A]\!]$.

Following convention, we shall call principles presented in terms of belief sets 'syntactic', and call 'semantic' those principles couched in terms of TPOs, denoting the latter by subscripting the corresponding syntactic principle with '$\preccurlyeq$'. Due to space considerations and for ease of exposition, we will largely restrict our focus to a semantic perspective on our problem of interest.

The AGM postulates do not entail that one's conditional beliefs are determined by one's beliefs—in the sense that, if $[\Psi] = [\Theta]$, then $[\Psi]_c = [\Theta]_c$—and there is widespread consensus that such determination would be unduly restrictive, with (Hansson 1992) providing supporting arguments. A fortiori, one should not identify conditional beliefs with beliefs in the corresponding material conditional. That said, there does remain a connection between $A \Rightarrow B \in [\Psi]_c$ and $A \supset B \in [\Psi]$. The following is well known:

Proposition 1. *Given AGM, (a) if $A \Rightarrow B \in [\Psi]_c$, then $A \supset B \in [\Psi]$, but (b) the converse does not hold.*

Indeed, (a) is simply equivalent, given (RT), to the AGM postulate of Inclusion, according to which $[\Psi * A] \subseteq \mathrm{Cn}([\Psi] \cup \{A\})$. This suggests the following catchline:

> 'Conditional beliefs are beliefs in material conditionals plus'

That is, conditional beliefs are beliefs in material conditionals that satisfy certain additional constraints.

Regarding the conditional beliefs resulting from single revisions, i.e. the beliefs resulting from sequences of two revisions, we assume an 'irrelevance of syntax' property, which, in its semantic form, is given by:

$(\mathrm{Eq}^*_\preccurlyeq)$ If $A \equiv B$, then $\preccurlyeq_{\Psi * A} = \preccurlyeq_{\Psi * B}$

Given this principle, we take the liberty to abuse both language and notation and occasionally speak of revision by a set of worlds S rather than by an arbitrary sentence whose set of models is given by S.

The DP postulates of (Darwiche and Pearl 1997) provide widely endorsed further constraints. We simply give them here in their semantic form:

$(\mathrm{C1}^*_\preccurlyeq)$ If $x, y \in [\![A]\!]$ then $x \preccurlyeq_{\Psi * A} y$ iff $x \preccurlyeq_\Psi y$

$(\mathrm{C2}^*_\preccurlyeq)$ If $x, y \in [\![\neg A]\!]$ then $x \preccurlyeq_{\Psi * A} y$ iff $x \preccurlyeq_\Psi y$

$(\mathrm{C3}^*_\preccurlyeq)$ If $x \in [\![A]\!]$, $y \in [\![\neg A]\!]$ and $x \prec_\Psi y$, then $x \prec_{\Psi * A} y$

$(\mathrm{C4}^*_\preccurlyeq)$ If $x \in [\![A]\!]$, $y \in [\![\neg A]\!]$ and $x \preccurlyeq_\Psi y$, then $x \preccurlyeq_{\Psi * A} y$

Importantly, while there appears to be a degree of consensus that these postulates should be strengthened, there is no agreement as to how this should be done. Popular options include the principles respectively associated with the operators of natural revision $*_\mathrm{N}$ (Boutilier 1996), restrained revision $*_\mathrm{R}$ (Booth and Meyer 2006) and lexicographic revision $*_\mathrm{L}$ (Nayak, Pagnucco, and Peppas 2003), semantically defined as follows:

Definition 1. *The operators $*_\mathrm{N}$, $*_\mathrm{R}$ and $*_\mathrm{L}$ are such that:*

> $x \preccurlyeq_{\Psi *_\mathrm{N} A} y$ *iff (1)* $x \in \min(\preccurlyeq_\Psi, [\![A]\!])$, *or (2)* $x, y \notin \min(\preccurlyeq_\Psi, [\![A]\!])$ *and* $x \preccurlyeq_\Psi y$

> $x \preccurlyeq_{\Psi *_\mathrm{R} A} y$ *iff (1)* $x \in \min(\preccurlyeq_\Psi, [\![A]\!])$, *or (2)* $x, y \notin \min(\preccurlyeq_\Psi, [\![A]\!])$ *and either (a)* $x \prec_\Psi y$ *or (b)* $x \sim_\Psi y$ *and* $(x \in [\![A]\!]$ *or* $y \in [\![\neg A]\!])$

> $x \preccurlyeq_{\Psi *_\mathrm{L} A} y$ *iff (1)* $x \in [\![A]\!]$ *and* $y \in [\![\neg A]\!]$, *or (2)* $(x \in [\![A]\!]$ *iff* $y \in [\![A]\!])$ *and* $x \preccurlyeq_\Psi y$.

The suitability of all three operators, which we will group here under the heading of '*elementary* revision operators' (Chandler and Booth 2019), has been called into question. Indeed, they assume that a state Ψ can be identified with its corresponding TPO $\preccurlyeq_\Psi$ and that belief revision functions map pairs of TPOs and sentences onto TPOs. (For this reason, we will sometimes abuse language and notation and speak, for instance, of the lexicographic revision of a TPO rather than of a state.) But this assumption has been criticised as implausible, with (Booth and Chandler 2017) providing a number of counterexamples.

Accordingly, (Booth and Chandler 2018; Booth and Chandler 2020) propose a strengthening of the DP postulates that is weak enough to avoid an identification of states with TPOs and is consistent with the characteristic postulates of both $*_\mathrm{R}$ and $*_\mathrm{L}$ (albeit not of $*_\mathrm{N}$). They suggest associating states with structures that are richer than TPOs: 'proper ordinal interval (POI) assignments'.

3   Conditional Revision

We now turn to our question of interest: How might one extend a revision operator to handle conditionals as inputs? We shall call such an extended operator, which maps pairs of states and consistent sentences in L_c onto states, a *conditional revision* operator.

In view of the considerable disagreement regarding revision that we noted in the previous section, it would be desirable to find a solution that abstracts from some of the details regarding how this problem is handled. In what follows, we shall propose a method that achieves just this. The idea that we will exploit is that the result of a conditional revision by a Ramsey Test conditional is determined by that of a plain revision by its corresponding material conditional.

More specifically, we will be suggesting the following kind of procedure for constructing $\preccurlyeq_{\Psi * A \Rightarrow B}$:

(1) Determine $\preccurlyeq_{\Psi * A \supset B}$.

(2) Remain as 'close' to this TPO as possible, while:

(a) ensuring that $A \Rightarrow B \in [\Psi * A \Rightarrow B]_{c}$, and

(b) retaining some of $\preccurlyeq_{\Psi * A \supset B}$'s relevant features.

Our proposal then is to derive $\preccurlyeq_{\Psi * A \Rightarrow B}$ from $\preccurlyeq_{\Psi * A \supset B}$, via distance minimisation under constraints. Importantly, this suggestion *does not tie us to any particular revision operator*, since it takes $\preccurlyeq_{\Psi * A \supset B}$ as its starting point, irrespective of how it is arrived at.

3.1 Distance-Minimisation Under Constraints

In an early paper on conditional revision, (Nayak et al. 1996) suggest that the task of conditional revision is no different from that of revision by the corresponding material conditional. Indeed, they note that, on their view of rational revision, whereby they identify $*$ with lexicographic revision $*_{L}$, revision by the material conditional is sufficient to ensure that the corresponding conditional is included in the resulting conditional belief set. In other words, identifying $*A \Rightarrow B$ with $*_{L} A \supset B$ is sufficient to secure the following desirable property of '*Success*' for conditional revisions:

$(S^{*}_{\preccurlyeq})$ $\quad \min(\preccurlyeq_{\Psi * A \Rightarrow B}, [A]) \subseteq [B]$

(S^{*}) $\quad A \Rightarrow B \in [\Psi * A \Rightarrow B]_{c}$

Since we have, in Section 2, rejected identifying rational revision with lexicographic revision, Nayak *et al's* proposal is not on the cards for us. But one might still wonder whether there exists a more acceptable conception of iterated revision that, like lexicographic revision, allows us to meet the requirement of Success by simply revising by the material conditional. But it is easy to find counterexamples to the inclusion $\min(\preccurlyeq_{\Psi * A \supset B}, [A]) \subseteq [B]$ for the best known strengthenings of the DP postulates (Figure 1 provides a case in point for restrained revision). In fact, we can easily show that, given mild conditions, lexicographic revision is the only operator that fits the bill:

Proposition 2. *If $*$ satisfies AGM, $(C1^{*}_{\preccurlyeq})$, $(C2^{*}_{\preccurlyeq})$, $(Eq^{*}_{\preccurlyeq})$, and the principle according to which, for all $A, B \in L$ and $\Psi \in S$, $A \Rightarrow B \in [\Psi * A \supset B]_{c}$, then $* = *_{L}$.*
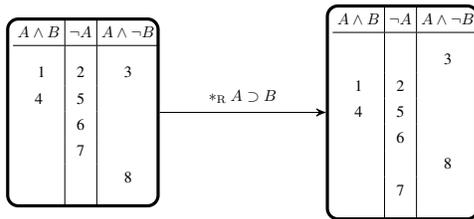

Figure 1: Illustration of $\min(\preccurlyeq_{\Psi * A \supset B}, [A]) \not\subseteq [B]$ with $* = *_{R}$. The relation $\preceq_{\Psi}$ orders the worlds—depicted by numbers—from bottom to top, with the minimal world on the lowest level. The columns group worlds according to the sentences that they validate. We can see that, here, $\min(\preccurlyeq_{\Psi *_{R} A \supset B}, [A]) = \{8\} \subset [\neg B]$.

Short of endorsing lexicographic revision, then, revision by the corresponding material conditional is not sufficient for the inclusion of a conditional in the resulting belief set. So just as conditional beliefs can be viewed as 'beliefs in material conditionals plus', we could say that:

'Conditional revision is revision by material conditionals plus.'

How, then, might we plausibly modify $\preccurlyeq_{\Psi * A \supset B}$ so as to arrive at a TPO $\preccurlyeq_{\Psi * A \Rightarrow B}$ that satisfies $(S^{*}_{\preccurlyeq})$?

Satisfaction of this principle, of course, will require some worlds in $[A \wedge B]$ to be promoted in the ranking, notably in relation to certain worlds in $[A \wedge \neg B]$. But we must be cautious as to how this is to take place. Plausibly, for instance, it should not occur at the expense of the worlds in $[\neg A]$. In fact, it seems quite reasonable that, more broadly, the internal ordering of $[A \supset B]$ should be left untouched. We therefore suggest supplementing $(S^{*}_{\preccurlyeq})$ with the following 'retainment' principle, which ensures the preservation of these features of $\preccurlyeq_{\Psi * A \supset B}$:

$(Ret1^{*}_{\preccurlyeq})$ $\quad$ If $x, y \in [A \supset B]$, then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$ iff
$x \preccurlyeq_{\Psi * A \supset B} y$

Its syntactic counterpart is given as follows:

Proposition 3. *Given AGM, $(Ret1^{*}_{\preccurlyeq})$ is equivalent to*

$(Ret1)$ $\quad$ *If $A \supset B \in Cn(C)$, then $[(\Psi * A \Rightarrow B) * C] = [(\Psi * A \supset B) * C]$*

Given the DP postulates, this constraint obviously translates into one that connects $\preccurlyeq_{\Psi}$ and $\preccurlyeq_{\Psi * A \Rightarrow B}$ and whose syntactic counterpart is easily inferable from Proposition 3:

Proposition 4. *Given $(C1^{*}_{\preccurlyeq})$, $(Ret1^{*}_{\preccurlyeq})$ is equivalent to:*

$(Ret1'^{*}_{\preccurlyeq})$ $\quad$ *If $x, y \in [A \supset B]$, then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$ iff
$x \preccurlyeq_{\Psi} y$*

That conditional revision does not affect the internal ordering of $[A \wedge B]$ or of $[\neg A]$ is in fact required by a set of principles for conditional revision proposed in (Kern-Isberner 1999), to which we shall return later. Our principle adds to these the constraint that conditional revision by $A \Rightarrow B$ does not affect the relative standing of worlds in $[\neg A]$ in relation to worlds in $[A \wedge B]$. This further restriction yields the correct verdict in the following scenario:

Example 1. *My friend and I have taken our preschoolers Akira and Bashir on holiday. They slept in bunkbeds last night. Since both beds were unmade by the morning, I initially believe that they did not choose to sleep in the same bed but suspend judgment as to which respective beds they did choose. Furthermore, in the event of coming to believe that they in fact did decide to share a bed, I would suspend judgment as to which bed they opted for. I then find out that, if Akira slept on top, then Bashir would have done so too (because he does not like people sleeping above him). What changes? Plausibly, my beliefs will change in the following respect: since I will still believe that they did not share a bed, I will now infer that Akira slept on the bottom bed and Bashir on the top. What of my conditional beliefs? Plausibly, we will have the following continuity: It will remain the*

case that, were to find out that they in fact decided to share a bed, I would suspend judgment as to which bed they chose.

Indeed, let A and B respectively stand for Akira and for Bashir's sleeping on the top bed and Ψ be my initial state. Assume for simplicity that the set of atomic propositions in L is $\{A, B\}$. Let $[\![A \wedge \neg B]\!] = \{x\}$, $[\![\neg A \wedge B]\!] = \{y\}$, $[\![A \wedge B]\!] = \{z\}$ and $[\![\neg A \wedge \neg B]\!] = \{w\}$. We then have $\preccurlyeq_\Psi$ plausibly given by $x \sim_\Psi y \prec_\Psi z \sim_\Psi w$. Since $z \sim_\Psi w$, our principle entails the plausible result that $z \sim_{\Psi * A \Rightarrow B} w$.

But unfortunately, $(\mathrm{S}^*_\preccurlyeq)$ and $(\mathrm{Ret1}^*_\preccurlyeq)$ are not generally jointly sufficient to have the TPO $\preccurlyeq_{\Psi * A \supset B}$ determine the TPO $\preccurlyeq_{\Psi * A \Rightarrow B}$. Our suggestion is to close the gap by means of distance minimisation. More specifically, we propose to consider the closest TPO that satisfies–or, in the event of a tie, some aggregation of the closest TPOs that satisfy–our two constraints.

In terms of measuring the distance between TPOs, a natural choice is the so-called Kemeny distance d_K:

Definition 2. $d_K(\preccurlyeq, \preccurlyeq') := |(\preccurlyeq - \preccurlyeq') \cup (\preccurlyeq' - \preccurlyeq)|.$

Informally, $d_K(\preccurlyeq, \preccurlyeq')$ returns the number of disagreements over relations of weak preference between the two orderings, returning the number of pairs that are in $\preccurlyeq$ but not in $\preccurlyeq'$ and vice versa. This measure is standard fare in the social choice literature. It was introduced there in (Kemeny 1959) and received an axiomatisation in terms of a set of prima facie attractive properties in (Kemeny and Snell 1962).

In the section that follows we shall show that there exists a unique d_K-closest TPO that meets the requirements $(\mathrm{S}^*_\preccurlyeq)$ and $(\mathrm{Ret1}^*_\preccurlyeq)$, which can be obtained from $\preccurlyeq_{\Psi * A \supset B}$ in a simple and familiar manner.

3.2 A Construction of the Posterior TPO

To outline our main result, we first need the following item of notation (see Figure 2 for illustration):

Definition 3. *For any sentence $A \in L$ and TPO $\preccurlyeq$, we denote by $D(\preccurlyeq, A)$ the* down-set *of the members of $\min(\preccurlyeq, [\![A]\!])$. It is given by $D(\preccurlyeq, A) := \{x \mid x \preccurlyeq z$, for some $z \in \min(\preccurlyeq, [\![A]\!])\}$.*

With this in hand, we propose:

Definition 4. *Let $*$ be a function from $\mathbb{S} \times L$ to $\mathbb{S}$. Then we denote by $\circledast$ an arbitrary extension of $*$ to the domain $\mathbb{S} \times L_c$,*
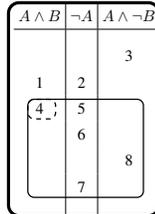

Figure 2: Down-set $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$ of the members of $\min(\preccurlyeq_{\Psi * A \supset B}, [\![A \wedge B]\!])$. The set $\min(\preccurlyeq_{\Psi * A \supset B}, [\![A \wedge B]\!])$, which here is a singleton, is marked by a dashed box. $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$ is marked by a solid box.
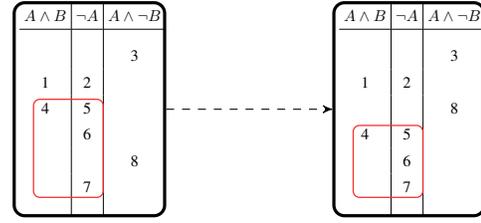

Figure 3: Relation between $\preccurlyeq_{\Psi * A \supset B}$ (depicted on the left) and $\preccurlyeq_{\Psi \circledast A \Rightarrow B}$ (depicted on the right). $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ is marked by a box.

such that $\preccurlyeq_{\Psi \circledast A \Rightarrow B}$ is given by the lexicographic revision of $\preccurlyeq_{\Psi * A \supset B}$ by $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$.[1]

The operator so-defined is illustrated in Figure 3, which depicts the resulting relation between $\preccurlyeq_{\Psi * A \supset B}$ and $\preccurlyeq_{\Psi \circledast A \Rightarrow B}$. Interestingly, in the special case of a Ramsey Test conditional with a tautologous antecedent, this transformation of $\preccurlyeq_{\Psi * A \supset B}$ amounts to its natural revision by the consequent. We propose to identify $\preccurlyeq_{\Psi * A \Rightarrow B}$ with $\preccurlyeq_{\Psi \circledast A \Rightarrow B}$. We do so on the basis of our main technical result, which is:

Theorem 1. *The unique TPO that minimises the distance d_K to $\preccurlyeq_{\Psi * A \supset B}$, given constraints $(\mathrm{S}^*_\preccurlyeq)$ and $(\mathrm{Ret1}^*_\preccurlyeq)$ is given by $\preccurlyeq_{\Psi \circledast A \Rightarrow B}$.*

As indicated above, short of endorsing lexicographic revision, which we do not want to do, the constraint of Success prevents us from having $[\Psi * A \Rightarrow B]_c = [\Psi * A \supset B]_c$ for all $A, B \in L$, $\Psi \in \mathcal{S}$. Having said that, a restricted version of this equality does hold for our proposal in the form of the following plausible 'Vacuity' postulate, which tells us that if revision by the material conditional leads to the conditional being accepted, then it is revision enough:

(V^*) If $A \Rightarrow B \in [\Psi * A \supset B]_c$, then
$[\Psi * A \Rightarrow B]_c = [\Psi * A \supset B]_c$

Furthermore, as a consequence of one of the results established in the proof of Theorem 1, we can also derive an interesting minimal change result with a more syntactic flavour:

Proposition 5. *Let $*$ be a function from $\mathbb{S} \times L$ to $\mathbb{S}$ and $*'$ an extension of $*$ to the domain $\mathbb{S} \times L_c$, satisfying $(\mathrm{S}^*_\preccurlyeq)$ and $(\mathrm{Ret1}^*_\preccurlyeq)$. Then, if $[\Psi *' A \Rightarrow B]_c$ agrees with $[\Psi * A \supset B]_c$ on all conditionals with a given antecedent C, so does $[\Psi \circledast A \Rightarrow B]_c$*

3.3 Some General Features

We have seen that our proposal to handle conditional revision using distance minimisation under constraints yields a unique TPO that can be obtained via lexicographic revision of $\preccurlyeq_{\Psi * A \supset B}$ by a particular proposition. In this section, we discuss some of its general consequences, including three additional retainment principles that it implies.

It is easy to establish the following:

Proposition 6. *Let $*$ be a function from $\mathbb{S} \times L_c$ to $\mathbb{S}$. Then, if $* = \circledast$, then $*$ satisfies:*

---

[1] In case we identify states with TPOs, there will exist only one such extension.

(Ret2$^*_\preccurlyeq$) *If $x, y \in [\![A \wedge \neg B]\!]$, then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$ iff*
  $x \preccurlyeq_{\Psi * A \supset B} y$

(Ret3$^*_\preccurlyeq$) *If $x \in [\![A \supset B]\!]$, $y \in [\![A \wedge \neg B]\!]$, and*
  $x \prec_{\Psi * A \supset B} y$, *then* $x \prec_{\Psi * A \Rightarrow B} y$

(Ret4$^*_\preccurlyeq$) *If $x \in [\![A \supset B]\!]$, $y \in [\![A \wedge \neg B]\!]$, and*
  $x \preccurlyeq_{\Psi * A \supset B} y$, *then* $x \preccurlyeq_{\Psi * A \Rightarrow B} y^2$

The conjunction of (Ret1$^*_\preccurlyeq$), with these three principles simply tells us that the only admissible transformations, when moving from $\preccurlyeq_{\Psi * A \supset B}$ to $\preccurlyeq_{\Psi * A \Rightarrow B}$, involve a doxastic 'demotion' of worlds in $[\![A \wedge \neg B]\!]$ in relation to worlds in $[\![A \supset B]\!]$, raising, in the ordering, the position of the former in relation to the latter. They have a similar flavour to that of the DP postulates, which tell us that that the only admissible transformations, when moving from $\preccurlyeq_\Psi$ to $\preccurlyeq_{\Psi * A}$, involve a demotion of worlds in $[\![\neg A]\!]$ in relation to worlds in $[\![A]\!]$.

We note the immediate implications of these principles, in the presence of the DP postulates:

Proposition 7. *Given* (C1$^*_\preccurlyeq$)–(C4$^*_\preccurlyeq$), (Ret2$^*_\preccurlyeq$) *holds iff:*

(Ret2$'^*_\preccurlyeq$) *If $x, y \in [\![A \wedge \neg B]\!]$, $x \preccurlyeq_{\Psi * A \Rightarrow B} y$ iff $x \preccurlyeq_\Psi y$*

and (Ret3$^*_\preccurlyeq$) *and* (Ret4$^*_\preccurlyeq$) *respectively entail:*

(Ret3$'^*_\preccurlyeq$) *If $x \in [\![A \supset B]\!]$, $y \in [\![A \wedge \neg B]\!]$, and $x \prec_\Psi y$,*
  then $x \prec_{\Psi * A \Rightarrow B} y$

(Ret4$'^*_\preccurlyeq$) *If $x \in [\![A \supset B]\!]$, $y \in [\![A \wedge \neg B]\!]$, and $x \preccurlyeq_\Psi y$,*
  then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$

but the converse entailments do not hold.

The syntactic counterparts of (Ret2$^*_\preccurlyeq$)–(Ret4$^*_\preccurlyeq$) are given in the following proposition, with the counterparts of (Ret2$'^*_\preccurlyeq$)–(Ret4$'^*_\preccurlyeq$) being easily inferable from these:

Proposition 8. *Given AGM,* (Ret2$^*_\preccurlyeq$)–(Ret4$^*_\preccurlyeq$) *are respectively equivalent to*

(Ret2) *If $A \wedge \neg B \in \mathrm{Cn}(C)$, then $[(\Psi * A \Rightarrow B) * C] = [(\Psi * A \supset B) * C]$*

(Ret3) *If $A \supset B \in [(\Psi * A \supset B) * C]$, then $A \supset B \in [(\Psi * A \Rightarrow B) * C]$*

(Ret4) *If $A \wedge \neg B \notin [(\Psi * A \supset B) * C]$, then $A \wedge \neg B \notin [(\Psi * A \Rightarrow B) * C]$*

In introducing (Ret1) above, we noted that, given (C1$^*_\preccurlyeq$), it strengthens, in a plausible manner, part of a principle proposed in (Kern-Isberner 1999). It turns out that, in the presence of the full set of DP postulates, (Ret1)–(Ret4) enable us to recover the trio of principles proposed by Kern-Isberner. These "KI postulates", originally named "(CR5)" to "(CR7)", are given semantically by:

(KI1$^*_\preccurlyeq$) *If $x, y \in [\![A \wedge B]\!]$, $x, y \in [\![\neg A]\!]$*
  or $x, y \in [\![A \wedge \neg B]\!]$, then $x \preccurlyeq_\Psi y$ iff
  $x \preccurlyeq_{\Psi * A \Rightarrow B} y$

(KI2$^*_\preccurlyeq$) *If $x \in [\![A \wedge B]\!]$, $y \in [\![A \wedge \neg B]\!]$ and $x \prec_\Psi y$,*
  then $x \prec_{\Psi * A \Rightarrow B} y$

(KI3$^*_\preccurlyeq$) *If $x \in [\![A \wedge B]\!]$, $y \in [\![A \wedge \neg B]\!]$ and $x \preccurlyeq_\Psi y$,*
  then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$

We can see that (KI1$^*_\preccurlyeq$) follows from (Ret1$^*_\preccurlyeq$) and (Ret2$^*_\preccurlyeq$), given (C1$^*_\preccurlyeq$) and (C2$^*_\preccurlyeq$). (KI2$^*_\preccurlyeq$) follows from the conjunction of (Ret3$^*_\preccurlyeq$) and (C3$^*_\preccurlyeq$), while (KI3$^*_\preccurlyeq$) follows from the conjunction of (Ret4$^*_\preccurlyeq$) and (C4$^*_\preccurlyeq$).[3]

In view of Theorem 1 and Proposition 6, it follows that, if a revision operator $*$ satisfies (C1$^*_\preccurlyeq$) to (C4$^*_\preccurlyeq$), then the conditional revision $\circledast$ operator that extends it in the manner described in Definition 4 satisfies (KI1$^*_\preccurlyeq$) to (KI3$^*_\preccurlyeq$).

(Ret3$^*_\preccurlyeq$) and (Ret4$^*_\preccurlyeq$) tell us that conditional revision by $A \Rightarrow B$ preserves any 'good news' for worlds in $[\![A \supset B]\!]$, compared to worlds in $[\![A \wedge \neg B]\!]$, that revision by $A \supset B$ would bring. Given (C3$^*_\preccurlyeq$) and (C4$^*_\preccurlyeq$), they notably add to (KI2$^*_\preccurlyeq$) and (KI3$^*_\preccurlyeq$) the idea that worlds in $[\![\neg A]\!]$ should not be demoted with respect to worlds in $[\![A \wedge \neg B]\!]$ in moving from $\preccurlyeq_\Psi$ to $\preccurlyeq_{\Psi * A \Rightarrow B}$. The appeal of this constraint is highlighted in the following case:

Example 2. *I am due to visit my hometown and would like to catch up with my friends Alex and Ben. Unfortunately, both of them moved away years ago and I doubt that I will see either. If I were to find out of either of them that he was going to be around, I would still believe that the other was not. Furthermore if I were to find out that exactly one of them would be back, I would not be able to guess which one of the two that would be. A friend now tells me that if Alex will be in town, then so will Ben. Very clearly, it should not be the case that, as a result of this new information, I would now take Alex to be a more plausible candidate for being the only one of my two friends that I will see (quite the contrary).*

Let A and B respectively stand for Alex and for Ben's being back in town and Ψ be my initial state. Assume for simplicity that the set of atomic propositions in L is simply $\{A, B\}$. Let $[\![A \wedge \neg B]\!] = \{x\}$, $[\![\neg A \wedge B]\!] = \{y\}$, $[\![A \wedge B]\!] = \{z\}$ and $[\![\neg A \wedge \neg B]\!] = \{w\}$. We then have $\preccurlyeq_\Psi$ plausibly given by $w \prec_\Psi x \sim_\Psi y \prec_\Psi z$. Since $y \preccurlyeq_\Psi x$, our principle entails that $y \preccurlyeq_{\Psi * A \Rightarrow B} x$, as it intuitively should be.

Aside from entailing the three further retainment principles that we have discussed, we also note that our postulates have the happy consequence of securing the following 'Doxastic Equivalence' principle, according to which conditional revisions are indistinguishable from revisions by material conditionals at the level of belief sets:

(DE$^*_\preccurlyeq$) $\min(\preccurlyeq_{\Psi * A \Rightarrow B}, W) = \min(\preccurlyeq_{\Psi * A \supset B}, W)$

More precisely, it is easy to show that:

---

[2]It turns out that this set of principles is sufficient to fully characterise our proposal when supplemented with (S$^*_\preccurlyeq$), (Ret1$^*_\preccurlyeq$) and the two following further postulates:

(Ret5$^*_\preccurlyeq$) If $x \in [\![A \wedge \neg B]\!]$, $y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$,
  $y \in [\![A \supset B]\!]$, and $x \prec_{\Psi * A \supset B} y$, then $x \prec_{\Psi * A \Rightarrow B} y$

(Ret6$^*_\preccurlyeq$) If $x \in [\![A \wedge \neg B]\!]$, $y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$,
  $y \in [\![A \supset B]\!]$, and $x \preccurlyeq_{\Psi * A \supset B} y$, then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$

These, however, are not immediately as clearly interpretable.

[3]We note that the KI postulates, in turn, subsume the DP postulates, which correspond to the special cases in which $A = \top$. Indeed, (KI1$^*_\preccurlyeq$) yields the conjunction of (C1$^*_\preccurlyeq$) and (C2$^*_\preccurlyeq$), while (KI2$^*_\preccurlyeq$) and (KI3$^*_\preccurlyeq$) give us (C3$^*_\preccurlyeq$) and (C4$^*_\preccurlyeq$), respectively.

Proposition 9. $(S^*_\preccurlyeq)$, $(\text{Ret}1^*_\preccurlyeq)$, $(\text{Ret}3^*_\preccurlyeq)$ *and* $(\text{Ret}4^*_\preccurlyeq)$ *collectively entail* $(\text{DE}^*_\preccurlyeq)$.

3.4 Elementary Conditional Revision Operators

A few interesting observations can be made regarding the more specific case in which $*$ is an elementary operator (i.e. belongs to the set $\{*_N, *_R, *_L\}$), which we illustrate in Figure 4. Having said that, we have noted above our significant reservations about identifying rational revision with any of these operators. This section is therefore addressed to those who are rather more optimistic.

First, we note that, in two of the special cases of interest, one of our two steps becomes superfluous.

If $* = *_L$, then the second step of our procedure is redundant. Indeed, for any x such that $x \in [\![A \wedge B]\!]$ and any $y \in [\![A \wedge \neg B]\!]$, we have $x \prec_{\Psi *_L A \supset B} y$. Hence every world that is in $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ is already strictly more minimal, in $\preccurlyeq_{\Psi * A \supset B}$, than any world that is not.

If $* = *_N$, then the first step of our procedure plays no role: we would obtain the same result by simply directly applying the second transformation to the initial TPO. This is apparent from the fact that natural revision by $A \supset B$ leaves unaffected the respective internal orderings of $\overline{D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)}$, $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ and
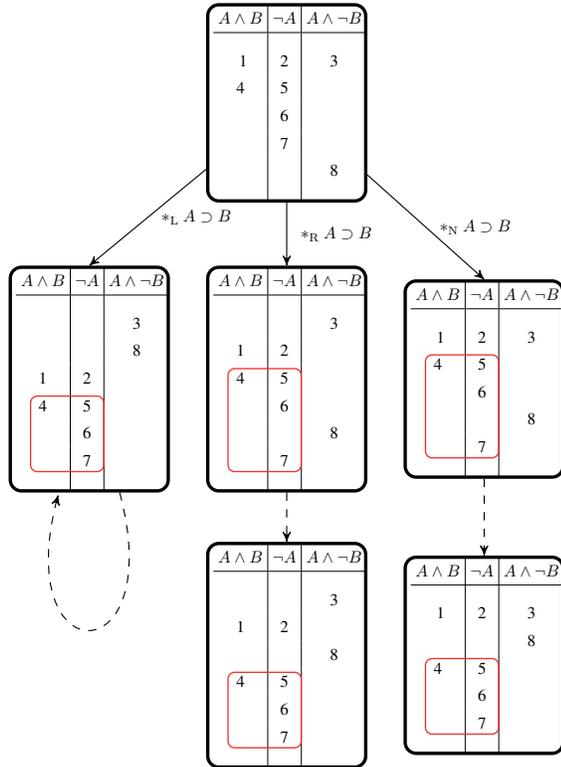

Figure 4: Two-step procedure for revision by $A \Rrightarrow B$ according to the respective proposed extensions of the elementary operators. $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ is marked by a box. The first step is denoted by a full arrow and the second by a dashed arrow.

$D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \wedge \neg B]\!]$, while the latter, on our proposal, jointly determine $\preccurlyeq_{\Psi * A \Rrightarrow B}$.

Secondly, in the case of elementary operators more generally, it can be shown that, on our proposal for $*A \Rrightarrow B$, the posterior internal ordering $\preccurlyeq_{\Psi * A \Rrightarrow B} \cap [\![A]\!]$ of the set of A-worlds is recovered by revising by B the restriction $\preccurlyeq \cap [\![A]\!]$ of the prior ordering to the A-worlds:[4]

Proposition 10. *If $*$ is an elementary revision operator, then $*$ and $\circledast$ satisfy:* $\preccurlyeq_{\Psi \circledast A \Rrightarrow B} \cap [\![A]\!] = (\preccurlyeq_\Psi \cap [\![A]\!]) * B$.

In other words: if one disregards the worlds in which the antecedent is false, the proposed transformation amounts to revision by the consequent.

Finally, in (Chandler and Booth 2019, Theorem 4), it was noted that there is an interesting connection between natural revision and the rational closure operator C_{rat} (Lehmann and Magidor 1992, Defs 20 and 21), which minimally extends any consistent set of conditionals to a set of conditionals corresponding to a rational consequence relation. This connection was that, if $\neg A \notin [\Psi]_c$, then $[\Psi *_N A]_c = C_{\text{rat}}([\Psi]_c \cup \{A\})$. This connection deepens on the proposed extension of natural revision to the conditional case. The proof of Chandler & Booth's theorem can be built upon to establish the following non-trivial result:

Proposition 11. *If $* = *_N$, then, if $A \Rrightarrow \neg B \notin [\Psi]_c$, then* $[\Psi \circledast A \Rrightarrow B]_c = C_{\text{rat}}([\Psi]_c \cup \{A \Rrightarrow B\})$

4 Related Research

We have already presented Kern-Isberner's trio of postulates for conditional revision and briefly discussed (and rejected) Nayak *et al*'s suggestion to treat conditional revision as lexicographic revision by a material conditional. In this section we turn to two further proposals that have been made in the literature and briefly compare them to ours. As we shall see, these both commit to identifying $*$ with $*_N$–which we have argued is undesirable–and exhibit further shortcomings.[5]

4.1 Hansson

(Hansson 1992) also takes a distance based approach, albeit an unconstrained one. He proposes to use the operator $*_H$:

Definition 5. $[\Psi *_H A \Rrightarrow B]_c := \bigcap [\Theta_i]_c$, *such that the Θ_i minimise the distance to Ψ, subject to the constraint that $A \Rrightarrow B \in [\Theta_i]_c$.*

The fate of this suggestion, of course, hinges on (i) one's view of the nature of states and (ii) the distance metric used. But if one equates states with TPOs and measures distance by means of d_K, then, first of all, rational revision coincides with natural revision: $\Psi * \top \Rrightarrow B = \Psi *_N B$. Indeed:

Proposition 12. *Let $\star$ be a revision operator that satisfies AGM. Then, if $\preccurlyeq_{\Psi *_N A} \neq \preccurlyeq_{\Psi \star A}$, then $d_K(\preccurlyeq_{\Psi *_N A}, \preccurlyeq_\Psi) < d_K(\preccurlyeq_{\Psi \star A}, \preccurlyeq_\Psi)$.*

---

[4]If, that is, we extend in the obvious manner the domain of $*$ to cover any TPO over some *subset* of W.

[5]We have left for a future occasion the comparison of our approach with the somewhat complex "c-revision" framework of (Kern-Isberner 2004), defined in terms of transformations of "conditional valuation functions", which include probability, possibility and ranking functions as special cases.

This, we take, is already not an appealing feature. Furthermore, Hansson's use of the intersection of a set of rational conditional belief sets should raise concerns, since it is well known that such intersections can fail to be rational. As it turns out, this worry is substantiated, and his suggestion is in fact inconsistent with at least one of the AGM postulates:

Proposition 13. *The operator* $*_{\text{H}}$ *does not satisfy*

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

An alternative way of aggregating the closest TPOs, which would guarantee an AGM-compliant output, would be to make use of an extension to the n-ary case of the binary TPO aggregation operator $\oplus_{\text{STQ}}$ of (Booth and Chandler 2019). We leave the study of this option to those who are more enthusiastic about the prospects of natural revision.

4.2 Boutilier & Goldszmidt

(Boutilier and Goldszmidt 1993) offer an alternative extension of $*_{\text{N}}$, which makes use of two further standard belief change operators: (i) the *contraction* operator $\div$, which returns the posterior state $\Psi \div A$ that results from an adjustment of Ψ to accommodate the retraction of A and (ii) the *expansion* operator $+$, which is similar to revision, save that consistency of the resulting beliefs needn't be ensured.

Like ours, their proposal involves a two-stage process, this time involving a first step of contraction by $A \Rightarrow \neg B$, then a step of expansion by $A \Rightarrow B$. In the case in which $A \Rightarrow B$ is not initially accepted, the contraction step involves moving the minimal $[\![A \wedge B]\!]$ worlds down to the rank r in which the minimal $[\![A \wedge \neg B]\!]$ worlds sit. The expansion step then has these minimal $[\![A \wedge \neg B]\!]$ worlds move up to a position immediately above r, while preserving their relations with any worlds that were strictly above or below them. Formally:

Definition 6. *The* Boutilier-Goldszmidt contraction operator $\div_{\text{BG}}$ *is such that*

    (1) If $x, y \notin \min(\preccurlyeq_{\Psi}, [\![A \wedge B]\!])$, then $x \preccurlyeq_{\Psi \div_{\text{BG}} A \Rightarrow \neg B} y$ iff $x \preccurlyeq_{\Psi} y$, and

    (2) If $x \in \min(\preccurlyeq_{\Psi}, [\![A \wedge B]\!])$, then

        (a) $x \preccurlyeq_{\Psi \div_{\text{BG}} A \Rightarrow \neg B} y$ iff $z \preccurlyeq_{\Psi} y$, for some $z \in \min(\preccurlyeq_{\Psi}, [\![A]\!])$, and

        (b) $y \preccurlyeq_{\Psi \div_{\text{BG}} A \Rightarrow \neg B} x$ iff $y \preccurlyeq_{\Psi} z$, for some $z \in \min(\preccurlyeq_{\Psi}, [\![A]\!])$

Definition 7. *The* Boutilier-Goldszmidt expansion operator $+_{\text{BG}}$ *is such that*

    (1) If $x \notin \min(\preccurlyeq_{\Psi}, [\![A \wedge \neg B]\!])$, then $x \preccurlyeq_{\Psi +_{\text{BG}} A \Rightarrow B} y$ iff $x \preccurlyeq_{\Psi} y$, and

    (2) If $x \in \min(\preccurlyeq_{\Psi}, [\![A \wedge \neg B]\!])$, then

        (a) if $y \in \min(\preccurlyeq_{\Psi}, [\![A \wedge \neg B]\!])$, then $y \preccurlyeq_{\Psi +_{\text{BG}} A \Rightarrow B} x$, and

        (b) if $y \notin \min(\preccurlyeq_{\Psi}, [\![A \wedge \neg B]\!])$, then $x \preccurlyeq_{\Psi +_{\text{BG}} A \Rightarrow B} y$ iff $x \preccurlyeq_{\Psi} y$ and there is no $z \in \min(\preccurlyeq_{\Psi}, [\![A \wedge B]\!])$ such that $y \preccurlyeq_{\Psi} z$[6]

---

[6]This is not quite the original formulation, which, in view of the informal description provided by the authors, appears to include a number of typographical errors.
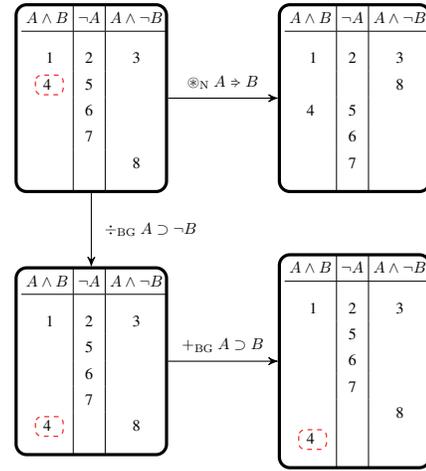

Figure 5: Illustrations of Boutilier & Goldszmidt's two-step proposal for extending $*_{\text{N}}$, contrasted with our own (denoted by $\circledast_{\text{N}}$). The set $\min(\preccurlyeq_{\Psi}, [\![A \wedge B]\!])$ is marked by a dashed box.

The corresponding revision operator is then defined as the composition of $\div_{\text{BG}} A \Rightarrow \neg B$ and $+_{\text{BG}} A \Rightarrow B$:

Definition 8. *The* Boutilier-Goldszmidt revision operator $*_{\text{BG}}$ *is given by* $\Psi *_{\text{BG}} A \Rightarrow B := (\Psi \div_{\text{BG}} A \Rightarrow \neg B) +_{\text{BG}} A \Rightarrow B$.

The operation $+_{\text{BG}} A \Rightarrow B$ bears some striking similarities to the second step in our construction of $* A \Rightarrow B$. In fact, it *coincides* with it in the kind of circumstances under which it is supposed to operate, i.e. on the heels of $\div_{\text{BG}} A \Rightarrow \neg B$.

Having said that, the introduction of the contraction step means that, overall, Boutilier & Goldszmidt's proposal quite clearly departs from the proposed extension of $*_{\text{N}}$ put forward in the previous section. In Figure 5, we see that it notably violates the requirement $(\text{Ret1}^{*}_{\preccurlyeq})$, according to which the ordering internal to $[\![A \supset B]\!]$ should be preserved (since, although $4, 7 \in [\![A \supset B]\!]$ and $7 \prec_{\Psi} 4$, we have $4 \prec_{\Psi *_{\text{BG}} A \Rightarrow B} 7$). This particular example is also an instance of the following feature of their revision operator:

Proposition 14. *If $A \in [\Psi]$, then $A \wedge B \in [\Psi *_{\text{BG}} A \Rightarrow B]$*

But this is a rather questionable property: it essentially precludes reasoning by Modus Tollens (aka denying the consequent). The following example highlights the counterintuitive character of this proscription:

Example 3. *I believe that the light in the bathroom next door is on (A), because the light switch in this room is down ($\neg B$). The owner of the house tells me that, contrary to what one might expect, when the bathroom light is on, that means that the switch in this room is* up*. So I revise by $A \Rightarrow B$. In doing so, I maintain my belief about the state of the switch ($\neg B$) and conclude that the bathroom light is off ($\neg A$).*

5 Concluding Comments

In what precedes we have offered a fresh approach to the problem of revision by conditionals, which imposes no constraints on the behaviour of the revision operator in relation

to non-conditional inputs. This independence was achieved by deriving the result of a revision by a conditional from the result of a revision by its material counterpart. This approach, we have argued, satisfies a number of attractive new properties and enjoys a number of distinctive advantages over existing alternative proposals.

Having said that, the scope of a number of results that we have established could perhaps be broadened.

Firstly, Proposition 10 shows that, at the level of the internal ordering of $[\![A]\!]$, conditional revision by $A \Rightarrow B$ operates like revision by B, for the special case of extensions of *elementary* revision operators. We do not know to what extent this generalises to a broader class of revision operators, such as the POI revision operators of (Booth and Chandler 2018; Booth and Chandler 2020).

Secondly, we establish in Proposition 11 that, if $* = *_N$ and $A \Rightarrow \neg B \notin [\Psi]_c$, then $[\Psi \circledast A \Rightarrow B]_c = C_{rat}([\Psi]_c \cup \{A \Rightarrow B\})$. This raises the following question: For $* = *_L$ or $* = *_R$, if $A \Rightarrow \neg B \notin [\Psi]_c$, do we have $[\Psi * A \Rightarrow B]_c = C([\Psi]_c \cup \{A \Rightarrow B\})$ for some suitable closure operator C?

Finally, at a number of points, we have made use of the distance metric d_K, noting that it was ubiquitous in the social choice literature. We are however aware of at least one alternative to this metric, proposed in (Duddy and Piggins 2012, Sec. 3.2), which coincides with d_K in the special case of linear orders. It would be interesting to assess the impact of this alternative choice on the proposal made here (another potential point of relevance concerns our assessment of Hansson's proposal, which also made use of d_K).

In addition to the question of the generalisability of certain results, we note that there is an extensive literature on a related issue for models of graded, rather than categorical, belief (esp. probabilistic models): how to update one's degrees of belief on information specifying a particular conditional degree of belief or presented in the form of a natural language indicative conditional. A natural approach here is to move to the posterior distribution that is "closest" to the prior one, on some appropriate distance measure, subject of the relevant informational constraint. However, an apparent issue with the use of the popular cross-entropy measure was presented in the classic Judy Benjamin example of (van Fraassen 1981), with similar observations being made in relation to two further measures in (van Fraassen, Hughes, and Harman 1986). For further discussions, see (Douven 2012), (Douven and Dietz 2011), (Douven and Romeijn 2011), (Eva, Hartmann, and Rad 2019), and (Grove and Halpern 1997). An examination of potential points of contact with the present work would be interesting to pursue.

6 Selected Proofs

Proposition 2. *If $*$ satisfies AGM, $(C1_{\preccurlyeq}^*)$, $(C2_{\preccurlyeq}^*)$, $(Eq_{\preccurlyeq}^*)$, and the principle according to which, for all $A, B \in L$ and $\Psi \in S$, $A \Rightarrow B \in [\Psi * A \supset B]_c$, then $* = *_L$.*

Proof: Given AGM, $(C1_{\preccurlyeq}^*)$, $(C2_{\preccurlyeq}^*)$, and $(Eq_{\preccurlyeq}^*)$, $*_L$ is characterised by the 'Recalcitrance' property

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

It will suffice to show that this property is entailed by:

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

So let $A \wedge B$ be consistent. Since $A \equiv (A \supset B) \supset A$, it suffices, by (Eq*) to show $A \in [(\Psi * (A \supset B) \supset A) * B]$. We know that for any AGM operator $*'$ and state Ψ', $A \in [(\Psi' *' B)]$ iff $A \in [(\Psi' *' A \supset B)]$. Hence it suffices to show $A \in [(\Psi * (A \supset B) \supset A) * A \supset B]$. Since $A \wedge B \equiv ((A \supset B) \supset A) \wedge (A \supset B)$ is consistent, we can apply (1) to recover the required result. □

Theorem 1. *The unique TPO that minimises the distance d_K to $\preccurlyeq_{\Psi * A \supset B}$, given constraints $(S_{\preccurlyeq}^*)$ and $(Ret1_{\preccurlyeq}^*)$ is given by $\preccurlyeq_{\Psi \circledast A \Rightarrow B}$.*

Proof: We prove the result in three lemmas. The first of these is the following:

Lemma 1. *For any function $*$ from $S \times L$ to S, $\circledast$ satisfies $(S_{\preccurlyeq}^*)$ and $(Ret1_{\preccurlyeq}^*)$.*

We note that, given the definition of $*_L$ in Definition 1, Definition 4 is equivalent to:

$x \preccurlyeq_{\Psi \circledast A \Rightarrow B} y$ iff
(1) $x \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ and $y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$, or
(2) $(x \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ iff $y \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!])$ and $x \preccurlyeq_{\Psi * A \supset B} y$

We now show that $(S_{\preccurlyeq}^*)$ is satisfied. By definition, $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \wedge B]\!] \neq \varnothing$. So let x be such that $x \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \wedge B]\!]$. Now consider $y \in [\![A \wedge \neg B]\!]$. By Definition 4, $x \prec_{\Psi \circledast A \Rightarrow B} y$, and hence $\min(\preccurlyeq_{\Psi \circledast A \Rightarrow B}, [\![A]\!]) \subseteq [\![B]\!]$, as required.

Regarding $(Ret1_{\preccurlyeq}^*)$: Assume $x, y \in [\![A \supset B]\!]$. For the left to right direction, assume that $x \preccurlyeq_{\Psi \circledast A \Rightarrow B} y$. From this one of either (1) or (2) holds. (2) immediately entails that $x \preccurlyeq_{\Psi * A \supset B} y$. (1) gives us: $x \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$ and $y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$. $x \preccurlyeq_{\Psi * A \supset B} y$ then follows from the definition of $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$. For the right to left direction, assume that $x \preccurlyeq_{\Psi * A \supset B} y$. If $x, y \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$ or $x, y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$, then we obtain $x \preccurlyeq_{\Psi \circledast A \Rightarrow B} y$ by (1). So assume that one of either x or y is in $D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$, while the other is not. From $x \preccurlyeq_{\Psi * A \supset B} y$, it must be the case that $x \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$ and $y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$. From (2), we then again recover $x \preccurlyeq_{\Psi \circledast A \Rightarrow B} y$.

Our next lemma states a general fact about lexicographic combinations of ordered pairs of TPOs, defined by

Definition 9. *The* lexicographic combination $\text{lex}(\preccurlyeq_1, \preccurlyeq_2)$ *of two TPOs $\preccurlyeq_1$ and $\preccurlyeq_2$ is given by the TPO $\preccurlyeq$ such that $x \preccurlyeq y$ iff (i) $x \prec_1 y$ or (ii) $x \sim_1 y$ and $x \prec_2 y$*

It is given as follows:

Lemma 2. *Let $\preccurlyeq_1$, $\preccurlyeq_2$ be two given TPOs and let $X(\preccurlyeq_2) = \{\preccurlyeq | \preccurlyeq$ is a TPO s.t. $\preccurlyeq \subseteq \preccurlyeq_2\}$. Then the TPO in $X(\preccurlyeq_2)$ that minimises the distance d_K to $\preccurlyeq_1$ is $\text{lex}(\preccurlyeq_1, \preccurlyeq_2)$.*

Let $\preccurlyeq' = \text{lex}(\preccurlyeq_1, \preccurlyeq_2)$. First we need to check that $\preccurlyeq' \in X(\preccurlyeq_2)$, i.e. $\preccurlyeq' \subseteq \preccurlyeq_2$. But this is clear from Definition 9. It remains to be shown that $d_K(\preccurlyeq_1, \preccurlyeq') \leq d_K(\preccurlyeq_1, \preccurlyeq'')$ for all $\preccurlyeq'' \in X(\preccurlyeq_2)$. To see this, we first reformulate d_K.

Definition 10. *A* hard conflict *between* $\preccurlyeq_1, \preccurlyeq'$ *is a 2-element set* $\{x, y\}$ *s.t.* $x \prec' y$ *and* $y \prec_1 x$. *Let* $\mathrm{Hard}(\preccurlyeq_1, \preccurlyeq')$ *denote the set of such hard conflicts.*

A soft conflict *between* $\preccurlyeq_1, \preccurlyeq'$ *is a 2-element set* $\{x, y\}$ *s.t. either (i)* $x \prec' y$ *and* $x \sim_1 y$ *or (ii)* $x \prec_1 y$ *and* $x \sim' y$. *Let* $\mathrm{Soft}(\preccurlyeq_1, \preccurlyeq')$ *denote the set of such soft conflicts.*

So $d_K(\preccurlyeq_1, \preccurlyeq') = 2 \times |\mathrm{Hard}(\preccurlyeq_1, \preccurlyeq')| + |\mathrm{Soft}(\preccurlyeq_1, \preccurlyeq')|$ and similarly for $d_K(\preccurlyeq_1, \preccurlyeq'')$. Hence, to show that $d_K(\preccurlyeq_1, \preccurlyeq') < d_K(\preccurlyeq_1, \preccurlyeq'')$ when $\preccurlyeq' \neq \preccurlyeq''$, it suffices to prove

(1) $\mathrm{Hard}(\preccurlyeq_1, \preccurlyeq') \subseteq \mathrm{Hard}(\preccurlyeq_1, \preccurlyeq'')$

(2) $\mathrm{Soft}(\preccurlyeq_1, \preccurlyeq') \subseteq \mathrm{Soft}(\preccurlyeq_1, \preccurlyeq'')$

Regarding (1): Let $\{x, y\} \in \mathrm{Hard}(\preccurlyeq_1, \preccurlyeq')$, i.e. $x \prec_1 y$ and $y \prec' x$. We must show $y \prec'' x$. By definition of $\preccurlyeq' = \mathrm{lex}(\preccurlyeq_1, \preccurlyeq_2)$, we have, from $y \prec' x$, (i) $y \prec_2 x$ or (ii) $y \sim_2 x$ and $y \prec_1 x$. We cannot have $y \prec_1 x$, since we already have $x \prec_1 y$. So $y \prec_2 x$. Hence, since $\preccurlyeq'' \in X(\preccurlyeq_2)$, i.e. $\preccurlyeq'' \subseteq \preccurlyeq_2$, we have $y \prec' x$ as well, as required.

Regarding (2): Let $\{x, y\} \in \mathrm{Soft}(\preccurlyeq_1, \preccurlyeq')$. We then have two cases to consider:

- $x \sim_1 y$ and $x \prec' y$: From $x \prec' y$, we get either (i) $x \prec_2 y$ or (ii) $x \sim_2 y$ and $x \prec_1 y$. The latter cannot occur, since we assume $x \sim_1 y$. Hence $x \prec_2 y$. Since $\preccurlyeq'' \subseteq \preccurlyeq_2$, we also then have $x \prec'' y$, so $\{x, y\} \in \mathrm{Soft}(\preccurlyeq_1, \preccurlyeq'')$, as required.
- $x \sim' y$ and $x \prec_1 y$: Impossible, since $x \sim' y$ entails that both $x \sim_1 y$ and $x \sim_2 y$ but we assume $x \prec_1 y$.

We now show that:

Lemma 3. *For any* $\preccurlyeq'$ *satisfying* $(\mathrm{S}^*_{\preccurlyeq})$ *and* $(\mathrm{Ret1}^*_{\preccurlyeq})$, *we must have* $\preccurlyeq' \subseteq \preccurlyeq_D$, *where* $\preccurlyeq_D$ *is defined as follows:*

$x \preccurlyeq_D y$ *iff* $x \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ *or* $y \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$

Let $\preccurlyeq'$ satisfy $(\mathrm{S}^*_{\preccurlyeq})$ and $(\mathrm{Ret1}^*_{\preccurlyeq})$. Suppose $y \preccurlyeq_D x$. We must show that $y \prec' x$. From $y \preccurlyeq_D x$, by definition of $\preccurlyeq_D$, we have $x \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$ and $y \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$. From $y \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$, we have $y \preccurlyeq_{\Psi * A \supset B} z$, where $z \in \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![A \wedge B]\!])$. We now consider two cases:

- $x \in [\![A \supset B]\!]$: Then, since $x \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B) \cap [\![A \supset B]\!]$, we have $x \notin D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$ and so, form this and $y \in D(\preccurlyeq_{\Psi * A \supset B}, A \wedge B)$, we get $y \preccurlyeq_{\Psi * A \supset B} x$. Since $x, y \in [\![A \supset B]\!]$, we get from this $y \prec' x$ by $(\mathrm{Ret1}^*_{\preccurlyeq})$, as required.
- $x \in [\![A \wedge \neg B]\!]$: By $(\mathrm{S}^*_{\preccurlyeq})$, we know that $u \prec' x$ for some $u \in [\![A \wedge B]\!]$. By the minimality of z, we know $z \preccurlyeq_{\Psi * A \supset B} u$. Hence $y \preccurlyeq_{\Psi * A \supset B} u$. Then, by $(\mathrm{Ret1}^*_{\preccurlyeq})$, we recover $y \prec' u$ and so $y \prec' x$, as required.

Putting together Lemmas 1, 2 and 3 yields the proof of the theorem. Lemma 1 tells us that $\preccurlyeq_{\Psi \circledast A \divideontimes B}$ satisfies $(\mathrm{S}^*_{\preccurlyeq})$ and $(\mathrm{Ret1}^*_{\preccurlyeq})$. Next, note that Definition 4 can be equivalently presented in terms of a lexicographic combination, so that $\preccurlyeq_{\Psi \circledast A \divideontimes B} = \mathrm{lex}(\preccurlyeq_D, \preccurlyeq_{\Psi * A \supset B})$.

In view of this, we can see that, by Lemma 2, $\preccurlyeq_{\Psi \circledast A \divideontimes B}$ minimises d_K to $\preccurlyeq_{\Psi * A \supset B}$ among all TPOs $\preccurlyeq$ s.t. $\preccurlyeq \subseteq \preccurlyeq_D$.

Finally, since all $\preccurlyeq$ that satisfy $(\mathrm{S}^*_{\preccurlyeq})$ and $(\mathrm{Ret1}^*_{\preccurlyeq})$ are such that $\preccurlyeq \subseteq \preccurlyeq_D$, by Lemma 3, $\preccurlyeq_{\Psi \circledast A \divideontimes B}$ must also minimise d_K among all TPOs satisfying $(\mathrm{S}^*_{\preccurlyeq})$ and $(\mathrm{Ret1}^*_{\preccurlyeq})$. $\qquad\square$

Proposition 5. *Let* $*$ *be a function from* $\mathbb{S} \times L$ *to* $\mathbb{S}$ *and* $*'$ *an extension of* $*$ *to the domain* $\mathbb{S} \times L_c$, *satisfying* $(\mathrm{S}^*_{\preccurlyeq})$ *and* $(\mathrm{Ret1}^*_{\preccurlyeq})$. *Then, if* $[\Psi *' A \divideontimes B]_c$ *agrees with* $[\Psi * A \supset B]_c$ *on all conditionals with a given antecedent* C, *so does* $[\Psi \circledast A \divideontimes B]_c$

Proof: We will prove the equivalent statement: if $[(\Psi \circledast A \divideontimes B) \circledast C] \neq [(\Psi \circledast A \supset B) \circledast C]$, then $[(\Psi *' A \divideontimes B) *' C] \neq [(\Psi *' A \supset B) *' C]$.

Suppose $[(\Psi \circledast A \divideontimes B) \circledast C] \neq [(\Psi \circledast A \supset B) \circledast C]$. Then $\mathrm{min}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, [\![C]\!]) \neq \mathrm{min}(\preccurlyeq_{\Psi \circledast A \supset B}, [\![C]\!])$. So we have two cases to consider, corresponding to the failures of each direction of subset inclusion.

Assume $\mathrm{min}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, [\![C]\!]) \nsubseteq \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!])$. Let $x \in \mathrm{min}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, [\![C]\!]) - \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!])$. Let $y \in \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!])$. Then $x \preccurlyeq_{\Psi \circledast A \divideontimes B} y$ (by minimality of x) and $y \preccurlyeq_{\Psi * A \supset B} x$. Since $x \preccurlyeq_{\Psi \circledast A \divideontimes B} y$, we know that either $x \sim_{\Psi \circledast A \divideontimes B} y$ or $x \prec_{\Psi \circledast A \divideontimes B} y$. By our construction of $\circledast$, it is not possible to have both $x \sim_{\Psi \circledast A \divideontimes B} y$ and $y \prec_{\Psi * A \supset B} x$. So we must have $x \prec_{\Psi \circledast A \divideontimes B} y$ and hence $\{x, y\} \in \mathrm{Hard}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$. Now, since we assume $*'$ to satisfy $(\mathrm{S}^*_{\preccurlyeq})$ and $(\mathrm{Ret1}^*_{\preccurlyeq})$, we know from our proof of Theorem 1 that $\mathrm{Hard}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B}) \subseteq \mathrm{Hard}(\preccurlyeq_{\Psi *' A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$. So $\{x, y\} \in \mathrm{Hard}(\preccurlyeq_{\Psi *' A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$ and hence $x \prec_{\Psi *' A \divideontimes B} y$. So we have $y \in \mathrm{min}(\preccurlyeq_{\Psi *' A \supset B}, [\![C]\!]) = \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!])$ (since $*'$ extends $*$, i.e. agrees with it on unconditional revisions) but $y \notin \mathrm{min}(\preccurlyeq_{\Psi *' A \divideontimes B}, [\![C]\!])$. So $\mathrm{min}(\preccurlyeq_{\Psi *' A \supset B}, [\![C]\!]) \neq \mathrm{min}(\preccurlyeq_{\Psi *' A \divideontimes B}, [\![C]\!])$, i.e. $[(\Psi *' A \divideontimes B) *' C] \neq [(\Psi *' A \supset B) *' C]$, as required

Assume $\mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!]) \nsubseteq \mathrm{min}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, [\![C]\!])$. Let $x \in \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!]) - \mathrm{min}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, [\![C]\!])$. Let $y \in \mathrm{min}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, [\![C]\!])$. Then $y \prec_{\Psi \circledast A \divideontimes B} x$ and $x \preccurlyeq_{\Psi * A \supset B} y$, from which $x \prec_{\Psi * A \supset B} y$ or $x \sim_{\Psi * A \supset B} y$.

If $x \prec_{\Psi * A \supset B} y$, then $\{x, y\} \in \mathrm{Hard}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$. Since, as noted above, $\mathrm{Hard}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B}) \subseteq \mathrm{Hard}(\preccurlyeq_{\Psi *' A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$, we get $y \prec_{\Psi *' A \divideontimes B} x$ and so $x \in \mathrm{min}(\preccurlyeq_{\Psi *' A \supset B}, [\![C]\!]) - \mathrm{min}(\preccurlyeq_{\Psi *' A \divideontimes B}, [\![C]\!])$, i.e. $[(\Psi *' A \divideontimes B) *' C] \neq [(\Psi *' A \supset B) *' C]$, as required.

If $x \sim_{\Psi * A \supset B} y$, then $\{x, y\} \in \mathrm{Soft}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$. Since, $\mathrm{Soft}(\preccurlyeq_{\Psi \circledast A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B}) \subseteq \mathrm{Soft}(\preccurlyeq_{\Psi *' A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$ (again, see proof of Theorem 1), we have $\{x, y\} \in \mathrm{Soft}(\preccurlyeq_{\Psi *' A \divideontimes B}, \preccurlyeq_{\Psi * A \supset B})$, so either $x \prec_{\Psi *' A \divideontimes B} y$ or $y \prec_{\Psi *' A \divideontimes B} x$. In the latter case, we get $[(\Psi *' A \divideontimes B) *' C] \neq [(\Psi *' A \supset B) *' C]$ as above. In the former case, we deduce $y \notin \mathrm{min}(\preccurlyeq_{\Psi *' A \divideontimes B}, [\![C]\!])$. But from $x \sim_{\Psi * A \supset B} y$ and $x \in \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!])$ we must have $y \in \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!])$. Hence $y \in \mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!]) - \mathrm{min}(\preccurlyeq_{\Psi *' A \divideontimes B}, [\![C]\!])$, so $\mathrm{min}(\preccurlyeq_{\Psi * A \supset B}, [\![C]\!]) \neq \mathrm{min}(\preccurlyeq_{\Psi *' A \divideontimes B}, [\![C]\!])$, which gives again $[(\Psi *' A \divideontimes B) *' C] \neq [(\Psi *' A \supset B) *' C]$, as required. $\qquad\square$

Acknowledgements

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(02):510–530.

Booth, R., and Chandler, J. 2017. The irreducibility of iterated to single revision. *J. Philosophical Logic* 46(4):405–418.

Booth, R., and Chandler, J. 2018. On strengthening the logic of iterated belief revision: Proper ordinal interval operators. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018.*, 210–219.

Booth, R., and Chandler, J. 2019. From iterated revision to iterated contraction: Extending the Harper Identity. *Artif. Intell.* 277.

Booth, R., and Chandler, J. 2020. On strengthening the logic of iterated belief revision: Proper ordinal interval operators. *Artif. Intell.* 285.

Booth, R., and Meyer, T. 2006. Admissible and restrained revision. *Journal of Artificial Intelligence Research* 26(1):127–151.

Boutilier, C., and Goldszmidt, M. 1993. Revision by conditional beliefs. In *Proceedings of the 11th National Conference on Artificial Intelligence. Washington, DC, USA, July 11-15, 1993.*, 649–654.

Boutilier, C. 1996. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic* 25(3):263–305.

Chandler, J., and Booth, R. 2019. Elementary iterated revision and the Levi Identity. In Blackburn, P.; Lorini, E.; and Guo, M., eds., *Logic, Rationality, and Interaction - 7th International Workshop, LORI 2019, Chongqing, China, October 18-21, 2019, Proceedings*, volume 11813 of *Lecture Notes in Computer Science*, 15–28. Springer.

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

Douven, I., and Dietz, R. 2011. A puzzle about Stalnaker's Hypothesis. *Topoi* 30(1):31–37.

Douven, I., and Romeijn, J. 2011. A new resolution of the Judy Benjamin problem. *Mind* 120(479):637–670.

Douven, I. 2012. Learning conditional information. *Mind & Language* 27(3):239–263.

Duddy, C., and Piggins, A. 2012. A measure of distance between judgment sets. *Social Choice and Welfare* 39(4):855–867.

Eva, B.; Hartmann, S.; and Rad, S. R. 2019. Learning from Conditionals. *Mind* 129(514):461–508.

Grove, A. J., and Halpern, J. Y. 1997. Probability update: Conditioning vs. cross-entropy. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, UAI'97, 208–214. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

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

Hansson, S. O. 1992. In defense of the Ramsey Test. *Journal of Philosophy* 89(10):522.

Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.

Kemeny, J., and Snell, J. 1962. *Mathematical Models in the Social Sciences*. Blaisdell, New York. Reprinted by MIT Press, Cambridge, 1972.

Kemeny, J. G. 1959. Mathematics without numbers. *Daedalus* 88(4):577–591.

Kern-Isberner, G. 1999. Postulates for conditional belief revision. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999. 2 Volumes, 1450 pages*, 186–191.

Kern-Isberner, G. 2004. A thorough axiomatization of a principle of conditional preservation in belief revision. *Ann. Math. Artif. Intell.* 40(1-2):127–164.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial intelligence* 55(1):1–60.

Nayak, A. C.; Pagnucco, M.; Foo, N. Y.; and Peppas, P. 1996. Learning from conditionals: Judy Benjamin's other problems. In *12th European Conference on Artificial Intelligence, Budapest, Hungary, August 11-16, 1996, Proceedings*, 75–79.

Nayak, A. C.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *Artificial Intelligence* 146(2):193–228.

van Fraassen, B.; Hughes, R. I. G.; and Harman, G. 1986. A problem for relative information minimizers, continued. *British Journal for the Philosophy of Science* 37(n/a):453–475.

van Fraassen, B. C. 1981. A problem for relative information minimizers in probability kinematics. *British Journal for the Philosophy of Science* 32(4):375–379.