Preservation, Commutativity and Modus Ponens: Two Recent Triviality Results

Jake Chandler
LMU Munich
jake.chandler@cantab.net

In a recent pair of publications, Richard Bradley has offered two novel no-go theorems involving the principle of 'Preservation' for conditionals, which guarantees that one's prior conditional beliefs will exhibit a certain degree of inertia in the face of a change in one's non-conditional beliefs.

We first note that Bradley's original discussions of these results—in which he finds motivation for rejecting Preservation, first in a principle of 'Commutativity', then in a doxastic analogue of the rule of Modus Ponens—are problematic in a significant number of respects.

We then turn to a recent u-turn on his part, in which he winds up rescinding his commitment to Modus Ponens, on the grounds of a tension with the rule of Import-Export for conditionals. Here we offer an important positive contribution to the literature, settling the following crucial question that Bradley leaves unanswered: Assuming that one gives up on full-blown Modus Ponens on the grounds of its incompatibility with Import-Export, what weakened version of the principle should one be settling for instead? Our discussion of the issue turns out to unearth an interesting connection between epistemic undermining and the apparent failures of Modus Ponens in McGee's famous counterexamples.

1 Triviality and the Ramsey test

For over two and a half decades, the so-called Ramsey test, which cashes out an agent's acceptance of a conditional of the form 'If A then B' in terms of her commitment to come to believe that B upon revising her beliefs so as to accommodate A, has been at the center of a lively debate regarding the appropriate formulation of the rules of rational belief dynamics.

A large proportion of this debate has been carried out against the backdrop of the AGM theory of belief change.[1] In this framework, an agent's state of mind is represented by a so-called belief set, a consistent, deductively closed set of sentences K, drawn from a conditional propositional language L.[2] '$K * A$' is used to denote the posterior belief set resulting from the revision of the prior belief set K by A: the rational adjustment of K to accommodate the acquisition of the belief that A. In this context, the Ramsey test can thus be succinctly formulated as follows:

(RAM) $A > B \in K$ iff $B \in K * A$[3]

where > is a conditional connective, standardly intended to correspond to the natural language indicative conditional.

[1] There is of course also an extensive discussion of a closely related issue in the Bayesian literature, initiated by the work of Adams (1975), followed by that of Lewis (1976) and many others. See Hájek & Hall 1994 for an excellent summary.

[2] In the literature, K is perhaps now more standardly known as a belief 'state' and represented by an upper-case Greek letter (e.g. Ψ). The term belief 'set' and the letter K are typically reserved for its factual subset.

[3] As an anonymous referee points out, it has been noted that on the present interpretation of the revision operator *, RAM yields apparently counterintuitive results in the context of agents capable of autodoxastic reasoning (see Hajek and Chalmers 2007 for a presentation of the difficulties, and Leitgeb 2011 for a reply). We shall however bracket this difficulty, since it is tangential to the particular issues at stake here.

A number of rationality constraints on the revision function have been offered, broadly guided by the idea that the adjustment of the prior belief set ought to be as conservative as possible. Among these, the so-called AGM postulates (see Alchourón *et al* 1985) provide fairly uncontroversial restrictions on the 'factual'—i.e. non-conditional—composition of the posterior belief set.

In the mid-1980's the publication of a now famous impossibility result due to Gärdenfors (1986), brought to general attention the difficulties that the Ramsey test gives rise to when the scope of the AGM postulates is extended so as to constrain not only factual, but also *conditional*, posterior beliefs. A number of variants of Gärdenfors' theorem can now be found in the literature. All uncover a tension between RAM, the unrestricted versions of a small subset of the AGM postulates and a fairly strong monotonicity condition on the consequence relation in L. In what follows, we shall sketch out a particularly perspicuous version, first presented in Rott 1989.

The AGM postulates involved include first of all the so-called Inclusion condition. Inclusion states that the posterior beliefs resulting from a revision do not go beyond the joint logical consequences of the prior belief set and the input to the revision:

(INC) If $B \in K * A$, then $B \in \mathrm{Cn}(K \cup \{A\})$

where Cn is a consequence function, retuning the set of consequences of any given subset Γ of L, and assumed to both be supraclassical—i.e. return at least the set of classical consequences of Γ—and satisfy the deduction theorem—i.e. be such that if $B \in \mathrm{Cn}(\Gamma \cup \{A\})$, then $A \supset B$ is in Γ.

Given closure of belief sets under logical consequence, the second constraint implicated, the Vacuity postulate, states that if the input to the revision is consistent with the prior belief state, then the joint logical consequences of the prior belief set and the input to the revision do not go beyond the posterior belief set:

(VAC) If $\neg A \notin K$ and $B \in \mathrm{Cn}(K \cup \{A\})$, then $B \in K * A$

The third constraint, the Success postulate, simply guarantees that any consistent input to revision is a member of the posterior belief set:

(SUC) If $\neg A \notin \mathrm{Cn}(\emptyset)$, then $A \in K * A$[4]

Finally, we have the aforementioned monotonicity condition, which ensures that the consequences of a given set of sentences are preserved when a member of the latter is replaced by a classically logically stronger counterpart:

(MON) If $B \in \mathrm{Cn}(\Gamma \cup \{A \vee C\})$, then $B \in \mathrm{Cn}(\Gamma \cup \{A\})$

It should be noted that all these principles can be weakened by restricting the membership of B to various subsets of L. We shall denote by L_0, the conditional-free fragment of L. L_{i+1} is then defined as the smallest language that includes all sentences in L_i alongside those conditionals of the form $A > B$, such that $A \in L_0$ and $B \in L_i$. Thus, in addition to factual sentences, L_1 also includes 'flat' conditionals, of the form '$A > B$', where 'A' and 'B' are both factual, and L_2 both flat and simple right-nested conditionals, of the form '$A > B$', where 'A' is factual and 'B' is a flat conditional. In what follows, we shall distinguish the various corresponding weakenings of the above principles, obtained by restricting B's membership to L_i, by means of the subscript i. Whilst the use of subscripting may, at first glance, strike the reader as somewhat cumbersome, the distinctions that it enables us to draw will prove to be of utmost importance. Indeed, many principles that are perfectly innocuous in their most restricted form turn out to be highly problematic when generalised.

[4] INC and VAC are also commonly referred to as K*3 and K*4, respectively. SUC corresponds to a slight weakening of what is known as K*2. The latter guarantees that A is in $K * A$ even in cases in which A is inconsistent. We assume here, however, that rational belief sets are consistent.

The general proof strategy involved in the impossibility results is the following. *Step 1*: Make a 'non-triviality' assumption about the set $\mathbb{K}$ of rationally permissible belief states, an assumption stating the existence of a particular kind of prior belief set K in $\mathbb{K}$. In the particular case of Rott's result, we have:

(NT) There exists a $K \in \mathbb{K}$ and three pairwise inconsistent sentences $A, B, C \in L_0$ such that $\neg A, \neg B, \neg C \notin K$[5]

Step 2: From this assumption, derive the result that there exist two particular ways of revising K, such that each of the resulting revised belief sets contains one of two conditionals with identical antecedents but jointly inconsistent consequents. Here, it is established that $\neg A > B \in K * A \vee B$ and $\neg A > C \in K * A \vee C$. *Step 3*: Show that, as a result of this, there exists a single revised belief set—in the present instance: $K * A$—that contains both of these conditionals. The existence of such a set, however, is prohibited by RAM, which entails the following principle of 'Conditional Contradiction', given closure and consistency of belief sets:

(CC) If $\neg C \in \mathrm{Cn}(B)$ and $A > B \in K$, then $A > C \notin K$

The derivation in Step 2 requires very little by way of assumptions: the right-to-left half of RAM_0, SUC, INC_0 and VAC_0.[6] None of these principles seem unduly strong.[7] Step 3 is considerably more controversial. It makes use of the far more substantial INC_1, VAC_1 and MON_1 to derive a restricted version of condition known as 'Monotonicity of Additions', guaranteeing that, in the circumstances, all L_1 members of $K * A \vee B$ or $K * A \vee C$, and in particular $A > B$ and $A > C$ respectively, will find their way into $K * A$:

(MA_1) If $\neg A \notin K$ and $C \in K * A \vee B$, then $C \in K * A$, where $A, B \in L_0$ and $C \in L_1$[8]

It is on MA_1 that the majority of contributors to the literature have been inclined to pin the blame, with Lindström and Rabinowicz (1995, p. 153) offering a convincing counterexample to the principle. Beyond this, it has remained somewhat unclear, in turn, which of INC_1, VAC_1 and MON_1 ought to be weakened and in what way.

In a noteworthy addition to the literature, however, Etlin (2009) has shown that if one grants the additional plausible principle according to which $\neg A > B$ and $\neg A > C$ are jointly inconsistent, then a retreat from VAC_1 to VAC_0 fails to get us out of trouble: either MON_1 or INC_1 must be relinquished. Indeed, he notes that the detour via MA_1 can be dispensed with: $\neg A > B$ and $\neg A > C$ needn't even make it into $K * A$ to create difficulties. On the aforementioned assumption that these conditionals are jointly inconsistent, trouble would ensue from the mere fact of their both winding up in $\mathrm{Cn}(K \cup \{A\})$.[9] And to derive this result from the fact that $\neg A > B \in K * A \vee B$ and $\neg A > C \in K * A \vee C$, we simply need INC_1 (to first get the conditionals into $\mathrm{Cn}(K \cup \{A \vee B\})$ and $\mathrm{Cn}(K \cup \{A \vee C\})$, respectively) and MON_1 (to then get them both into $\mathrm{Cn}(K \cup \{A\})$).

[5]Note that, since $\neg A, \neg B, \neg C \notin K$ and the sentences A, B and C are pairwise inconsistent, we also have $A, B, C \notin K$, by virtue of closure and consistency of belief sets.

[6]See Appendix, Observation 1.

[7]It should be noted that a number of people have considered the option of weakening RAM_0. Gärdenfors (1987) shows that trouble still ensues from a number of these proposals, including the ones proposed by Rott (1986). Lindström and Rabinowicz (1992) do offer a weakening that escapes Gärdenfors' later result, but it is one that imposes extremely stringent conditions on beliefs about conditionals. Rabinowicz (1996) has also cast aspersions on VAC_0, although, for reasons that space considerations preclude us from outlining here, we do not find his discussion entirely convincing.

[8]See Appendix, Observation 2.

[9]Indeed, K is assumed to be consistent and such that $\neg A \notin K$. $\mathrm{Cn}(K \cup \{A\})$ is therefore itself consistent.

Whilst Etlin's observation clearly alleviates the pressure placed on VAC by the Gärdenfors/Rott theorems, Bradley (2007; 2012) has recently noted, in a well-received pair of papers, that the principle is implicated in two further impossibility results. And these, he claims, *do* show that the principle is untenable in its unrestricted form.

Bradley's immediate target is a slight weakening of VAC, which goes under the name of 'Preservation':

(PRES) If $\neg A \notin K$ and $B \in K$, then $B \in K * A$

This principle equivalently states that if $B \in K$ but $B \notin K * A$, then $\neg A \in K$, that is, that one takes to be false any proposition whose truth one would take to undermine one's current beliefs.

The proofs essentially follow the three-step template outlined above, although Bradley does not assume that $\neg A > B \in K * A \vee B$ and $\neg A > C \in K * A \vee C$ *follow* from NT. Instead he makes an assumption NT* such that there exists a $K \in \mathbb{K}$ that both has this property and satisfies the conditions spelled out in NT:

(NT*) There exists a $K \in \mathbb{K}$ and three pairwise inconsistent sentences $A, B, C \in L_0$ such that $\neg A, \neg B, \neg C \notin K$ *and* (ii) both $\neg A > B \in K * A \vee B$ and $\neg A > C \in K * A \vee C$

This enables him to do away with INC_0, which, for reasons that we shall not discuss here (see Bradley 2012, pp. 153–54), he deems objectionable.[10]

In order to derive the result that $\neg A > B$ and $\neg A > C$ are in $K * A$, neither INC_1 nor MON_1 are appealed to. This important feature of the proofs is what allows Bradley to apply renewed pressure on VAC, via Preservation. In the first result, he makes use of the relatively modest PRES_1 and the following restricted version of a principle of 'Commutativity' which states that, under certain conditions, the order in which a pair of successive revisions take place does not make a difference to the contents of the resulting belief set:

(COM_1) If $\neg A, \neg C \notin K, \neg C \notin K*A$, and $\neg A \notin K*C$, then $B \in (K*A)*C$ iff $B \in K*A\&C$, where $B \in L_1$

In the second result, the work is done by the stronger PRES_2, alongside the following principle of closure of belief sets under Modus Ponens:

(MP_1) If $A, A > B \in K$, then $B \in K$, where $A \in L_0$ and $B \in L_1$

In both cases, Bradley initially pins the blame on Preservation, arguing for two successive weakenings of the principle.

Bradley's first result is the subject of section 2. We first outline a correction to his proof, noting that, contrary to what he claims, a weak version of INC, namely INC_0, *would* in fact appear to be required to derive the relevant result. We then suggest that COM_1 is excessively strong, noting that it is inconsistent with an uncontroversial weakening of a popular principle of 'iterated' belief revision.

In section 3, we then turn to Bradley's second result. As we have just mentioned, he initially intends to resolve the tension by weakening PRES_2. We point out in subsection 3.1, however, that the principle that he winds up recommending is both too strong and too weak. In a later paper he opts for rejecting MP_1, on the grounds of its incompatibility with the rule of Import-Export. He does not, however resolve the issue

[10] It should perhaps be noted that Bradley also makes use of a somewhat stronger *logical* principle of 'Conditional Contradiction' rather than a doxastic one, such as the one presented above (see Bradley 2007, p. 5). The idea is essentially to declare as jointly *inconsistent*, rather than simply jointly unbelievable, those conditionals with identical, consistent antecedents and jointly inconsistent consequents. Something similar to our principle CC is then recovered via the consistency requirement on belief sets.

of the extent to which it ought to be weakened. This is the focus of subsection 3.2, in which we offer the key positive contribution of the present paper. We note that a slightly restricted version of MP_1, that we derive from weak principles of belief dynamics, would provide a suitable replacement. This suggestion turns out to highlight an interesting potential connection between epistemic undermining and counterexamples to Modus Ponens.

In order to not overload the main body of the article, a number of proofs have been relegated to the appendix.

2 Triviality from Commutativity

As mentioned in the previous section, Bradley offers what he takes to be a proof that, in the presence of CC and NT*, PRES_1 is inconsistent with COM_1:

> [L]let $\{A, B, C\}$ be a set of mutually contradictory…sentences and suppose that an agent's initial epistemic state is given by a set K containing none of $A, \neg A, B, \neg B, C$ or $\neg C$. Suppose also that $\neg A > B$ belongs to $K * A \vee B$ and that $\neg A > C$ belongs to $K * A \vee C$. Now consider the sets $(K * A \vee B) * A$ and $(K * A \vee C) * A$, representing the epistemic states resulting from learning or supposing that A when in the states represented by $K * A \vee B$ and $K * A \vee C$. Since $\neg A$ does not belong to either $K * A \vee B$ or $K * A \vee C$, it follows from [PRES_1] that $\neg A > B$ belongs to $(K * A \vee B) * A$ and that $\neg A > C$ belongs to $(K * A \vee C) * A$ and from [COM_1] that [$\neg A > B$ belongs to $(K * A \vee B) * A$ iff it belongs to $K * A$ and $\neg A > C$ belongs to $(K * A \vee C) * A$ iff it belongs to $K * A$]. Hence both $\neg A > B$ and $\neg A > C$ belong to $K * A$, contrary to the hypothesis that it is not permissible to believe both that $\neg A > B$ and that $\neg A > C$. (Bradley 2007, p. 12)

The upshot of this, Bradley tells us, is that, assuming the legitimacy of CC and NT*, we must give up either COM_1 or PRES_1. He aims to keep the first, but only retain a substantially weakened version of the second. He recommends a retreat to

$$(\mathrm{PRES}_1^*) \quad \text{If } \neg A \notin K \text{ and } B > C \in K, \text{ then } B > C \in K * A, \text{ so long as } \neg B \notin \mathrm{Cn}(A)$$

which states that conditional beliefs are preserved under non-belief-contravening revisions so long as their antecedent is consistent with the revision-inducing input. Bradley remarks that, in the absence of any further background assumptions, a shift from PRES_1 to PRES_1^* allows one to steer clear of a conflict between NT* and COM_1: since $\neg\neg A \in \mathrm{Cn}(A)$, PRES_1^* is too weak to enable the derivation of $\neg A > B \in (K * A \vee B) * A$ and $\neg A > C \in (K * A \vee C) * A$ in the proof above.

Before turning to more substantial issues, we first note a lacuna in the proof sketched out above. Indeed, the claim that '$\neg A$ does not belong to either $K * A \vee B$ or $K * A \vee C$', which plays an essential role in the proof, remains unsubstantiated. And it would appear that to derive it, we require INC_0. Here is how the amended proof might go: Assume NT*. Assume for reductio that $\neg A \in K * A \vee B$. By INC_0, we then have $\neg A \in \mathrm{Cn}(K \cup \{A \vee B\})$. Since Cn is assumed to satisfy the deduction theorem, it follows that $(A \vee B) \supset \neg A \in K$. Since $\neg A \in \mathrm{Cn}((A \vee B) \supset \neg A)$, we then have, by closure of belief sets, $\neg A \in K$, in contradiction with our initial assumption. By parallel reasoning, we can establish that $\neg A \notin K * A \vee C$. QED.[11]

[11] In a second version of the proof, Bradley eschews the aforementioned claim altogether. Instead, he appeals to the observation that 'A does not contradict either $A \vee B$ or $A \vee C$' to derive $\neg A > B$ and $\neg A > C$'s memberships of $(K * A \vee B) * A$ and $(K * A \vee C) * A$ respectively (see Bradley 2012, p. 152). PRES, however, does not license this inference.

Let us now turn to the moral that Bradley would have us draw from the result, namely that COM_1 ought to be retained at the expense of PRES_1.

Now COM_0, the restriction of COM to factual sentences, is both rather weak and rather plausible. It is widely endorsed in the belief revision community, and follows from a subset of the AGM and so-called DP postulates.[12] COM_1, however, is substantially stronger. It is not just inconsistent with its being *mandatory* that both $\neg A > B$ and $\neg A > C$ are preserved in the transitions from $K * A \vee B$ to $(K * A \vee B) * A$ and from $K * A \vee C$ to $(K * A \vee C) * A$, respectively, as PRES_1 entails. As Bradley's result demonstrates, it is also inconsistent with its being merely *permissible*.

This strikes us as being too restrictive. To illustrate, consider an example discussed by Bradley himself. You have in front of you three urns, one of which contains a prize. You are initially agnostic as to the location of the prize. Scenario 1: You obtain information that suggests that the prize is either in urn A or in urn B and thus come to hold that if the prize is not in urn A, it is in urn B. Scenario 2: You obtain information that suggests that the prize is either in urn A or in urn C, and thus come to hold that if the prize is not in urn A, it is in urn C. Now, in either scenario, upon obtaining further information suggesting that the prize is in urn A, it would not seem intrinsically irrational for you to preserve your commitment to the relevant conditional. In Scenario 1, for instance, coming to believe that the prize is in urn A, seems rationally compatible with maintaining that if it turns out that the prize is not in A after all, it is in urn B. COM_1, however, says otherwise: the conditional would have to be dropped in at least one of the two situations.

Furthermore, it should be noted that even the comparatively strong proposition that the preservation of $\neg A > B$ and $\neg A > C$ is *required* by rationality might turn out to enjoy some support in the belief revision community. Indeed, the claim would follow from the conjunction of RAM and

(WDP2) If C is equivalent to $\neg A$, then $B \in K * A$ iff $B \in (K * C) * A$, where $A, B, C \in L_0$ [13]

WDP2 is a substantial weakening of a principle satisfied by the vast majority of proposals regarding so-called iterated revision, namely Darwiche and Pearl's postulate DP2, which states that if $\neg C \in \mathrm{Cn}(A)$, then $B \in K * A$ iff $B \in (K * C) * A$, where $A, B, C \in L_0$[14] Those familiar with the literature may wish to note that WDP2 remains weak enough to avoid the objection leveled at DP2 by Cantwell (1999, p. 73), who asks us to consider the following case:

> You initially believe that George is a mass murderer [(B)] and therefore not a law-abiding citizen [$(\neg C)$] but suspend judgment as to whether or whether or not he is a shoplifter [(A)]. You then come to believe that he is in fact a law abiding citizen, and hence not a mass murderer. But then the news comes in that George is in fact a shop-lifter, and hence not a law-abiding citizen after all.

On the assumption that you would initially have still believed that George is a mass murderer had you found out that he is a shop-lifter ($B \in K * A$), since $\neg C \in \mathrm{Cn}(A)$, DP2 entails that you ultimately recover your belief that DP2 George is a mass murderer ($B \in (K * C) * A$). This, however, is a result that Cantwell deems counterintuitive: DP2, according to him, must go. WDP2, however, does not license

[12] See Appendix, Observation 3. The DP postulates, introduced by Darwiche and Pearl (1997) and widely endorsed in the belief revision community, extend the constraints already imposed by the AGM postulates.

[13] If we assume the plausible principle of 'Extensionality' (EXT), according to which respective revisions by two logically equivalent sentences yield the same posterior belief set, WDP2 can be more simply recast as: $B \in K * A$ iff $B \in (K * \neg A) * A$.

[14] DP2 is notably satisfied by lexicographic revision (Nayak 1994), restrained revision (Booth & Meyer 2006) and natural revision (Boutillier 1996).

the conclusion that $B \in (K * C) * A$, since the requirement that $A \in \mathrm{Cn}(\neg C)$ is not satisfied (some law-breakers have respect for the retail industry).

3 Triviality from Modus Ponens

In a second no-go theorem, Bradley effectively shows that, in the presence of SUC, CC, the right-to-left direction of RAM_1 and NT*, PRES_2 is inconsistent with MP_1 (Bradley 2007, pp. 8–9). Indeed, assume NT*. From the fact that $\neg A > B \in K * A \vee B$ and $\neg A > C \in K * A \vee C$, we recover, by the right-to-left direction of RAM_1, the result that $(A \vee B) > (\neg A > B)$ and $(A \vee C) > (\neg A > C)$ are both in K. This, in turn, in conjunction with the fact that $\neg A$ is not in K, yields, by PRES_2:

$$(A \vee B) > (\neg A > B) \in K * A \text{ and } (A \vee C) > (\neg A > C) \in K * A \tag{1}$$

By SUC it follows that A is in $K * A$ and hence, by closure of belief sets, that

$$A \vee B \in K * A \text{ and } A \vee C \in K * A \tag{2}$$

But it then follows from (1) and (2), by MP_1, that

$$\neg A > B \in K * A \text{ and } \neg A > C \in K * A \tag{3}$$

in contradiction with CC. QED

3.1 *Weakening Preservation*

Bradley is initially inclined to see in this result a motivation for rejecting PRES_2 (Bradley 2007, p. 12) and offers a suggestion as to how the principle should be weakened so as to avoid the derivation of proposition (1).[15] The proposal is that, $(A \vee B) > (\neg A > B)$ and $(A \vee C) > (\neg A > C)$ might fail to be in $K * A$ because the antecedent of the embedded conditional, namely $\neg A$, is inconsistent with the sentence with which K is being revised, namely A. More precisely, the suggestion would be, where B, C, D are in L_0:

(PRES$_2^*$) If $B > (C > D) \in K$ then $B > (C > D) \in K * A$ so long as (i) $\neg A \notin K$ and (ii) $\neg C \notin \mathrm{Cn}(A)$[16]

[15] Another option would of course be to reject the right-to-left direction of RAM_1, so as to block the derivation of $(A \vee B) > (\neg A > B), (A \vee C) > (\neg A > C) \in K$. We concur with Bradley (2007, p. 9), however, that it is not the inclusion of these nested conditionals in K that is problematic here.

[16] This is not in fact quite what he proposes in print. What he does is to restrict Preservation in such a way that a conditional $B > C$ is not guaranteed to be preserved in the transition from K to $K * A$ if it is not the case that $C \in (K + B) * A$, where '$K + B$' denotes $\mathrm{Cn}(K \cup \{B\})$, $B \in L_0$ and $C \in L_1$ (Bradley 2012, p. 158). In the context of his already weakened Preservation condition PRES_2^*, this yields:

If (i) $\neg A \notin K$, (ii) $\neg B \notin \mathrm{Cn}(A)$ and (iii) $C \in (K+B)*A$, then if $B > C \in K$, then $B > C \in K*A$

The thought is that, in relation to the triviality result, whilst conditions (i) and (ii) are met, clause (iii) is not satisfied: $\neg A > B \notin (K + A \vee B) * A$. The reasoning goes as follows: either $\neg A > B$ is in $K + A \vee B$ or it is not. If it is not, the above principle clearly cannot yield the result that $\neg A > B$ is in $(K + A \vee B) * A$, since there is no conditional to preserve in the first place. But neither can it yield this result if it is in $K + A \vee B$, since the input to revision (A) is inconsistent with the antecedent of the conditional ($\neg A$), yielding a failure of condition (ii).

Setting aside worries that $K + B$ might fail to constitute a legitimate belief set, and thus $(K + B) * A$ fail to be defined (see Rott 1989 and Hansson 1992, also cited in Bradley 2007 p. 10), it is not particularly clear why Bradley requires C to be in $(K + B) * A$, rather than merely in $K * A$. The constraint is certainly not motivated by the triviality result, since the above reasoning carries over, mutatis mutandis, to the variant of the proposal in which (iii) is replaced by the requirement that $C \in K * A$.

In correspondence, Bradley has suggested that PRES_2^* might better capture his original intention.

But, in retreating to PRES_2^*, Bradley appears to have both given up too much and given up too little, endorsing a principle that is both too weak and too strong.

Regarding excessive strength, consider first the following example. My brother has just joined us on holiday, where we are planning to take diving classes. Let 'A' stand for the proposition that his wristwatch looks very much like a brand new Luminox from the Navy Seal range, 'B' for the proposition that the watch is the genuine article, 'C' for the proposition that he bought the watch from a street vendor when transiting through Bangkok and 'D' for the proposition that the watch will survive the week's nautical antics. Let K be my state of belief prior to seeing his recent horological acquisition, in which I am completely agnostic as to whether or not A, B, C, D, or any contingent Boolean combination thereof. I would take A to provide a sufficient but defeasible reason to believe B, and B, in turn, to bear the same relation to the belief that D. So B is in $K * A$ and hence $(K * A) * B$ is identical to $K * A$. But I also would take C to undercut the support that A provides for B. So, whilst D is in $K * A$, as I infer B from A and then D from B, D is not in $(K * A) * C$ and hence not in $((K * A) * B) * C$: C undercuts the probative force of the only grounds I had to hold that B, which, in turn, constituted my only reason to hold that D. Since I do not take C to undercut the support that B provides for D (a Navy Seal watch is water-resistant, whatever its place of purchase), and B remains in $(K * B) * C$, we have the result that D is in $(K * B) * C$. So, to recapitulate, D is in $(K * B) * C$, but *not* in $((K * A) * B) * C$: in terms of conditionals, given the Ramsey test, we would have $B > (C > D)$ in K but not in $K * A$. However, since, in addition, $\neg C$ is not a consequence of A and $\neg A$ is not in K, Bradley's preservation principle would judge my doxastic dispositions to be irrational.

There is also a technical result that could be taken to show that PRES_2^* is too strong. In a recent paper, Rott (2011, Corollary 6) shows that the somewhat more demanding principle PRES_2 has some very unwelcome consequences in the presence of the AGM postulates and RAM_1. Interestingly, the proof that he provides can be tweaked to show that even PRES_2^* itself sits unhappily alongside these assumptions. More specifically, it can be proven that PRES_2^* is inconsistent with NT in the presence of INC_0, VAC_0, SUC and RAM_1: hanging on to Bradley's suggested principle would plausibly require giving up on the Ramsey test for iterated conditionals.[17]

Regarding excessive weakness, PRES_2^* leaves us unable, given the Ramsey test, to predict the preservation of a number of right-nested conditionals. Suppose, for instance, that Eeyore believes that if he and Pooh decide to have a picnic the following Sunday (A), then if it proves to be a rainy day (B), then they'll be sorry (C): $A > (B > C) \in K$. This belief ought to be preserved in the event that his outlook on the weather takes an optimistic turn: $A > (B > C) \in K * \neg B$. PRES_2^* fails to predict this, since $\neg C$ is not a consequence of A.

3.2 Weakening Modus Ponens: MP for import-exporters

It should be noted, however, that Bradley appears to have now backtracked on his original diagnosis of the second triviality result. Indeed, in his most recent discussion of the result (Bradley 2012, p. 152), he expresses doubts regarding MP_1, suggesting that the scope of MP should be restricted to cases in which $B \in L_0$. In his current view, 'Modus Ponens is invalid when applied to nested conditionals and for this reason there is good reason for rejecting the argument [from the above no-go theorem to the falsity of PRES_2].'

So why the change of heart? As it happens, there exists a well-known alternative derivation of proposition (1) above. Since this line of argument does not appeal to PRES_2, it potentially provides independent grounds for the rejection of MP_1. Indeed, proposition (1) can be quickly established using RAM_0, along-

[17] Whether or not this would be too great a price to pay is up for discussion, of course. See Appendix, Observation 4 for a derivation of the result.

side the following principle of doxastic Import-Export:

(IE) $A > (B > C) \in K$ iff $(A \& B) > C \in K$.[18]

The upshot of this is that, to the extent that IE is deemed more plausible than MP_1, the latter must be rejected.[19] Furthermore, as Bradley (2007, pp. 11–2) notes, rejecting MP_1 would mean that we would have to also jettison the left-to-right direction of RAM_1, in turn, at least to the extent that we are committed to a rather intuitive 'Identity' principle, which states that revision by a sentence that is already believed yields no change in beliefs, factual or otherwise:

(ID) If $A \in K$, then $K * A = K$[20]

Indeed, assume that $A > B \in K$, where $B \in L_1$. By the left-to-right direction of RAM_1, we have $B \in K * A$. Assume that $A \in K$. By ID, we have $K * A = A$, and hence $B \in K$. QED.

Bradley, however, never ends up addressing a rather crucial question, namely: If one does insist on IE, exactly *how much* of MP_1 and RAM_1 should one be giving up? As it turns out, there is a fairly substantial restriction on the extent to which these principles can be weakened, if one assumes RAM_0 and IE. Indeed, given these two principles and a number of standard constraints on belief dynamics, it can be shown that the following weakening of the left-to-right direction of RAM_1 is recovered:

($\mathrm{WRAM}_1^{\rightarrow}$) If $A \in (K * A) * B$ or $B > C \in K * \neg A$ then, if $A > (B > C) \in K$, then $C \in (K * A) * B$

where $A, B, C \in L_0$.[21] Given ID, this in turn yields the following restricted version of MP_1:

(WMP_1) If $A \in (K * A) * B$ or $B > C \in K * \neg A$ then, if $A, A > (B > C) \in K$, then $B > C \in K$

again, where $A, B, C \in L_0$. Whilst clearly stronger than MP_0, this last principle remains weak enough to avoid the derivation of (3) from the conjunction of (1) and (2) in the triviality result. For the inference to go through, one would need to establish at least one of two claims, namely that: (a) $A \vee B \in ((K * A) * A \vee B) * \neg A$ and $A \vee C \in ((K * A) * A \vee C) * \neg A$ or (b) $\neg A > B \in K * \neg(A \vee B)$ and $\neg A > C \in K * \neg(A \vee C)$. Both (a) and (b), however, can be shown to be inconsistent with a set of extremely weak assumptions.[22]

That is not all. Indeed, one might still wonder whether WMP_1 is the strongest weakening of MP_1 that can be recovered, in the sense that the *falsity* of (a) and (b) would then suffice for the *failure* of MP_1. It turns out that this question can be answered in the positive: one can show that, given RAM_0 and a set of standard assumptions regarding belief dynamics, for all $A, B, C \in L_0$, if $A \notin (K * A) * B$ and

[18] In fact, only the right-to-left direction of this principle, the so-called rule of Export, is needed here. See Appendix, Observation 5.

[19] This rejection of MP_1 on the basis of IE is of course closely related to McGee's well-known objection to the former's logical counterpart (see McGee 1985).

[20] A restricted version of this principle, namely ID_0, in which the identity is required to hold only between the L_0-subsets of $K * A$ and K, follows from the conjunction of INC_0 and VAC_0 (proof omitted, but trivial). It should be noted that the full-blown version is violated, in the presence of the Ramsey test, by a number of proposals for handling so-called 'iterated' revision, including lexicographic revision and restrained revision. Such a violation is regrettable, however, in our view, and the aforementioned proposals can easily be modified so as to accommodate the principle.

[21] See Appendix, Observation 6.

[22] See Appendix, Observation 7.

$B > C \notin K * \neg A$, then $B > C \notin K$, even when $A, A > (B > C) \in K$.[23,24] Note in passing that one upshot of this is that, when one's belief set K is such that $A \notin (K * A) * B$, one cannot be properly said to *infer* the conclusion $B > C$ from the minor premise A, assuming the major premise $A > (B > C)$, since the conclusion will only be rationally believable to the extent that it would be so even if the minor premise were disbelieved.

These results are certainly of technical interest, but it is also worth taking a closer look at the first disjunct of the antecedent of WMP_1, namely $A \in (K * A) * B$, since it admits a rather interesting intuitive informal interpretation.

In a recent article (Chandler 2013), I have argued in favour of a Ramsey-test-style principle, not for beliefs about conditionals, but for beliefs about *reasons for belief*. The proposal was that, to check whether or not one is committed to the belief that A is a reason to believe B (notation: $A \rhd B \in K$), one should first hypothetically retract one's commitment to B, then hypothetically commit oneself to A, upon which one should verify whether or not a commitment to B is recovered in the process. In other words, one ought to check whether or not $B \in (K \dot{-} B) * A$. The episode of contraction by B prior to revision by A, which distinguishes the proposal from the Ramsey test, is crucial to capturing the idea that coming to believe A should be *relevant* to acquiring the belief that B. Indeed, as is well known, under mild conditions, the use of an exact analogue of the Ramsey test for evidential beliefs would entail a 'centering' condition, according to which, for *any* A and B, no matter how intuitively unrelated, if one believes both A and B, then one thereby believes that A is a reason to believe B (see Gärdenfors 1979). Taking into consideration the appropriately contracted belief set blocks the derivation of this result.

But this suggestion regarding reasons to believe suggests a very natural parallel proposal regarding beliefs about reasons to *not* believe—'underminers', as one could call them: In order to verify whether or not one is committed to the claim that A provides a reason to not believe that B (let us denote this by '$A \bowtie B \in K$'), one should first hypothetically *commit* oneself to B, then hypothetically commit oneself to$A$, upon which one should verify whether or not one's commitment to B *disappears* in the process. In other words, one ought to check whether or not $B \notin (K * B) * A$. Again, the episode of revision by B prior to revision by A is essential to the proposal. Indeed, were one to simply propose to check whether or not $B \notin K * A$, we would, under mild conditions, recover the unhappy result that everything that one believes is taken to be a reason to not believe anything that one does not believe.[25] Again, taking into consideration the appropriately revised belief set blocks the derivation of this result, securing a requirement of relevance between A and B.[26]

Returning to our weakening of MP_1, we can see that it can be plausibly interpreted as stating that MP_1 holds on condition that either (a) the agent does not take the antecedent of the embedded conditional to undermine—i.e. to provide a reason to not believe—the antecedent of the whole conditional ($B \bowtie A \notin K$), or (b) the agent *does* take such a relation of undermining to obtain ($B \bowtie A \in K$) but would remain committed to the embedded conditional *even if* he or she were to come to believe that the antecedent of the whole conditional is false ($B > C \in K * \neg A$). Similarly, our sufficient condition for failure of MP_1 can be interpreted as stating that the latter fails when the relevant undermining relation is taken to obtain, but the embedded conditional would not be endorsed upon coming to believe the falsity of the antecedent of the whole conditional ($B > C \notin K * \neg A$). This is rather intuitive. Indeed, assume that the assessment

[23] See Appendix, Observation 8. I am grateful to an anonymous referee for pressing me to settle this.

[24] Note that this also demonstrates that $\mathrm{WRAM}_1^{\rightarrow}$ is also the strongest weakening of the left-to-right direction of RAM_1.

[25] See Appendix, Observation 9.

[26] In fact, one can show that such a result, when coupled with our proposal would entail, under weak conditions, that if $A \in K$ and $B \notin K$, then $A \notin K * B$. See Appendix, Observation 10.

of $B > C$ goes by checking that C is present upon revising one's beliefs by B. If one takes takes B to undermine A, then a revision of one's beliefs by B will leave us unable to count on the support of A when it comes to ensuring that C. For $B > C$ to be accepted, B is going to have to stand on it's own two feet, so to speak, and be able to deliver C whether A is taken to be true or taken to be false.

This interpretation of whether or not $A \in (K * A) * B$ in terms of whether or not B is taken to undermine A is noteworthy, since it would offer an intuitive diagnosis of the apparent failure of Modus Ponens in McGee's widely discussed election case. McGee (1985, p. 462) writes:

> Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the poll results believed, with good reason:
>
> > If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.
> >
> > A Republican will win the election.
>
> Yet they did not have reason to believe
>
> > If it's not Reagan who wins, it will be Anderson.

Clearly, since those apprised of the poll will only hold that a Republican will the election on the basis of a belief that Reagan will win the election, it follows that they would take the news that Reagan will not win as undermining their belief in a Republican victory. Furthermore, it is plausible that were these people to come to believe that Reagan will not win the election, they would not believe that if Reagan loses, then Anderson will win. So the inference to McGee's final conditional would not licenced on the mere basis of the previous two.

Furthermore, it is rather plausible that, were such people to come to revise their beliefs so as to expect a Republican defeat, they would not believe that if Reagan will lose, then Anderson will win. If this is the case, according to the above analysis, not only would they not have a reason to believe the final conditional, as McGee suggests, but they would in fact be rationally prohibited from doing so.

4 Concluding remarks

As we have seen, it would appear that Bradley's first result fails to have the probative force advertised. The jury therefore remains out as to whether the weaker form of Preservation for conditionals, PRES_1, is a sensible constraint on belief dynamics.

Regarding the implications of the second result, pertaining to the stronger principle PRES_2, things are perhaps less clear cut. Much depends on the relative plausibilities of IE and MP_1. What is certainly true is that, to the extent that one wishes to endorse MP_1 in its unqualified incarnation, the proposed retreat to PRES_2^* is not a promising move. A better proposal is needed. Furthermore, for those who wish to endorse IE, belief revision theory offers clear guidance as to how MP_1 should be accordingly weakened, enabling us to derive a principle that also provides an intuitive account of the apparent failure of Modus Ponens in McGee's election case.[27]

[27] I am very grateful to Igor Douven, David Etlin, Hannes Leitgeb and Greg Wheeler for their valuable comments on earlier drafts of this paper, as well as to Richard Bradley for his helpful correspondence. Three anonymous referees for this journal also provided extremely detailed feedback that has resulted in a much improved end result. As pointed out above, I owe some important steps in the derivation of Observation 6 to Richard Booth. The research underpinning this article was funded by, and carried out at, the Center for Mathematical Philosophy of LMU Munich. I would like to thank Hannes Leitgeb and Stephan Hartmann for providing such a truly exceptional and stimulating working environment.

5 Appendix

In what follows, we make occasional use of the following further principles, in addition to some of the ones already outlined in the main body of the paper:

(SUPEX) If $B \in K * (A\&C)$, then $B \in \text{Cn}((K * A) \cup \{C\})$

(SUBEX) If $\neg C \notin K * A$ and $B \in \text{Cn}((K * A) \cup \{C\})$, then $B \in K * (A\&C)$

(DP1) If $C \in \text{Cn}(A)$, then $B \in K * A$ iff $B \in (K * C) * A$

(DP3) If $B \in K * A$, then $B \in (K * B) * A$

(DP4) If $\neg B \notin K * A$, then $\neg B \notin (K * B) * A$

The first two 'subexpansion' and 'superexpansion' postulates are also individually known as 'K*7' and 'K*8' respectively, and collectively as the 'supplementary' AGM postulates. Again, we use the subscript i to denote the restriction of these principles to B's that are members of L_i. The last three conditions are non-controversial members of the Darwiche-Pearl postulates, mentioned above in footnote 12. They are only assumed to hold for $B \in L_0$.

Observation 1. *Given the right-to-left half of RAM$_0$, INC$_0$, VAC$_0$ and SUC, if $K \in \mathbb{K}$ is such that, for pairwise inconsistent $A, B, C \in L$, we have $\neg A, \neg B, \neg C \notin K$, then $\neg A > B \in K * A \vee B$ and $\neg A > C \in K * A \vee C$*

Proof of Observation 1. We show that $\neg A > B \in K * A \vee B$. The proof that $\neg A > C \in K * A \vee C$ is analogous. Given that A and B are mutually exclusive, $(A \vee B) \supset A$ is equivalent to $\neg B$. Since $\neg B \notin K$, it then follows by closure of belief sets that $(A \vee B) \supset A \notin K$. Since Cn is assumed to satisfy the deduction theorem, we therefore have $A \notin \text{Cn}(K \cup \{A \vee B)\})$. Hence, by INC$_0$, we have: (1) $A \notin K * A \vee B$. Since $A \vee B$ is consistent, by SUC, we have $A \vee B \in K * A \vee B$ and hence, by (1) and VAC$_0$, we have: (2) $A \vee B \in (K * A \vee B) * \neg A$. Since $\neg A$ is consistent, by SUC, we have $\neg A \in (K * A \vee B) * \neg A$ and hence, by (2) and closure of belief sets, we recover: (3) $B \in (K * A \vee B) * \neg A$. Finally, from (3) and right-to-left half of RAM$_0$, we then obtain $A > B \in K * A \vee B$, as desired. □

Observation 2. *INC$_1$, VAC$_1$ and MON$_1$ jointly entail MA$_1$*

Proof of Observation 2. Consider an arbitrary $C \in L_1$ and $K \in \mathbb{K}$, such that $\neg A \notin K$ and $C \in K * A \vee B$. From the latter, by INC$_1$, it follows that $C \in \text{Cn}(K \cup \{A \vee B\})$. By MON$_1$, we have $C \in \text{Cn}(K \cup \{A \vee B, A\})$. If we assume that the consequence relation is supraclassical, the latter is equal to $\text{Cn}(K \cup \{A\})$. Since, by hypothesis, $\neg A \notin K$ and hence, by closure of belief sets, $\neg(A \vee B) \notin K$, we then have, by VAC$_1$, the result that $C \in K * A$. □

Observation 3. *VAC$_0$, SUC, SUPEX$_0$, SUBEX$_0$, and DP1 jointly entail COM$_0$*

Proof of Observation 3. Let B be a sentence in L_0 and simply assume that $\neg C \notin K * A$. By SUC, $A \in K * A$. It then follows by VAC$_0$ that $A \in (K * A) * C$. By SUPEX$_0$ and SUBEX$_0$, it follows that $B \in (K * A) * C$ iff $B \in (K * A) * A\&C$. But, since $A \in \text{Cn}(A\&C)$, DP1 entails that $B \in (K * A) * A\&C$ iff $B \in K * A\&C$. Hence $B \in (K * A) * C$ iff $B \in K * A\&B$, as required. □

Observation 4. *PRES$^*_2$ and NT are jointly inconsistent in the presence of INC$_0$, VAC$_0$, SUC and RAM$_1$.*

Proof of Observation 4 (adapted from the proof of Corollary 6 of Rott 2011). Assume NT. We will show that, given INC_0, VAC_0, SUC and RAM_1, (i) $(A \vee B) > (\neg A > B)$ is in K but (ii) it is not in $K * A \vee C$, in contradiction with PRES_2^*.

As we have already shown in Observation 1, given the right-to-left direction of RAM_0, INC_0, VAC_0 and SUC, we have (1) $\neg A > B \in K * A \vee B$ and (2) $\neg A > C \in K * A \vee C$.

Regarding (i): From (1), by the right-to-left direction of RAM_1, we then recover $(A \vee B) > (\neg A > B) \in K$, as desired.

Regarding (ii): Assume for reductio that $\neg(A \vee B) \in K * A \vee C$. By INC_0, $\neg(A \vee B) \in \text{Cn}(K \cup \{A \vee C\})$. By the deduction theorem, $(A \vee C) \supset \neg(A \vee B) \in K$. By closure of belief sets, we then have $\neg A \in K$, contrary to our assumption. So (3) $\neg(A \vee B) \notin K * A \vee C$. Since, by assumption, $\neg C \notin K$, it follows by closure of belief sets that C is consistent and hence that (4) $\neg(B \vee C) \notin \text{Cn}(A \vee C)$. From (2), (3) and (4), it then follows by PRES_2^* that $\neg A > C \in (K * A \vee C) * A \vee B$ and hence, given the left-to-right direction of RAM_0, that (5) $C \in ((K * A \vee C) * A \vee B) * \neg A$. Now, finally, assume for reductio that $(A \vee B) > (\neg A > B) \in K * A \vee C$. It follows by the left-to-right direction of RAM_1 that $B \in ((K * A \vee C) * A \vee B) * \neg A$. Since B and C are jointly inconsistent, by (5) and closure and consistency of belief sets, we have a contradiction. Hence $(A \vee B) > (\neg A > B) \notin K * A \vee C$, as desired. □

Observation 5. *Given RAM_0, SUC, and the right-to-left direction of IE, if, for pairwise inconsistent $A, B, C \in L_0$, we have $\neg A, \neg B, \neg C \notin K$ then $(A \vee B) > (\neg A > B) \in K * A$ and $(A \vee C) > (\neg A > C) \in K * A$*

Proof of Observation 5. Since $\{\neg A, A \vee B\}$ and $\{\neg A, A \vee C\}$ are consistent and respectively entail B and C, it follows, by RAM_0, SUC and closure of belief sets, that $((A \vee B) \& \neg A) > B$ and $((A \vee C) \& \neg A) > C$ are in $K * A$, and hence, by Export, so too are $(A \vee B) > (\neg A > B)$ and $(A \vee C) > (\neg A > C)$. □

Observation 6. *RAM_0, IE, SUPEX_0, SUBEX_0 and DP1-DP4 collectively entail $\text{WRAM}_1^{\rightarrow}$*

Proof of Observation 6. We divide the proof into two parts: (I) a proof that if $A \in (K * A) * B$, then if $A > (B > C) \in K$, then $C \in (K * A) * B$, and (II) a proof that if $A \notin (K * A) * B$ and $B > C \in K * \neg A$, then if $A > (B > C) \in K$, then $C \in (K * A) * B$ (where $A, B, C \in L_0$).

Regarding (I), assume $A \in (K * A) * B$ and $A > (B > C) \in K$ for conditional proof. From the latter, by the left-to-right half of IE, we have $A \& B > C \in K$. By RAM_0, this yields $C \in K * A \& B$. By DP1, we then recover $C \in (K * A) * A \& B$ and hence, by SUPEX_0, $C \in \text{Cn}((K * A) * B \cup \{A\})$. But by closure of belief sets and the fact that $A \in (K * A) * B$, $\text{Cn}((K * A) * B \cup \{A\})$ is itself equal to $(K * A) * B$. Hence $C \in (K * A) * B$, as required.

Regarding (II), we first establish two useful lemmas. The first lemma is:

Lemma 1: If $C \in (K * A) * B$, then either $C \in K * B$ or $C \in K * A \& B$

To see why this holds, assume $C \in (K * A) * B$ and hence (1) $C \in \text{Cn}((K * A) * B \cup \{\neg A\})$ and (2) $C \in \text{Cn}((K * A) * B \cup \{A\})$. We now consider two cases: $\neg A \in (K * A) * B$ and $\neg A \notin (K * A) * B$.

Assume $\neg A \in (K * A) * B$. It follows by consistency and closure of belief sets that $\neg\neg A \notin (K * A) * B$. From this, in conjunction with (1) and SUBEX_0, we recover $C \in (K * A) * \neg A \& B$. By DP2, since $\neg A \in \text{Cn}(\neg A \& B)$, it then follows that $C \in K * \neg A \& B$, and hence, by SUPEX_0, that (3) $C \in \text{Cn}(K * B \cup \{\neg A\})$. Furthermore, from the fact that $\neg A \in (K * A) * B$, we also have, by DP4, $\neg A \in K * B$ and hence $\text{Cn}(K * B \cup \{\neg A\}) = K * B$. The fact that $C \in K * B$ then immediately follows from (3) by the equality just established.

Assume that $\neg A \notin (K * A) * B$. It then follows from (2) and SUBEX_0 that $C \in (K * A) * A \& B$. By DP1, since $A \in \text{Cn}(A \& B)$, we then have $C \in K * A \& B$. This completes the proof of Lemma 1.

The second lemma is:

Lemma 2: If both $C \in K * B$ and $C \in K * A\&B$, then $C \in (K * A) * B$

To see why this holds, assume that both $C \in K * B$ and $C \in K * A\&B$. We now consider two cases: $A \in K * B$ and $A \notin K * B$.

Assume $A \in K * B$. It follows by DP3 that $A \in (K * A) * B$ and hence that $\mathrm{Cn}((K * A) * B \cup \{A\}) = (K * A) * B$. Since we have assumed that $C \in K * A\&B$, it therefore follows by SUPEX$_0$ that $C \in (K * A) * B$, as required.

Assume $A \notin K * B$ and hence, by closure of belief sets, that $\neg\neg A \notin K * B$. Since $C \in K * B$, and hence $C \in \mathrm{Cn}(K * B \cup \{\neg A\})$, it then follows from SUBEX$_0$ that $C \in K * \neg A\&B$. But we have also assumed $C \in K * A\&B$. Furthermore, $A \in \mathrm{Cn}(A\&B)$ and $\neg A \in \mathrm{Cn}(\neg A\&B)$. It therefore follows, by DP1 and DP2, that $C \in (K * A) * \neg A\&B \cap (K * A) * A\&B$. From SUPEX$_0$, we then recover the desired result that $C \in (K * A) * B$. This completes the proof of Lemma 2.

With all this in hand, assume $A > (B > C) \in K$ for conditional proof. The conjunction of left-to-right half of IE with RAM$_0$ then yields $C \in K * A\&B$. Now assume $B > C \in K * \neg A$, again for conditional proof. By the left-to-right direction of RAM$_0$, we have $C \in (K * \neg A) * B$. By Lemma 1 we therefore have either $C \in K * B$ or $C \in K * \neg A\&B$. Either way, we have $C \in K * B$, since the latter also follows from $C \in K * \neg A\&B$ by SUPEX$_0$ and the fact that $C \in K * A\&B$. So both $C \in K * A\&B$ and $C \in K * B$. It therefore follows, by Lemma 2, that $C \in (K * A) * B$, as desired. □

Observation 7. *Given INC$_0$, SUC, WDP2 and ID, if, for pairwise inconsistent $A, B, C \in L_0$, we have $\neg A, \neg B, \neg C \notin K$ then (i) neither $A \vee B \in ((K * A) * A \vee B) * \neg A$ nor $A \vee C \in ((K * A) * A \vee C) * \neg A$. Furthermore, (ii) neither $\neg A > B \in K * \neg(A \vee B)$ nor $\neg A > C \in K * \neg(A \vee C)$*

Proof of Observation 7. Assume for conditional proof that, for pairwise inconsistent $A, B, C \in L_0$, we have $\neg A, \neg B, \neg C \notin K$.

Regarding (i): We simply prove that $A \vee B \notin ((K * A) * A \vee B) * \neg A$ (the proof that $A \vee C \in ((K * A) * A \vee C) * \neg A$ is analogous). Since both A and B are inconsistent with C and $\neg C \notin K$, it follows by closure of belief sets that $A \vee B \notin K$.

Now assume for reductio that $A \vee B \in K * \neg A$. By INC$_0$, we have $A \vee B \in \mathrm{Cn}(K \cup \{\neg A\})$ and hence, by the deduction theorem, $\neg A \supset A \vee B \in K$. But, by closure of belief sets, this would mean that $A \vee B \in K$. Contradiction. Hence $A \vee B \notin K * \neg A$.

Assume for reductio again that $A \vee B \in ((K * A) * A \vee B) * \neg A$. Since $A \vee B \in \mathrm{Cn}(A)$, it follows, by SUC and closure of belief sets, that $A \vee B \in K * A$. By ID, we then recover $(K * A) * A \vee B = K * A$ and hence $A \vee B \in (K * A) * \neg A$. We then have, by WDP2, $A \vee B \in K * \neg A$. Contradiction. Hence $A \vee B \notin ((K * A) * A \vee B) * \neg A$.

Regarding (ii): We simply prove that $\neg A > B \notin K * \neg(A \vee B)$ (the proof that $\neg A > C \in K * \neg(A \vee C)$ is analogous). Assume $\neg A > B \in K * \neg(A \vee B)$ for reductio. By RAM$_0$, we have $B \in (K * \neg(A \vee B)) * \neg A$. Since $\neg(A \vee B)$ is consistent, given our initial assumption, we have, by SUC, $\neg(A \vee B) \in K * \neg(A \vee B)$, and hence, by closure of belief sets, $\neg A \in K * \neg(A \vee B)$. By ID, it then follows that $(K * \neg(A \vee B)) * \neg A = K * \neg(A \vee B)$. So, since $B \in (K * \neg(A \vee B)) * \neg A$, we have $B \in K * \neg(A \vee B)$ and hence, by closure and consistency of belief sets, $\neg(A \vee B) \notin K * \neg(A \vee B)$. Contradiction. □

Observation 8. *SUPEX$_0$, SUBEX$_0$, DP1, DP2, DP3 and RAM$_0$ collectively entail that if $A \notin (K * A) * B$ and $B > C \notin K * \neg A$, then $B > C \notin K$*

Proof of Observation 8. We prove the logically equivalent proposition that if $A \notin (K * A) * B$ and $B > C \in K$, then $B > C \in K * \neg A$. Assume that $A \notin (K * A) * B$. From this fact, it follows by DP3, that $A \notin K * B$, and hence, by closure of belief sets, that $\neg\neg A \notin K * B$. Assume that $B > C \in K$.

It follows by RAM_0 that $C \in K * B$ and hence that $C \in \mathrm{Cn}(K * B \cup \{\neg A\})$. By SUBEX_0, we then have $C \in K * \neg A \& B$. From this, in conjunction with the fact that $C \in K * B$, it follows by Lemma 2 of Observation 6 that $C \in (K * \neg A) * B$, and hence, by RAM_0, that $B > C \in K * \neg A$, as required. □

Observation 9. *Given INC_0, if (i) $A \rhd\!\lhd B \in K$ iff $B \notin K * A$, then (ii) if $A \in K$ and $B \notin K$, then $A \rhd\!\lhd B \in K$*

Proof of Observation 9. Assume that $A \in K$. It follows by closure of belief sets that $K = \mathrm{Cn}(K \cup \{A\})$. Now further assume that $B \notin K$. It follows that $B \notin \mathrm{Cn}(K \cup \{A\})$. By the contrapositive of INC_0, we then have $B \notin K * A$, and hence, by the right-to-left half of (i), $A \rhd\!\lhd B \in K$ □

Observation 10. *SUC and ID jointly entail that, for non-contradictory $A, B \in L_0$, if (i) $A \rhd\!\lhd B \in K$ iff $B \notin (K * B) * A$ and (ii) if $A \in K$ and $B \notin K$, then $A \rhd\!\lhd B \in K$, then (iii) if $A \in K$ and $B \notin K$, then $A \notin K * B$.*

Proof of Observation 10. Assume (ii). On the assumption that (i), this equivalent to the claim that if $A \in K$ and $B \notin K$ then $B \notin (K * B) * A$. Now assume that for some non-contradictory $A, B \in L_0$, we have $A \in K$ and $B \notin K$, and hence $B \notin (K * B) * A$. Now assume, for reductio, that $A \in K * B$. By ID, we have $(K * B) * A = K * B$. It follows from the fact that B is non-contradictory, in conjunction with SUC, that $B \in K * B$ and hence, by the previous equality, that $B \in (K * B) * A$. Contradiction. □

References

Adams, Ernest 1975: *The Logic of Conditionals*. Dordrecht: Reidel.

Alchourrón, Carlos, Peter Gärdenfors and David Makinson 1985: 'On the Logic of Theory Change: Partial Meet Contraction and Revision Functions'. *Journal of Symbolic Logic*, 50(2), pp. 118–39.

Booth, Richard and Thomas Meyer 2006: 'Admissible and Restrained Revision'. *Journal of Artificial Intelligence Research*, 26(1), pp. 127–51.

Boutilier, Craig 1996: 'Iterated Revision and Minimal Change of Conditional Beliefs'. *Journal of Philosophical Logic*, 25(3), pp. 263–305.

Bradley, Richard 2007: 'A Defence of the Ramsey Test'. *Mind*, 116(461), pp. 1–21.

——2012: Restricting Preservation: A Response to Hill. *Mind*, 121(481), pp. 147–59.

Cantwell, John 1999: 'Some Logics of Iterated Belief Change'. *Studia Logica: An International Journal for Symbolic Logic*, 63(1), pp. 49–84.

Chalmers, David and Alan Hajek 2007: 'Ramsey + Moore = God'. *Analysis*, 67(72), pp. 170–72.

Chandler, Jake 2013: 'Transmission Failure, AGM style'. *Erkenntnis*, 78(2), pp. 383–98.

Crocco, Gabriella, Luis Fariñas del Cerro & Andreas Herzig (eds.) 1995: *Conditionals: From Philosophy to Computer Science*. Oxford: Clarendon Press.

Eells, Ellery and Brian Skyrms (eds.) 1994: *Probability and Conditionals*. Cambridge: Cambridge University Press.

Etlin, David 2009: 'The Problem of Noncounterfactual Conditionals'. *Philosophy of Science*, 76(5), pp. 676–88.

Fuhrmann, André and Hans Rott (eds.) 1996: *Logic, Action, and Information*, de Gruyter: Berlin & New york.

Gärdenfors, Peter 1979: 'Conditionals and Changes of Belief'. In Niiniluoto and Tuomelo 1979, pp. 381–404.

——1986: 'Belief Revisions and the Ramsey Test for Conditionals'. *The Philosophical Review*, 95(1), pp. 81–93.

——1987: 'Variations on the Ramsey Test: More Triviality Results'. *Studia Logica*, 46(4), pp. 321–327.

Hájek, Alan and Ned Hall 1994: 'The Hypothesis of the Conditional Construal of Conditional Probability'. In Eells and Skyrms 1994, pp. 75–113.

Hansson, Sven Ove 1992: 'In Defense of the Ramsey Test'. *Journal of Philosophy*, 89(10), pp. 499–521.

Jin, Yi and Michael Thielscher 2007: 'Iterated Belief Revision, Revised'. *Artificial Intelligence*, 171(1), pp. 1–18.

Leitgeb, Hannes 2011: 'God – Moore = Ramsey (A Reply to Chalmers and Hajek)'. *Topoi*, 30(1), pp. 47–51.

Lewis, David 1976: 'Probabilities of Conditionals and Conditional Probabilities'. *Philosophical Review*, 85(3), pp. 297–315.

Lindström, Sten and Wlodzimierz Rabinowicz 1992: 'Belief Revision, Epistemic Conditionals and the Ramsey Test'. *Synthèse*, 91(3), pp. 195–237.

——1995: 'The Ramsey Test Revisited'. In Crocco *et al* 1995, pp. 131–82.

McGee, Van 1985: 'A Counterexample to Modus Ponens'. *Journal of Philosophy*, 82(9), pp. 462–471.

Nayak, Abhaya C. 1994: 'Iterated Belief Change Based on Epistemic Entrenchment'. *Erkenntnis*, 41(3), pp. 353–90.

Niiniluoto, Ikka, and Raimo Tuomelo (eds.) 1979: *The Logic and Epistemology of Scientific Change*, North-Holland, Amsterdam, *Acta Philosophica Fennica*, 30(2-4).

Rabinowicz, Wlodzimierz 1996: 'Stable Revision, or Is Preservation Worth Preserving?' In Fuhrmann and Rott 1996, pp. 101–28.

Rott, Hans 1986: 'Ifs, Though, and Because'. *Erkenntnis*, 25(3), pp. 345–70.

——1989: 'Conditionals and Theory Change: Revisions, Expansions, and Additions'. *Synthèse*, 81(1), pp. 91–113.

——2011: 'Reapproaching Ramsey: Conditionals and Iterated Belief Change in the Spirit of AGM'. *Journal of Philosophical Logic*, 40(2), pp. 155–91.