

ANTHONY S. GILLIES

## WHAT MIGHT BE THE CASE AFTER A CHANGE IN VIEW

Received 4 April 2004; received in revised version 27 June 2005

### 1. INTRODUCTION

When it comes to belief dynamics – how a rational agent ought to change her full-beliefs after learning something new – doxastic conservatism seems to have a lot going for it. Information is not gratuitous, the slogan goes, and so changes in view should always be *minimal* changes in view. This is the core of conservatism, and looks like just the kind of constraint that rationality imposes on belief revision. The conservative's slogan is naturally codified by requiring that revision be *preservative*:

**PRESERVATION.** If you do not believe  $\neg\varphi$  in a prior state, then revising that state with  $\varphi$  should result in a posterior state that is at least as strong as – carries as many commitments – as the prior state.

And Preservation, to put it mildly, has an impressive history.<sup>1</sup>

When it comes to epistemic modals – what, in view of the information I have, *might* and *must* be the case – another thesis has just as much going for it. Suppose I am agnostic about whether it is currently raining: neither the conclusive nor the defeasible information I have about the current weather decides the issue of its raining outside now. If my epistemic state is configured in this way, then, in view of what I take myself to know, it might be raining. Conversely, if (again, in view of what I take myself to know) it might be raining, then my epistemic state had better be structured in something like this way – in particular, I had better not think it is not raining. If and only if an agent's epistemic state does not commit her to believing that it is not raining can we count *It might be raining* as among her rational epistemic commitments – things her current epistemic state commits her to accepting.<sup>2</sup> The thesis is that rational belief is reflective in just this way:

**REFLECTIVITY.** Your epistemic state commits you to *It might be that p* iff it does not commit you to  $\neg p$ . (And dually for *must*).

To the extent that we are interested in agents who can introspect their first-order beliefs, we ought to think that something along these lines is a hard constraint, not an optional extra, for belief revision models.

The surprising fact is that these constraints – Preservation and Reflectivity – are, on pain of triviality, inconsistent: introspective agents can't be doxastically conservative in belief change.<sup>3</sup> And so it looks like we are faced with a dilemma, and must choose between being ideally reflective and being doxastically conservative in belief revision.

In order to appreciate the proper force of the dilemma, we should be clear at the outset about two simplifying assumptions. First, we will only be considering the sorts of beliefs which might better be called *acceptances* – the agent in question takes herself to know something, *p*, irrespective of the defeasible expectations she may have.<sup>4</sup> Second, although our topic is epistemic change and although we are investigating epistemic change in the context of solipsistic epistemic modals like *might*, let us assume that the impetus for such change is always non-modal. So although agents will have epistemic commitments like *It might be raining*, they only ever revise their picture of the world with respect to “plain facts” like *It is raining*.

What makes dilemmas like this interesting – more than a mere puzzle – is not the mere fact that some set of constraints, each of which we antecedently find plausible and attractive, turns out (modulo triviality) to be inconsistent. Nor is it interesting to learn that it is some particular constraint that ought to go. That much follows from the inconsistency plus some process of reflective equilibrium to find the constraint with the least plausibility. The interest instead rides on getting the right diagnosis of the difficulty since that, ideally anyway, buys us an *explanation* for why it really is that one of the constraints is something we should never have wanted in the first place. But without the right diagnosis, we have little more than a puzzle.

Our triviality result leaves us but with two options: give up on ideally reflective agents or give up on preservative belief revision. Each path is fairly well-known – the near-consensus is to opt for the latter at the expense of the former.<sup>5</sup> Preservation, I agree, is the culprit. But the orthodox diagnosis is all wrong. I aim to set that straight.

Rationality constraints in belief dynamics – like Preservation and Reflectivity – are only as good as the relations of epistemic commitment they are built on. Such relations should *a fortiori* be an explicit part of our modeling, not hidden in the background, since the properties they have will have non-trivial effects in our revision models. It is in getting straight about them, in the context of *might* and *must*, that we will see

just what has gone wrong, and why we never should have wanted Preservation in the first place.

## 2. PRELIMINARIES

Belief revision models should specify how agents ought to change their beliefs in the face of new information. To formalize this, we need to specify four things. First, what epistemic states are. Second, what the language of “epistemic inputs” is, and what the language of epistemic commitments is – i.e., what language encodes the information which is the impetus for epistemic change and what language encodes the information which expresses the beliefs of the agents. These languages may coincide, of course, but we should make room for allowing a difference. Third, we need a canonical notion of epistemic commitment, a relation – i.e., a consequence relation – that associates an epistemic state with a set of commitments. Finally, a belief revision model must specify a revision function, or a family of such functions, over the set of epistemic states (with the language of epistemic inputs as domain).

Of interest to us will be models in which we insist that inputs are confined to formulas in classical propositional logic (CPL),<sup>6</sup> but allow commitments to include epistemic modals like *It might be raining*. In order to adequately model not only “objective” beliefs – what an agent is committed to believing about ordinary facts – but also what epistemic modal commitments such an agent has in virtue of her objective beliefs, we will need a language a bit richer than CPL. Define:

DEFINITION 2.1. Let  $L^+$  be the smallest set containing CPL such that if  $\varphi \in \text{CPL}$ , then  $\diamond\varphi, \neg\diamond\varphi \in L^+$ .

Disjunction ( $\vee$ ) and the material conditional ( $\rightarrow$ ) can be introduced in the usual way in the classical fragment; the unary modality  $\square$ , intuitively expressing the epistemic *must*, abbreviates  $\neg\diamond\neg$  in the modal fragment.<sup>7</sup>

Given just the four broad constraints above on what it takes to be a revision model (and some seemingly innocuous assumptions about how they are fulfilled) plus a robust commitment to the idea that rational agents have views on what, according to them, might and might not be the case, it looks like we buy ourselves an awful lot of trouble. But these broad constraints can be implemented in two ways, and the escape routes from the trouble look a bit different in the two cases. So I will present the problem in its normal guise first. Then I will present the problem in a slightly different guise, one that makes clear just what I think has gone

wrong, and will lend itself to investigating the properties of the consequence relation implicated in the revision models.

### 3. ONE WAY

One way of seeing the trouble – indeed the way one is accustomed to seeing it – is to view the problem “AGM-wise.”<sup>8</sup> AGM epistemic states (at least for our purposes) are belief sets, i.e., theories – sets of sentences of  $L^+$  closed under Cn.  $\mathbb{K}$  is the set of such states. Even before saying just *which* theories will count as belief sets, and so before saying just what the sets in  $\mathbb{K}$  look like, we can put constraints on how an agent ought to move from one such equilibrated state to another. Minimally, we want to insist that revisions be *successful* and *consistent*:

AGM SUCCESS ( $S_{AGM}$ ). For any  $\varphi \in CPL$ :  $\varphi \in K \star \varphi$ .

AGM CONSISTENCY ( $C_{AGM}$ ). If  $\neg\varphi \notin Cn(\emptyset)$ , then  $K \star \varphi$  is consistent.<sup>9</sup>

( $S_{AGM}$ ) requires that posterior states carry commitments to the information inducing the change; ( $C_{AGM}$ ) insists that revision should be consistency preserving and, where this does not conflict with ( $S_{AGM}$ ), consistency restoring. These constraints are, for our purposes, background constraints on revision models – when it comes to our dilemma, they are non-negotiable. Preservation can be very simply put:

AGM PRESERVATION ( $P_{AGM}$ ). If  $\neg\varphi \notin K$ , then  $K \subseteq K \star \varphi$ .

And ( $P_{AGM}$ ) is meant to codify our intuitions about information preservation – keep believing as much as you can after a change of view (Harman, 1984; Gärdenfors, 1986, 1988). Together these three constraints form the basic kernel of the AGM theory. Thus: “The central rationality criterion on revisions is that the revision of  $K$  by  $A$  be the minimal change of  $K$  that is consistent and includes  $A$ ” (Gärdenfors, 1988, p. 16). So let us call a revision model *basic* iff it satisfies these three constraints.

Such constraints govern the space of revision models, saying just which ways of moving from one equilibrated state to another ought to count. But we have said nothing yet about the relevant sense of equilibrium we want such states to be in. A belief set is a theory that reflects our interest in agents who believe the (classical) consequences of what they believe. But not just any theory will do as a belief set. In the

case of epistemic modals, we also want agents to be introspective, their belief sets encoding not only objective beliefs but also being closed under what the agent considers a serious possibility – under what, in view of what she believes about the objective facts, might and might not be the case.

DEFINITION 3.1. (Levi (1979); Fuhrmann (1989)). Let  $K$  be a belief set and  $\varphi$  be any formula in CPL. Then  $\text{Poss}(K)$  is the smallest set such that:

- (1) if  $\varphi \in K$ , then  $\Box \varphi \in \text{Poss}(K)$ ; and
- (2) if  $\neg \varphi \notin K$ , then  $\Diamond \varphi \in K$ .

An AGM belief set  $K$  is *closed under Poss* iff  $\text{Poss}(K) \subseteq K$ .

Reflectivity can now be put as the constraint that  $\mathbb{K}$  contains only Poss-closed belief sets:

AGM REFLECTIVITY ( $R_{\text{AGM}}$ ).  $K \in \mathbb{K}$  only if  $K$  is closed under Poss.

Finally, let us say that an AGM revision model is *trivial* iff every state in the model is fully opinionated about objective matters of fact, ruling out objective uncertainty altogether. And so a model is non-trivial just in case it does not rule out such uncertainty:

AGM NON-TRIVIALITY ( $\text{NT}_{\text{AGM}}$ ). There is a belief set  $K \in \mathbb{K}$  and  $\varphi \in \text{CPL}$  such that  $\varphi \notin K$  and  $\neg \varphi \notin K$ .

The Fuhrmann triviality result is just this. If an AGM revision model concerns only belief sets which are closed under Poss and that model satisfies the Gärdenfors idea of the core, then that model must be trivial in the sense that it requires that rationality rules out uncertainty. Put a slightly different way, there is no non-trivial revision model faithful to both the Gärdenfors idea of the core and the idea of ideally reflective agents.<sup>10</sup>

PROPOSITION 3.1 (Fuhrmann (1989)). *If  $\langle \mathbb{K}, \star \rangle$  is basic, it is trivial.*

*Proof.* Suppose otherwise. By ( $\text{NT}_{\text{AGM}}$ ), there is a  $\varphi \in \text{CPL}$  and a  $K$  such that  $\varphi \notin K$  and  $\neg \varphi \notin K$ . Now,  $\text{Poss}(K) \subseteq K$  and hence  $\Diamond \neg \varphi \in K$ . Consider  $K \star \varphi$ . Since  $\neg \varphi \notin K$ , by ( $\text{P}_{\text{AGM}}$ )  $K \subseteq K \star \varphi$ . So, since  $\Diamond \neg \varphi \in K$ ,  $\Diamond \neg \varphi \in K \star \varphi$ . By ( $\text{S}_{\text{AGM}}$ ),  $\varphi \in K \star \varphi$ . And so, by closure under Poss,  $\Box \varphi \in K \star \varphi$ .  $K \star \varphi$  is thus inconsistent, whence by ( $\text{C}_{\text{AGM}}$ ) it follows that  $\neg \varphi \in \text{Cn}(\emptyset)$ . But all belief sets are closed under Cn, so  $\neg \varphi$  must be in  $K$ . Contradiction.  $\square$

## 4. ANOTHER WAY

The tradition is to pose this problem about epistemic modals AGM-wise. But this hides an important element in the background. One score on which we hope a belief revision model to inform us is what the relation of epistemic commitment is, and how sets of commitments generated by that relation undergo changes as agents learn new things. Rather than keeping such a relation lurking in the background, let's bring it to the fore and see what its properties ought to be. That's the second way of seeing the trouble about epistemic modals.

As before, we need suitable concepts of epistemic states, and of the right kind of relations of epistemic commitment. Intuitively, we can think of an epistemic state as a pair: a set of possible worlds constituting (a subset of) the space of possibilities, and an ordering of implausibility over the space of possibilities (Grove, 1988; Spohn, 1988). In such models the kinematics of the revision function over epistemic states is determined largely by the properties of the ordering. For our purposes here, however, we will assume very little about the structure of revision functions, and so we need not specify the nature of the orderings in any detail. In fact, for now at least, we can take epistemic states simply to be subsets of the space of possible worlds. Call such states basic states.

**DEFINITION 4.1 (BASIC STATES).** Fix a space  $W$  of possible worlds.  $s$  is a basic epistemic state iff  $s \subseteq W$ .  $I$  is the set of such  $s$ 's.

Clearly these states have no expressive advantage over belief sets – there is an obvious mapping between the AGM representation and this one. But thinking of epistemic states as sets of worlds does make it plain that the relation between states and sets of commitments – a consequence relation for our chosen language – does quite a lot of the heavy lifting in a belief revision model. And that is where I think we will make some progress.

So we will only be concerned with the class  $\mathbb{M}$  of revision models (for  $L^+$ ) that take  $I$  as the set of epistemic states. We also need to specify the kinds of commitment relations, the kinds of consequence relations  $\models \subseteq I \times L^+$ , that will concern us and constrain our revision functions by appealing to rationality postulates that implicate them. (Of course, such constraints will represent a significant interaction between the revision function and the commitment relation.) But even before saying anything at all about what kinds of consequence relations will enter into our models, and so without saying exactly what models make up  $\mathbb{M}$ , we can constrain the models by constraining how agents ought to move from one

equilibrated state to another, by insisting that our revision functions  $\circ$  be *successful* and *consistent*. Success is transparent:

SUCCESS (S). For any  $s \in I$ ,  $\varphi \in \text{CPL}$ :  $s \circ \varphi \models \varphi$ .

A basic state  $s$  will carry with it a set of commitments – formulas  $\varphi \in L^+$  such that  $s \models \varphi$ . And such a state will be consistent just in case that set of commitments is. So whether a state is consistent or not is parametric on our choice for a consequence relation, for it is only with respect to such a choice that a state carries any set of commitments – consistent or otherwise – at all.

DEFINITION 4.2. Fix a consequence relation  $\models \subseteq I \times L^+$ . A state  $s$  is inconsistent (with respect to  $\models$ ),  $s = \perp$ , iff for some  $\varphi \in L^+$  it is the case that  $s \models \varphi$  and  $s \models \neg\varphi$ . Otherwise  $s$  is consistent.

And now we can insist that revision be consistency preserving:

CONSISTENCY (C). For any  $s \in I$ ,  $\varphi \in \text{CPL}$ : if  $\llbracket \varphi \rrbracket \neq \emptyset$ , then  $s \circ \varphi \neq \perp$ .

As before, these two constraints are non-negotiable. Preservation just is the thought that, in the limiting case of revision, prior commitments should be included in posterior commitments. Letting  $B_s = \{\varphi \in L^+ : s \models \varphi\}$ , this becomes:

PRESERVATION (P). For any  $s \in I$ ,  $\varphi \in \text{CPL}$ : if  $s \not\models \neg\varphi$  then  $B_s \subseteq B_{s \circ \varphi}$ .

But not just any revision model over  $I$  will do – we are interested in models for introspective agents, and so we need to select models with reasonable consequence relations. Partition an agent's commitments in a given state into those expressible in  $\text{CPL}$  and those that are only expressible in  $L^+$ . With respect to  $\text{CPL}$ , we want rational commitment to be identified with classical satisfaction. With respect to the modal fragment, we want agents to be ideally reflective about what might and might not be the case, in view of their commitments expressible in  $\text{CPL}$ . *Might* is a reflective or autoepistemic modality, and we expect rational agents to have epistemic commitments which are reflective in this way about basic, non-modal commitments.

DEFINITION 4.3. Let  $\llbracket \cdot \rrbracket$  be the classical interpretation function over  $\text{CPL}$ . A relation  $\models \subseteq I \times L^+$  is *basically reflective* iff, for any  $\varphi \in \text{CPL}$  and  $\psi \in L^+$ :

- (1)  $s \subseteq \llbracket \varphi \rrbracket$  iff  $s \models \varphi$ ;

- (2) if  $s \models \varphi$ , then  $s \models \Box \varphi$ ;
- (3) if  $s \not\models \neg \varphi$ , then  $s \models \Diamond \varphi$ ;
- (4) if  $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket$ , then  $s \models \psi$  iff  $s \models \psi[\alpha/\beta]$ , where  $\psi[\alpha/\beta]$  is  $\psi$  with zero or more occurrences of  $\alpha$  in  $\psi$  replaced by  $\beta$ .

Being basically reflective brings some nice properties in its wake: such relations are supraclassical over the non-modal fragment, and in the modal fragment they are reflective about those non-modal commitments. Here is a contrived sort of example of such a relation:

EXAMPLE 1 (BASIC COMMITMENT). Let  $s$  be any state in  $I$ ,  $\varphi$  be any formula in  $\text{CPL}$ , and  $\psi$  be any formula in  $L^+$ . Let  $K_s$  be the smallest set such that:

- (1)  $s \subseteq \llbracket \varphi \rrbracket$  iff  $\varphi \in K_s$ ;
- (2) if  $\varphi \in K_s$ , then  $\Box \varphi \in K_s$ ;
- (3) if  $\neg \varphi \notin K_s$ , then  $\Diamond \varphi \in K_s$ ;
- (4) if  $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket$ , then  $\psi \in K_s$  iff  $\psi[\alpha/\beta] \in K_s$ .

Let the *basic commitment relation*  $\models^+ \subseteq I \times L^+$  be specified as follows:  $s \models^+ \varphi$  iff  $\varphi \in K_s$ .

This relation really is contrived: all we have done is take a state and form a stable autoepistemic theory (in  $L^+$ ) from it, and then define a consequence relation in terms of that autoepistemic theory.<sup>11</sup> Clause (1) ensures that that relation respect  $\text{CPL}$  with respect to the non-modal fragment. Clause (4) ensures that, inside the scope of the *might* operator, truth-functionally equivalent subformulas can be swapped. And clauses (2) and (3) track closure under Poss.

There are three natural desiderata that we might want met when we are thinking about modeling rational belief in the context of reflective modalities. We want to make sure our chosen consequence relation predicts that the epistemic modals *really are* reflective, that it predicts that agents are *completely opinionated* about what might and might not be the case (in view of the information they have), and we want it to predict that commitment with respect to the modals is entirely *grounded in* objective beliefs, in the sense that any two (consistent) states which share exactly the same objective beliefs share exactly the same commitments. In turn:

PROPOSITION 4.1. *Let  $\models$  be a basically reflective relation, and  $s, s'$  be any states in  $I$ . Then:*

- (1) If  $s \neq \perp$ , then  $s \models \Box \varphi$  iff  $s' \models \varphi$ .



- (2) For any  $\varphi \in \text{CPL}$ , either  $s \models \diamond \varphi$  or  $s \models \neg \diamond \varphi$ .
- (3) Let  $s, s'$  be any consistent states. If  $\{\varphi \in \text{CPL} : s \models \varphi\} = \{\varphi \in \text{CPL} : s' \models \varphi\}$ , then  $\{\varphi \in L^+ : s \models \varphi\} = \{\varphi \in L^+ : s' \models \varphi\}$ .

*Proof.*

- (1) For the left-to-right direction, suppose (toward a reductio) that  $s \models \Box \varphi$  but  $s \not\models \varphi$ . Since  $s \not\models \varphi$  it follows that  $s \models \diamond \neg \varphi$ , and so  $\{\psi : s \models \psi\}$  is inconsistent, and so  $s = \perp$  after all-completing the reductio. For the other direction, suppose  $s \models \varphi$ . And so clearly  $s \models \Box \varphi$ , as required.
- (2) Consider any  $\varphi \in \text{CPL}$ . Clearly, either  $s \subseteq \llbracket \neg \varphi \rrbracket$  or  $s \not\subseteq \llbracket \neg \varphi \rrbracket$ . Suppose the former. Then  $s \models \neg \varphi$ , and so  $s \models \neg \diamond \neg \neg \varphi$ , i.e.,  $s \models \neg \diamond \varphi$ . So suppose the latter. If  $s \not\subseteq \llbracket \neg \varphi \rrbracket$ , then  $s \not\models \neg \varphi$ , and so  $s \models \diamond \neg \neg \varphi$ , i.e.,  $s \models \diamond \varphi$ .
- (3) Here we show that  $\{\varphi \in L^+ : s \models \varphi\} \subseteq \{\varphi \in L^+ : s' \models \varphi\}$  (the other direction being symmetric). The modal fragment only contains formulas of the form  $\diamond \varphi$  and  $\neg \diamond \varphi$ , and so we have two cases to consider. For the first case, suppose that  $s \models \diamond \varphi$ . Since  $s \neq \perp$ ,  $s \not\models \neg \diamond \varphi$ , i.e.,  $s \not\models \Box \neg \varphi$ . And so, since  $\models$  is basically reflective,  $s \not\models \neg \varphi$ . But  $s$  and  $s'$  support all and only the same  $\text{CPL}$  formulas, so  $s' \not\models \neg \varphi$ . Hence  $s' \models \diamond \neg \neg \varphi$ , i.e.,  $s' \models \diamond \varphi$ . Now, for the second case, suppose  $s' \not\models \neg \diamond \varphi$ , that is, that  $s' \not\models \Box \neg \varphi$ . So, given that  $\models$  is basically reflective,  $s' \not\models \neg \varphi$ . Thus, since  $s$  and  $s'$  support all and only the same objective formulas,  $s \not\models \neg \varphi$ . It then follows that  $s \models \diamond \varphi$ . And since  $s \neq \perp$ , we have that  $s \not\models \neg \diamond \varphi$ , as required.  $\square$

Definition 4.3 does what we wanted: we want to insist that agents are reflective about what plain facts they are and are not committed to. And this makes the class of revision models with basically reflective commitment relations of interest.

Now we can see the problem posed by epistemic modals. Suppose, as before, that the only impetus for epistemic change is expressible in non-modal language: that is, let us restrict our attention to revision functions which take states and formulas of  $\text{CPL}$  to states. And suppose we only consider models in which the revision function takes states in  $I$  and formulas of  $\text{CPL}$  to states. Consider the class of revision models (for  $L^+$ ) that are built from basically reflective relations between states and commitments. That is, we insist that

REFLECTIVITY (R). A revision model  $M = \langle I, \circ, \models \rangle$  is in  $\mathbb{M}$  only if  $\models$  is basically reflective.<sup>12</sup>

Say that a model is *basic* iff it satisfies (S), (C), and (P); and call it *trivial* iff every state in the model is fully opinionated about objective matters of fact. And so a model is non-trivial just in case it does not rule out such uncertainty:

NON-TRIVIALITY (NT). For some  $s \in I$ ,  $\varphi \in \text{CPL}$ :  $s \not\models \varphi$  and  $s \not\models \neg\varphi$ .

Fuhrmann's Impossibility Theorem can now be stated rather simply: being a basic model in  $\mathbb{M}$  entails being trivial.

PROPOSITION 4.2. *Let  $M$  be any model in  $\mathbb{M}$ . If  $M$  is basic, then  $M$  is trivial.*

*Proof.* Suppose otherwise. By (NT), there is an  $s$  and  $\varphi$  (in CPL) such that  $s \not\models \varphi$  and  $s \not\models \neg\varphi$ . Since  $\models$  is basically reflective, it follows that  $s \models \diamond\neg\varphi$ . Consider  $s \circ \varphi$ . By (S),  $s \circ \varphi \models \varphi$ . Since  $s \not\models \neg\varphi$ , by (P) we have that  $\{\psi : s \models \psi\} \subseteq \{\psi : s \circ \varphi \models \psi\}$ . Hence, since  $s \models \diamond\neg\varphi$ ,  $s \circ \varphi \models \diamond\neg\varphi$ . But since  $s \circ \varphi \models \varphi$ , we know that  $s \circ \varphi \models \Box\varphi$ . Thus  $\{\psi : s \circ \varphi \models \psi\}$  is inconsistent, and so  $s \circ \varphi = \perp$ . By (C) it follows that  $\llbracket \neg\varphi \rrbracket = W$ . Since  $\models$  is supraclassical it then follows that  $s \models \neg\varphi$ . Contradiction.  $\square$

Triviality results like this need escape routes. There are some obvious nonstarters, but given that (S) and (C) are nonnegotiable, our options are rather constrained: give up on ideally reflective agents by giving up (R) or give up on preservative belief revision by giving up (P). Orthodoxy tells us to hold on to (R) at the expense of (P). But the orthodox case for this is all wrong.

Before getting to that, however, consider a strategy for retaining Preservation at the expense of Reflectivity.<sup>13</sup> Levi denies that we ever have beliefs of the form *It might be that  $p$* . Strictly speaking, the story goes, we have beliefs which are truth-evaluable (expressible in CPL) and then we have judgments of serious possibility (expressible only in the modal fragment). The former are contained in an agent's corpus of beliefs, whereas the latter are confined to the agent's meta-corpus (an AGM-style belief set which is closed under Poss). Beliefs and judgments of serious possibility are not of the same ilk, and so it is just confused to insist that an agent's *beliefs* are introspective in the way that Reflectivity seems to require.

When the Fuhrmann result is put AGM-wise, it is easy enough to see how Levi's way out does indeed provide a way out – he denies (R<sub>AGM</sub>) by denying that such modal commitments can belong to corpora in the

first place (Levi, 1988, pp. 56–58). They *are* commitments – agents in a given state *are* committed to judging this or that as seriously possible; it is just that those commitments expressing those judgments are not real beliefs, and so don't enter into the constraints on the revision functions in our models.<sup>14</sup>

But this reason for rejecting Reflectivity will not help at all for the version of the trouble I am most concerned with. The tension the Fuhrmann result brings out – witness the second way of seeing the trouble – is independent of any important distinction between an agent's beliefs and her other epistemic commitments. Perhaps only some of an agent's commitments can be properly called “beliefs”, and maybe those are just those that Levi would be happy attributing to an agent's corpus. But it does not really matter. Our constraints have not made mention of beliefs at all, but have instead been about how epistemic states (sets of worlds) and an agent's rational commitments in those states (determined by a consequence relation between states and formulas in our modal language) interact with a revision operator. And Levi clearly and rightly recognizes that there are rationality constraints on how states commit agents to judgments of serious possibility, namely, that those judgments ought to be reflective in the way that Definition 4.3 says they are. What Proposition 4.2 shows is that this concession, in the presence of our other assumptions, is enough to get the triviality result, if not for an agent's beliefs, then for the class of her epistemic commitments. So, quite apart from how we may want to further subdivide the class of epistemic commitments, christening some “beliefs” and others “judgments of serious possibility”, rejecting reflectivity in this way is not going to do the trick.<sup>15</sup>

Orthodoxy, on the other hand, wants no part of denying that agents' corpora can contain proper beliefs which bear epistemic modals.<sup>16</sup> So suppose we resolve the problem in favor of rational commitment being systematically autoepistemic at the expense of Preservation. Then the claim is that we have a choice to make since Preservation is entailed by two near-platitudes about belief dynamics.

Suppose we think that revisions go by way of (some version or other of) the Levi Identity—revision with respect to  $\varphi$  decomposes: first take the minimal weakening of that state to give up of the commitment to  $\neg\varphi$ , then expand the resulting state with respect to  $\varphi$ . A minimal weakening of a state with respect to  $\varphi$  should be the minimal change to that state such that in the resulting state the agent is no longer committed to  $\varphi$ .<sup>17</sup> And suppose that weakening idles on non-belief: if a state does not commit an agent to  $\varphi$  then weakening that state with respect to  $\varphi$  should

produce no change at all. The orthodox diagnosis notes that these constraints jointly entail ( $P_{AGM}$ ), and so we have to give up either the Levi Identity or the vacuity constraint on weakenings.<sup>18</sup> There are, it turns out, two senses of “weakening”. One, a “mind-opening” sense, is tied-up with the Levi Identity-weakening a state with respect to  $\neg\varphi$  in order to take on the commitment with respect to  $\varphi$ . But there is also a “minimal loss” sort of weakening, and this is the sort that the vacuity constraint is about. The theorem points to a tension between the two senses. It is hard to muster enthusiasm for getting rid of the minimal loss kind, so the Levi Identity has to go. Put another way, the diagnosis is that the impossibility result turns on a subtle equivocation about epistemic weakenings, and it is only the minimal loss sense which should have any purchase on our intuitions.

But this diagnosis is wide of the mark, the Levi Identity and equivocations on “weakening”, it will turn out, are red herrings. Orthodoxy has it right that Preservation is the culprit here, but it gets the diagnosis wrong and so mislocates what is pressing about the Fuhrmann result.

## 5. PRESERVATION AND PERSISTENCE

Preservation – and so both ( $P_{AGM}$ ) and its more model-theoretic counterpart ( $P$ ) – is meant to capture some sort of information economy principle. But, and this is most clearly seen in ( $P$ ), such a principle is only as good as the consequence relation upon which it is based. Glossing quite a bit, what we will see is that the Levi Identity just is not an important player in our drama. Instead it is the connection between rationality constraints – ( $P$ ) in particular – and consequence relations that is most important. And this dependence runs deep – ( $P$ ) is forced on us (plus or minus just a bit), if only the consequence relation has a certain structure. And so giving up ( $P$ ) will require moving away from consequence relations with that certain structure. If we are lucky, we will have *independent* reason to move away from such relations in the first place. If we are *very* lucky, we can find an independently motivated consequence relation that predicts exactly the shortcomings of those with the problematic structure. I will make a case for thinking we are lucky in this section, and for our being *very* lucky in the next.

Now we have to make the gloss precise. First, nothing at all turns on whether we talk about weakenings (as orthodoxy does) or revisions (as I have). Each way has its own constraint on vacuous changes in view, and each has a (well-known) way of representing the other. At a minimum,

let us require that a *downdating* function,  $\downarrow: I \times_{\text{CPL}} \rightarrow I$ , weakens commitments by increasing a state:  $s \subseteq s \downarrow \varphi$ . To require that downdates idle on non-belief, and that seems a platitude, is to require that they satisfy Easy Weakening:

**DEFINITION 5.1 (EASY WEAKENING).** Let  $M$  be a revision model for  $L^+$ .  $M$  satisfies *easy weakenings* (EW) iff for any  $s \in I$  and  $\varphi \in \text{CPL}$ : if  $s \not\models \varphi$  then  $s \downarrow \varphi = s$ .

The Levi Identity requires revision to decompose into a downdate followed by an update:

**DEFINITION 5.2 (LEVI IDENTITY).** Let  $M = \langle I, \models, \circ \rangle$  be a revision model for  $L^+$ , and  $\downarrow$  a *downdating* function.  $M$  satisfies the Levi Identity (LI) iff: for any  $s \in I$  and  $\varphi \in \text{CPL}$ ,  $s \circ \varphi = (s \downarrow \neg \varphi) \cap \llbracket \varphi \rrbracket$ .<sup>19</sup>

Just as there is an easy (limit) case of weakening, there is an easy (limit) case of revision. Suppose we want to revise our view of the world to take on a commitment which is fully compatible with our prior view. In such fortunate circumstances, the revision should reduce to just a simple sort of learning:

**DEFINITION 5.3 (EASY REVISIONS).** Let  $M = \langle I, \models, \circ \rangle$  be a revision model for  $L^+$ .  $M$  satisfies *easy revisions* (ER) iff: for any  $s \in I$  and  $\varphi \in \text{CPL}$ , if  $s \models \neg \varphi$  then  $s \circ \varphi = s \cap \llbracket \varphi \rrbracket$ .

(LI) informs us of one direction of fit that might obtain between revision functions and operations of epistemic weakening. Indeed, it is what allows one to construct a revision function on the basis of a construction for weakenings. But one might just as well go the other direction, starting with a revision function and defining downdates in terms of it. This direction for the link between revision and downdates is known, in the belief revision lore, as the Harper Identity:

**DEFINITION 5.4 (HARPER IDENTITY).** Let  $M = \langle I, \models, \circ \rangle$  be a revision model for  $L^+$ . Where  $\downarrow$  is a *downdating* function,  $M$  satisfies the Harper Identity (HI) iff: for any  $s \in I$  and  $\varphi \in \text{CPL}$ ,  $s \downarrow \varphi = s \cup (s \circ \neg \varphi)$ .

There is a close, and well-known, relationship between these two casts of characters: assuming (LI), (EW) implies (ER); and assuming (HI), (ER) implies (EW).<sup>20</sup> So there is a broad and natural sense in which

our constraints on easy weakenings and easy revisions come to just the same thing. For completeness, I will reproduce the facts here:

**PROPOSITION 5.1.** *Let  $M = \langle I, \models, \circ \rangle$  be a revision model for  $L^+$ . If  $M$  satisfies (LI) then it satisfies (EW) only if it satisfies (ER). And if  $M$  satisfies (HI) then it satisfies (ER) only if it satisfies (EW).*

*Proof.* Assume  $M$  satisfies (LI) and (EW). Suppose  $s \not\models \neg\varphi$ , and consider  $s \circ \varphi$ . By (LI),  $s \circ \varphi = (s \downarrow \neg\varphi) \cap \llbracket \varphi \rrbracket$ . But, by (EW),  $(s \downarrow \neg\varphi) = s$ , and so  $s \circ \varphi = s \cap \llbracket \varphi \rrbracket$ .

Now assume that  $M$  satisfies (HI) and (ER). Suppose  $s \not\models \neg\varphi$ , and consider  $s \downarrow \neg\varphi$ . By (HI),  $s \downarrow \neg\varphi = s \cup (s \circ \varphi)$ . But, by (ER),  $s \circ \varphi = s \cap \llbracket \varphi \rrbracket$ , and so  $s \downarrow \neg\varphi = s$ .  $\square$

The orthodox diagnosis has it that (P) is entailed by (LI) and (EW).<sup>21</sup> But this just is not so. Here is a simple counter-model. Take the (admittedly contrived) relation  $\models^+$  from Example 1.<sup>22</sup> Suppose  $s = \{w_1, w_2\}$  such that  $p$  is the case at  $w_1$  but not at  $w_2$ . Clearly  $s \not\models^+ \neg p$ . So, assuming both (LI) and (EW),  $s \circ p = s \cap \llbracket p \rrbracket = \{w_1\}$ . But while  $s \models^+ \diamond \neg p$ , not so for the posterior:  $s \circ p = \{w_1\} \not\models^+ \diamond \neg p$ . And this violates (P).

Fuhrmann and I agree on the culprit – it is (P) – but disagree on the mastermind. It is not (LI), so let us look more closely at the properties our consequence relations bear. Suppose we adopt (ER), a transparent property for vacuous revision. Then (P) will be satisfied by revision models in which the consequence relation is, in the sense defined below, *persistent*.<sup>23</sup>

**DEFINITION 5.5 (PERSISTENCE).** Fix a consequence relation  $\models \subseteq I \times L$  for a language  $L$ . A formula  $\varphi \in L$  is *persistent with respect to*  $\models$  iff for any  $s, s' \in I$ : if  $s \models \varphi$  and  $s' \subseteq s$ , then  $s' \models \varphi$ . The relation  $\models$  is *persistent* iff all formulas in  $L$  are persistent with respect to it.

**PROPOSITION 5.2.** *Let  $M = \langle I, \models, \circ \rangle$  be any revision model for  $L^+$ . If  $\models$  is persistent and  $M$  satisfies (ER), then  $M$  satisfies (P). As a corollary: if  $\models$  is persistent and  $M$  satisfies (LI) and (EW), then  $M$  satisfies (P).*

*Proof.* Suppose the hypothesis, and consider any  $\varphi \in \text{CPL}$  such that  $s \not\models \varphi$ . By (ER),  $s \circ \varphi = s \cap \llbracket \varphi \rrbracket$ . Now consider any  $\psi \in L^+$  such that  $s \models \psi$ . We have to show that  $s \circ \varphi \models \psi$ . But since  $s \circ \varphi = s \cap \llbracket \varphi \rrbracket$ , it follows that  $s \circ \varphi \subseteq s$ . Whence it follows by the persistence of  $\models$  that  $s \circ \varphi \models \psi$ . Thus,  $\{s \circ \varphi \models \psi\} \subseteq \{s \models \psi\}$ , as required. The corollary then follows immediately by Proposition 5.1.  $\square$

The assumption that the consequence relation is persistent is critical. We showed that (LI) and (EW) do not jointly entail (P) by exploiting  $\models^+$  (from Example 1). A bit of checking verifies that such a relation is not, in fact, persistent.

So, in the presence of a near-platitudinous constraint on easy revisions, persistence of the consequence relation forces (P) upon us. And although (P) is meant to capture intuitions about the rationality of information economy, in the context of epistemic modals it carries implausible predictions. For suppose that my state is characterized by just two possibilities,  $w_1$  and  $w_2$ , such that  $p$  is the case at the first but not the second of these. It seems clear that in such a state I am committed to  $\diamond\neg p$ , and any suitably reflective consequence relation will bear this out. In such a state, moreover, I am not committed to  $\neg p$ . Now consider the set of commitments after I revise this state with the new fact that  $p$ . Should it contain, properly or otherwise, the set of commitments from the earlier state? (P) requires it, but this is certainly not acceptable, for in the prior state I was committed to  $\diamond\neg p$  but in the posterior state I had better not be, lest my commitments run inconsistent (assuming that commitment is suitably reflective I will have a posterior commitment to  $\neg\diamond\neg p$ ). So we should not have wanted (P) in the first place. But, given that (ER) and persistence entail (P), we cannot jettison (P) without also abandoning either (ER) or the persistence of our chosen consequence relation.

This may seem an uncomfortable choice since (ER) seems utterly obvious and persistence generally makes for well-behaved notions of consequence. But we are lucky. We have independent reason to think that insisting on persistent relations of rational commitment is a bad idea. The argument parallels the one above: an agent in the state  $s = \{w_1, w_2\}$ , where  $p$  is true at  $w_1$  but not at  $w_2$ , ought to be committed to  $\diamond\neg p$ . But an agent who has properly more information, say an agent in state  $s' = \{w_1\}$ , clearly ought *not* be so committed. And this fact – a fact that has more to do with the structure of rational *belief* than it does about the structure of rational *belief dynamics* – is squarely at odds with the constraints placed on us by a persistent consequence relation.<sup>24</sup>

I agree that (P) is the culprit in the impossibility result, but the mastermind is not as orthodoxy would have us believe. The trouble is with persistence. It is a persistent relation of epistemic commitment, not the Levi Identity plus some platitudes about vacuous changes in view, that commits us to (P) and its implausible predictions. And we have reason apart from issues in belief dynamics to demand a different kind of consequence relation to model rational epistemic commitment in a satisfactory way anyway. Moreover, this diagnosis offers a unified

explanation of the trouble: the reasons for rejecting (P) are completely seamless with the reasons for rejecting persistence as a property of rational commitment in the context of the autoepistemic *might*.

## 6. EPISTEMIC COMMITMENT

Rationality constraints like (P) are only as good as the consequence relations they implicate. Assuming some near-platitudes about vacuous changes in view, a poor choice of a consequence relation – opting for a persistent one in modal contexts – straightaway leads to triviality. We do have an example of a non-persistent relation ( $\models^+$ ) but it is, to say the least, inelegant. We can do better. And doing better will pay off: once we get such a consequence relation, it is easy to see how to couple it with virtually any off-the-shelf story about revision functions to get a non-trivial model with modal commitments.

The most obvious model for epistemic commitment is the “propositional containment” analysis. The idea for the non-modal fragment is familiar: a state  $s$  commits an agent to  $\varphi$  just in case  $s \subseteq \llbracket \varphi \rrbracket$ . For in that case the information that  $\varphi$  carries is already present in  $s$ . But this is a story that can equally well be told by appeal to a fixed-point:  $s$  commits an agent to  $\varphi$  iff *adding* the information that  $\varphi$  carries to  $s$  would induce no change at all. If we think that adding the information of a non-modal  $\varphi$  goes by way of intersection with  $\llbracket \varphi \rrbracket$ , this is just to require that  $s \cap \llbracket \varphi \rrbracket = s$ . And, of course, in that case our two stories coincide since  $s \cap \llbracket \varphi \rrbracket = s$  iff  $s \subseteq \llbracket \varphi \rrbracket$ .

To extend this story to reflective modals, the most natural path would seem to be to generalize the notion of a “ $\varphi$ -world” so that we can pick out, in addition to  $\varphi$ -worlds, the worlds where  $\varphi$  might be the case, the set of  $\diamond\varphi$ -worlds. But there is an apparent catch: there seems to be no set answering to the description “set of worlds where  $\varphi$  might be the case”. For suppose otherwise. Take again a state  $s$  containing just two possibilities,  $w_1$  and  $w_2$ , where the first is a  $p$ -world and the second is not. Then an agent in  $s$  ought to be committed to  $\diamond p$ . But that means that both  $w_1$  and  $w_2$  are in the set of  $\diamond p$ -worlds. In particular,  $w_2$  must be. Now add the information that  $\neg p$  to  $s$ . This ought result in  $s' = \{w_2\}$ , a state that surely does not carry  $\diamond p$  as a commitment. But in that case  $w_2$  had better not be in the set of  $\diamond p$ -worlds, which is rather unfortunate since we already decided otherwise. Contradiction.

Now, what this simple argument does *not* do – and is not intended to do – is give a general proof that epistemic modals do not express propositions in the normal sense. But what it definitely *does* do is



provide a little motivation for thinking about commitment in a slightly different way. The intuition that it is pumping is that epistemic modals tell us more about an agent's information she has about the world than it does about the world itself.<sup>25</sup> And so by generalizing what exactly we mean by “the information a modal carries”, we can bootstrap our way into a story about commitment. The general idea – not a new one, but new to this way of thinking about problems in the dynamics of belief – is to identify the information that  $\varphi$  carries with the potential that  $\varphi$  has to change a state if we *update* that state with  $\varphi$ . Updating is meant here to be simple-minded: it corresponds to the simplest kind of epistemic change in which new information is added without any safeguards to prevent collapse into inconsistency.<sup>26</sup> So the information that  $\varphi$  carries is its update-profile. What is left is to define the update function.

Up to this point I have been content to follow tradition and focus on epistemic modalities which are restricted to depth at most one. But this restriction is not obviously required, and a more elegant account of commitment should be able to do without it. Let's now lift this restriction.

DEFINITION 6.1. Let  $L^\diamond$  be the smallest set including  $\text{CPL}$  and such that if  $\varphi, \psi \in L^\diamond$  then  $\neg\varphi, (\varphi \wedge \psi), \diamond\varphi \in L^\diamond$ .

As before, disjunction, the material if then, and the box for *must* are introduced in the usual way. Clearly,  $L^+ \subset L^\diamond$ , so any commitment relation for the latter is also a commitment relation for the former.

The general idea is alarmingly simple. Assume that updating is a species of simple-minded learning: no care is taken to prevent collapse into inconsistency. But remember that we are also thinking of *might* as both subjective and solipsistic: it is telling us more about an agent's information she has than it is about the world directly. Then updating a state with a modal like *It might be raining* should either do nothing to that state (if, in view of the information in that state, it really might be raining) or else reduce that state to absurdity (if, in view of the information in that state, it is not raining).<sup>27</sup> Generalizing things a bit:

DEFINITION 6.2. Let  $s$  be any state in  $I$ ,  $p$  be any atom, and  $\varphi, \psi$  be any formulas in  $L^\diamond$ . The *basic state update function*,  $\uparrow: I \times L^\diamond \rightarrow I$  is defined by the following recursion:

- (1)  $s \uparrow p = \{w \in s : w \in \llbracket p \rrbracket\}$
- (2)  $s \uparrow \neg\varphi = s \setminus (s \uparrow \varphi)$
- (3)  $s \uparrow (\varphi \wedge \psi) = (s \uparrow \varphi) \uparrow \psi$
- (4)  $s \uparrow \diamond\varphi = \{w \in s : \uparrow \varphi \neq \emptyset\}$

Updating a state with an atom  $p$  just eliminates from that state all not- $p$  possibilities. Negation is just set subtraction: updating with  $\neg\varphi$  eliminates from a state just those possibilities which would survive an update with  $\varphi$ . Conjunction is functional composition, i.e., it amounts to a sequential update, taking the conjuncts in order. And updating with *might* invites a test of a state  $s$ , returning either all of  $s$  (if the test is successful) or none of it (if it is not).

This more generalized update function on epistemic states has some interesting properties. For our purposes the point is that it does not at all have the same (trivial) dynamic properties that  $\llbracket \cdot \rrbracket$  has, and so when we use it to define a commitment relation we will get something much better.

DEFINITION 6.3.  $f$ , a function on sets with domain  $X$ , is *eliminative* with respect to a structure  $\langle X, \leq \rangle$  iff for any  $x \in X$ :  $f(x) \leq x$ .  $f$  is *distributive* (or *continuous*) over  $X$  iff for any  $x \in X$ :  $\bigcup_{a \in x} f(\{a\}) = f(x)$ ,

PROPOSITION 6.1. *The update function  $\uparrow$  is eliminative with respect to  $\langle I, \subseteq \rangle$  but is not distributive over  $I$ .*

*Proof.* A routine induction suffices to show that  $\uparrow$  is eliminative with respect to  $\langle I, \subseteq \rangle$ . To see that  $\uparrow$  does not distribute over  $I$ , let  $s = \{w_1, w_2\}$  where  $w_1 \in \llbracket p \rrbracket$  but  $w_2 \notin \llbracket p \rrbracket$ . Then  $s \uparrow \diamond p = \{w \in s : s \uparrow p \neq \emptyset\}$ , that is  $s \uparrow \diamond p = s$ . Now,  $\{w_1\} \uparrow \diamond p = \{w_1\}$  but  $\{w_2\} \uparrow \diamond p = \emptyset$ . Thus,  $s \uparrow \diamond p \neq \bigcup_{w \in s} \{w\} \uparrow \diamond p$ .  $\square$

That  $\uparrow$  is not distributive, and that the failure surrounds the behavior of the modals, means that those modals express non-local properties of epistemic states. By way of contrast,  $\llbracket \cdot \rrbracket$  is distributive. And this reflects the kind of trouble we saw above in trying to get a handle on a set answering to the description “the set of  $\diamond\varphi$ -worlds”. Our reasons for thinking there is no such set to be found were just that *might* expresses a *global* property of a state, and those are just the kinds of properties that  $\llbracket \cdot \rrbracket$  is insensitive to. Since we are thinking of commitment as a fixed-point of the update function, the nondistributivity of the latter will have dynamic effects on former.

DEFINITION 6.4 (COMMITMENT). Let  $\varphi, \psi$  be any formulas in  $L^\diamond$ . Then:

- (1) An agent in  $s$  is committed to  $\varphi$ ,  $s \Vdash \varphi$ , iff  $s \uparrow \varphi = s$ .
- (2)  $\varphi$  entails  $\psi$ ,  $\varphi \Vdash \psi$ , iff for any  $s$ :  $s \uparrow \varphi \Vdash \psi$ .
- (3)  $\varphi$  and  $\psi$  are equivalent,  $\varphi \Leftrightarrow \psi$ , iff for any  $s$ :  $s \uparrow \varphi = s \uparrow \psi$ .

This way of relating states to commitments is basically reflective. In fact, we can say something a bit stronger:  $\Vdash$  is *completely reflective* – it’s reflectivity does not respect the boundaries between the modal formulas and the non-modals.<sup>28</sup>

DEFINITION 6.5. A relation  $\models \subseteq I \times L^\diamond$  is *completely reflective* iff for any  $s \in I$ :

- (1) for any  $\varphi \in \text{CPL}$ ,  $s \models \varphi$  iff  $s \subseteq \llbracket \varphi \rrbracket$ ;
- (2) for any  $\varphi \in L^\diamond$ , if  $s \models \varphi$  then  $s \models \Box \varphi$ ;
- (3) for any  $\varphi \in L^\diamond$ , if  $s \not\models \varphi$  then  $s \models \Diamond \neg \varphi$ .

PROPOSITION 6.2.  $\Vdash$  is *completely reflective*.

*Proof.*

- (1) First, note that for any  $\varphi \in \text{CPL}$ ,  $s \uparrow \varphi = s \cap \llbracket \varphi \rrbracket$ ,  $s \Vdash \varphi$  iff  $s \uparrow \varphi = s$ . Since  $s \uparrow \varphi = s \cap \llbracket \varphi \rrbracket$  in this case, we have that  $s \uparrow \varphi = s$  iff  $s \cap \llbracket \varphi \rrbracket = s$ , i.e., iff  $s \subseteq \llbracket \varphi \rrbracket$ .
- (2) Suppose  $s \Vdash \varphi$ . Thus  $s \uparrow \varphi = s$ . But note that  $s \Vdash \Box \varphi$  iff  $s \uparrow \Box \varphi = s$ , i.e., iff  $\{w \in s : s \uparrow \varphi = s\} = s$ , and this iff  $s \uparrow \varphi = s$ , as required.
- (3) Suppose  $s \not\models \Diamond \neg \varphi$ . (We need to see that  $s \Vdash \varphi$ .) So  $s \uparrow \Diamond \neg \varphi \neq s$ . That is,  $\{w \in s : s \uparrow \neg \varphi \neq \emptyset\} \neq s$ , which implies that  $s \uparrow \neg \varphi = \emptyset$ . Whence it follows that  $(s \setminus (s \uparrow \varphi)) = \emptyset$ , and hence  $s \uparrow \varphi = s$ . And so  $s \Vdash \varphi$ .  $\square$

COROLLARY 6.3.  $\Vdash$  is *basically reflective*.

So  $\Vdash$  is reflective about commitments expressible in CPL and completely opinionated about what might and might not be the case (according to a given state). And it is genuinely a dynamic consequence relation: the non-distributivity of the update function makes  $\Vdash$  non-persistent.

PROPOSITION 6.4.  $\Vdash$  is *not persistent*.

*Proof.* Let  $s = \{w_1, w_2\}$  where  $w \in \llbracket p \rrbracket$  and  $w_2 \notin \llbracket p \rrbracket$ . Then  $s \Vdash \Diamond p$  since  $s \uparrow p \neq \emptyset$ . Now consider  $s' = s \uparrow \neg p = \{w_2\}$ .  $s' \uparrow p = \emptyset$  and so  $s' \uparrow \Diamond p = \emptyset$ . Thus  $s' \not\models \Diamond p$  even though clearly  $s' \subset s$ , violating persistence.  $\square$

Non-persistence has the consequence that entailment defined via  $\Vdash$  is not reflexive – for some choices of  $\varphi \in L^\diamond$  we have that  $\varphi \not\models \varphi$ . For let  $\varphi = \Diamond p \wedge \neg p$  and consider an  $s$  like that in the proof above. But, you might say, surely this is a mistake!

It isn't, and the reason why it isn't reveals something significant. Reflexivity of a consequence relation is intimately tied to both persistence and the properties of our update function. Our chosen consequence relation is defined as a fixed-point of our chosen update function. And our chosen update function is not distributive – it is the modals that are responsible for this. That means that the test-like behavior of the modals, whether or not they are supported in a state, is a global property of the state. Since we have taken to thinking of conjunction as functional composition we have allowed for the possibility that two conjuncts might introduce different information into a state. And that is just the possibility that gets exploited by formulas like  $\diamond p \wedge \neg p$ . The first conjunct invites a global test on a state which, if passed, gets undercut by the second conjunct.

So there is something distinctive about formulas like  $\diamond p \wedge \neg p$ . Distinctive, but not defective. We might very easily define a concept of consistency as follows:  $\varphi \in L^\diamond$  is *consistent* iff there is a state  $s \in I$  such that  $s \uparrow \varphi \neq \emptyset$ .<sup>29</sup> Inconsistency is certainly the mark of defectiveness. But the case we have been considering is not of this sort at all. Suppose I have yet to open the blinds on a particular morning in London. Given the facts I have, I believe that it might be raining out. Then I open the blinds to see that it isn't raining at all. Assuming that we are content enough to treat sequences of sentences as cases of inter-sentential conjunction, then we have just the sort of situation we are after: I first believe  $\diamond p$  and then, after learning that  $\neg p$ , I update accordingly. But notice that this required me to have acquired a bit of new information along the way. This looks to be a case of monotonic information growth of the simplest kind and so should not be lumped with either a proper revision or with the inconsistent formulas.

But we do have the expressive resources to mark the distinctive feature we are after, neither lumping  $\diamond p \wedge \neg p$  with the inconsistent formulas nor with run of the mill conjunctions like  $p \wedge q$ . The idea is simple: what formulas like  $\diamond p \wedge \neg p$  require is a change of the epistemic landscape midway through, and that shift prevents the resulting state from supporting the modal in the first conjunct. And that means that there is no *single* non-empty state which can support the whole conjunction at once. More formally: say that a formula  $\varphi \in L^\diamond$  is *cohesive* iff there is a non-empty state  $s$  such that  $s \Vdash \varphi$ , i.e. such that  $s \uparrow \varphi = \varphi$ .<sup>30</sup> And so something like  $\diamond p \wedge \neg p$ , not being cohesive, isn't really the kind of thing that an agent can have as a commitment. Strictly speaking, it can only be a *sequence* of commitments. So, there is no mistake here; we shouldn't expect, nor want,  $\Vdash$  to be reflexive in the general case.

Starting with two pretty simple ideas – that simple-minded updates can inform us in getting straight about epistemic commitment, and that these updates might have a non-trivial dynamics of their own, we end up with a consequence relation with all the right features for modeling rational epistemic commitment for ideally reflective agents: it predicts that agents really are reflective in their commitments, that modals express global properties of their states, and that their commitments based on such global properties do not always persist. What is left is to put such a consequence relation to work in belief dynamics.

## 7. OFF-THE-SHELF NON-TRIVIALITY

So far I have argued that we should reject the doxastic conservative's (P) by rejecting the idea that rational epistemic commitment in the context of *might* is persistent. This, I think, is the favored escape route to our problem. In place of a persistent consequence relation we have an independently motivated dynamic relation which gets us what we should want from epistemic modals. We can, and with striking ease, put such a dynamic consequence relation to work in revision models. In fact, since the specific details of a revision function run largely orthogonal to our main concern here, we can point to a rather broad class of revision models which, with the help of our dynamic notion of consequence, avoid triviality.

The recipe will be to take any one of a broad class of extant revision models for CPL. Identify its consequence relation, drop it from the model, and put  $\models$  in its place. The result will be a model for  $L^\diamond$  that is non-trivial, preservative in its non-modal fragment, and completely reflective.

There is a range of candidate revision functions, and motivations for them, to be found (see, e.g., Spohn, 1988; Grove, 1988; Katsuno and Mendelzon, 1991). Since such functions are reminiscent of the Stalnaker–Lewis semantics for conditionals, I call them “broadly conditional” revision functions.

**DEFINITION 7.1.** Consider any  $s \in I$ . A partial ordering  $\leq_s$  over  $\mathcal{W}$  is an *s-implausibility ordering* iff for any  $w \in s$ :

- (1) for any  $w' \in \mathcal{W}$ ,  $w \leq_s w'$ ; and
- (2) for any  $w' \notin s$ ,  $w' \not\leq_s w$ .

Let  $\min(\varphi, \leq_s)$  be the set of worlds  $w \in \llbracket \varphi \rrbracket$  such that no  $w' \in \llbracket \varphi \rrbracket$  is strictly less *s-implausible* than  $w$ . A revision function  $\circ : I \times_{\text{CPL}} \rightarrow I$  is

*broadly conditional* iff for any  $s$  and  $\varphi \in \text{CPL}$  there is an  $s$ -implausibility ordering  $\leq_s$  such that  $s \circ \varphi = \min(\varphi, \leq_s)$ .

It is easy to see that revision models based on such broadly conditional revision functions, coupled with our dynamic consequence relation, satisfy (S), (ER), and (C). The consequence relation ensures that the modals are reflective and that the model satisfies (R) on that score. Provided the space of possibilities  $W$  is non-trivial, they also satisfy the requirement codified in (NT) that rationality does not rule out uncertainty.

**PROPOSITION 7.1.** *Let  $M = \langle I, \Vdash, \circ \rangle$  be a revision model for  $L^\diamond$ , where  $\circ$  is a broadly conditional revision function, and consider any  $\varphi \in \text{CPL}$  and  $s \in I$ .*

- (1) (S)  $s \circ \varphi \Vdash \varphi$ .
- (2) (ER) If  $s \not\Vdash \neg\varphi$  then  $s \circ \varphi = s \cap \llbracket \varphi \rrbracket$ .
- (3) (C) If  $\llbracket \varphi \rrbracket \neq \emptyset$  then  $s \circ \varphi \neq \perp$ .
- (4) (NT) If  $W$  is non-trivial, i.e., if there are two distinct possibilities, then  $M$  is non-trivial.

*Proof.* For reference, here are the relevant details (omitting (4)).

- (1) Note that  $s \circ \varphi = \min(\varphi, \leq_s)$  and so  $s \circ \varphi \subseteq \llbracket \varphi \rrbracket$ . Since  $\varphi \in \text{CPL}$  we know that for any  $s'$  whatever  $s' \uparrow \varphi = s' \cap \llbracket \varphi \rrbracket$ , and so clearly  $(s \circ \varphi) \uparrow \varphi = s \circ \varphi$ , and thus  $s \circ \varphi \Vdash \varphi$ .
- (2) Suppose  $s \not\Vdash \neg\varphi$ , for an arbitrary  $\varphi \in \text{CPL}$ . Since  $\varphi \in \text{CPL}$ , this implies that  $s \cap \llbracket \varphi \rrbracket \neq \emptyset$ . Now,  $s \circ \varphi = \min(\varphi, \leq_s)$  for some  $s$ -implausibility ordering  $\leq_s$ . We first show that  $\min(\varphi, \leq_s) \subseteq s \cap \llbracket \varphi \rrbracket$ . Assume  $w \notin s \cap \llbracket \varphi \rrbracket$ . If  $w \notin \llbracket \varphi \rrbracket$ , then it follows straightaway that  $w \notin \min(\varphi, \leq_s)$ . So suppose  $w \in \llbracket \varphi \rrbracket$  but  $w \notin s$ . We need to show that there is a  $w' \in \llbracket \varphi \rrbracket$  such that  $w' <_s w$ . Since  $s \cap \llbracket \varphi \rrbracket \neq \emptyset$ , let  $w' \in s \cap \llbracket \varphi \rrbracket$ . By construction of  $\leq_s$ , it then follows that  $w' <_s w$ , as required. To see that  $s \cap \llbracket \varphi \rrbracket \subseteq \min(\varphi, \leq_s)$  consider an arbitrary  $w \in s \cap \llbracket \varphi \rrbracket$ . Again, the construction of  $\leq_s$  gives us that  $w$  is minimal in  $\leq_s$  (since  $w \in s$ ), and so must be in  $\min(\varphi, \leq_s)$ .
- (3) Suppose  $\llbracket \varphi \rrbracket \neq \emptyset$ . Since, for  $\varphi \in \text{CPL}$ ,  $s \circ \varphi = \min(\varphi, \leq_s)$  it is immediate that  $s \circ \varphi \neq \emptyset$ . (And, it doesn't take much to see that, if a state  $s \neq \emptyset$ , then  $\{\varphi : s \Vdash \varphi\}$  is consistent).  $\square$

(P) should have no purchase on our intuitions precisely because the natural way of thinking about the contours of rational commitment, when it comes to reflective agents, forces us to model commitment with a

consequence relation that is not persistent. But pairing any such broadly conditional revision function with  $\Vdash$  does give us a model preservative in its non-modal fragment. That is, where  $\circ$  is broadly conditional the following holds:

$$\begin{aligned} (\text{P}_{\text{CPL}}) : \text{For any } s \in I, \varphi \in \text{CPL} : \text{if } s \Vdash \neg\varphi \text{ then } B_s \cap \text{CPL} \\ \subseteq B_{s \circ \varphi} \cap \text{CPL} \end{aligned}$$

The reason is simple:  $\circ$  is preservative over  $\text{CPL}$ , and  $\Vdash$  respects classical satisfaction over the non-modal fragment.

This is a diagnosis, moreover, with some real teeth. There is, it turns out, a rather large and varied class of revision functions which, when coupled with a sensible notion of commitment, make for non-trivial revision models in the presence of *might*. And we have a clear prediction for when, and explanation for why, (P) leads us astray. This is progress.

## 8. CODA: CONTRACTING EXPANSIONS AND PRESERVING PRESERVATION

The lesson I want to draw from the Fuhrmann Impossibility Theorem is that epistemic commitment concerning modals is a dynamic affair, and this dynamics is a difference which makes all the difference when we put such a relation to work in belief revision models. Further, we have been able to squeeze quite a bit out of a single phenomenon about *might*. Four antecedently plausible theses – persistence, the doxastic conservative's (P), distributivity of simple-minded updates, and the natural extension of the propositional containment analysis of epistemic commitment – fell in one stroke. While this is not quite the same feat as *seven* in one blow, it does mean that my way with the Fuhrmann result has a robustly unified flavor.

When our topic is a triviality result, lessons are tied to escape routes. And since my escape route is not the only one possible, my lesson is not the only one on offer. I want to briefly sketch two other lessons in the vicinity, and say just a bit about how they relate to mine. Both alternative lessons have something right in them. But that kernel is best got at by way of the story I have been telling.

The triviality result we have been considering forces us to rethink some issues about the relationship between the doxastic conservative's (P) and the rational constraints on ideally reflective agents. The first of the alternative lessons, advocated primarily by Hans Rott (1989), suggests that the deep problem the Fuhrmann result reveals is lurking beneath (P). We have assumed all along, he says, that the easy (limit) cases of revision reduce to simple-minded updates of the sort codified by AGM-

style belief set expansion and set intersection of a state with a proposition. Such an assumption straightaway yields a version of Preservation, and so has got to go when we have introspectively rational agents. What this reveals is an *incoherence* of the concept of an “expansion” and “contraction”. He argues that “it does not make good sense any more to speak of ‘expansions’ and ‘contractions’....[G]enuine expansions and contractions simply do not exist. The only kinds of belief or theory change are revisions” (Rott, 1989, p. 109). He goes on to argue that we should define the simple-minded update with  $\varphi$  as the *revision* by  $\Box\varphi$ , and define the contraction with respect to  $\varphi$  as the *revision* by  $\Diamond\neg\varphi$ .

This lesson is not altogether the right one. First, it is just not so that if the easy (limit) case of revision reduces to simple-minded updating we get (P). We saw above that (ER) is not only a near platitude, but that it can only lead to (P) and the Fuhrmann problem if we think that commitment is persistent. But it just can’t be in the context of *might*. Second, there is ample room in conceptual space for a taxonomy of epistemic change operations that recognizes weakenings, simple-minded updates (in which there is no attempt to maintain/restore consistency should things go awry), and genuine revisions (or, if you like, non-simple-minded updates). Part of the task of a theory of epistemic change is to investigate what various instantiations of these broad categories may and must look like. We ought to be suspicious of any lesson which denies the existence of one or more of these categories. Third, when we turn to a fully general story about epistemic change with epistemic modals we will want our revision function defined for inputs like  $\Diamond\varphi$ . And when we do we will surely want to have that a state revised by  $\Diamond\varphi$  amounts to a weakening of that state with respect to  $\neg\varphi$ . But this should not be *definitional*, it should rather be a *consequence* of such a theory.

But there is something right in Rott’s lesson. Much of the moral he wants to draw turns on distinguishing AGM expansion from what he calls “additions”, revisions in the limiting case, or as we might say, consistent revisions. We know that expansion (or set intersection, in worlds-talk) induces only a trivial dynamic. But it is an open and substantive question what additions should look like, and maybe they won’t induce a trivial dynamic. This is close to the truth of the matter, I think, since a good story about simple-minded-updates in our modal context won’t give us a trivial dynamic either. And with such a story, we showed how we could get ourselves a more reasonable concept of commitment. But we need not – indeed, ought not – abandon (ER) to get it.

Now to the second alternative lesson.<sup>31</sup> One might well wonder what all the fuss about *might* really comes to. The moral is simple: we can,



and ought to, preserve Preservation – our preservation condition (P) is just fine so long as we understand that it is meant only to apply to beliefs expressible in CPL. So restricted, of course, we have no threatening triviality result. The lesson is that the modals are merely epiphenomena, and if we take care to cast our revision model carefully with respect to the non-modal fragment, then the modals should remain well-behaved.

There is a sense in which I think this is just right. The behavior of the modals should be determined by the behavior of the non-modals, and we have shown that that is one thing that makes basically reflective relations well-behaved is that they are grounded in this sense. What I have done is given this suggestion an independent motivation and codification in terms of a dynamic consequence relation. And, if we pair that relation with an off-the-shelf revision function, we get a model satisfying something like Preservation for the non-modal fragment. It is better, I say, to *predict* the shortcomings of (P) by getting clear about the properties of the consequence relation meant to model epistemic commitment than to *stipulate* its restricted scope.

But there is another sense in which I think this moral is not quite right. Lurking behind it seems to be the idea that restricting Preservation is obviously the right move to make, and it was just a mistake to ever think it applied more widely than that. Put less diplomatically: the Fuhrmann result is a non-problem. The trouble is, I think, that restricting (P) to CPL in this way drains it of much of its philosophical significance. It is meant to make precise the doxastic conservative's slogan for rational changes in view.<sup>32</sup> Without such a substantive interpretation of (P), it is hard to see the doxastic conservative as advancing something interesting – without such a substantive (P), for instance, Harman's arguments against foundationalist belief revision lose their normative force (Harman, 1984). I am no fan of either (P) or Preservation – for lots of reasons – but I think those who are fans are staking themselves to a substantive claim. And so a proper reaction to the Fuhrmann result ought not have the consequence that (P) is wrong for *trivial* reasons. It is better, I say, to draw the moral that doxastic conservatism ought not have the sort of wholesale purchase on our intuitions that many have thought.

#### ACKNOWLEDGEMENTS

Thanks to Kai von Fintel, Jim Joyce, Isaac Levi, David Makinson, Wlodek Rabinowicz, Hans Rott, Krister Segerberg, Rich Thomason,

Madison Williams, and an anonymous *JPL* referee for very helpful comments on earlier versions of this paper. It has also benefited from discussions with audiences at the Prague International Colloquium for Formal Epistemology (September 2002) and the Formal Epistemology Workshop (May 2005).

#### NOTES

<sup>1</sup> It is, of course, at the center of the AGM model of belief revision (Alchourrón et al., 1985; Gärdenfors, 1988) and probabilistic versions of it at the center of the Bayesian tradition (Gärdenfors, 1988, Chapter 5).

<sup>2</sup> Two quick terminological stipulations here. First, the relevant sense of *might* is both epistemic and solipsistic: it is an expression of relative possibility – of what, in view of the information I have, might be the case. With that said, we can drop the modifying “in view of” phrase in what follows. This sense of *might* corresponds to what Levi (1979) calls “serious possibility.” It is not, and should not be confused with, an expression of metaphysical possibility like *It didn’t rain today, but it might have*. Second, talk of “commitments”, rather than just beliefs, is a way of making sure we don’t beg any questions about whether these modals can or cannot be the proper objects of belief. Levi, for one, does not think that such modals are the kind of things that bear truth-values and so cannot be the object of belief; Gibbard seems to have a similar view (he thinks such modal expressions, like indicative conditionals, do not express propositions in the normal sense). But when I hold, in view of what information I have, that *It might be raining*, it seems all too belief-like to me not to count as a belief. So I am happy to say that they are beliefs and then adjust what we must mean by “belief” if need be. But others are less permissive. Hence the talk of commitment.

<sup>3</sup> See Levi (1988); Fuhrmann (1989); Rott (1989). I will follow Hansson (1999) and call this triviality result the “Fuhrmann Impossibility Theorem.”

<sup>4</sup> Of course, since our topic is epistemic *change* any (contingent) belief is in principle a belief which may be given up, and so is in a sense defeasible. But this muddies our waters. Defeasibility ought, at least in this context, be thought of as a property marked by the way a belief is justified, not its status with respect to possible revision. An agent has a defeasible belief in  $p$ , an *expectation* that  $p$ , if she believes  $p$  on the basis of a defeasible rule (e.g., *Normally p*). Our ordinary concept of belief is, I think, ambiguous between acceptances and expectations – and it is a difference that makes a difference in belief dynamics generally. See, for example, Rott (2001) and Gillies (2004b) for two (very different) views on the matter. But this hidden structure in our ordinary belief talk won’t enter into things here. So, having acknowledged the distinction, I propose that we ignore it.

<sup>5</sup> The classic references: Fuhrmann (1989); Rott (1989). Levi (1988) is the notable exception who goes the other direction.

<sup>6</sup> Let CPL be generated from a fixed set of atomic formulas  $\{p, q, \dots\}$  plus conjunction ( $\wedge$ ) and negation ( $\neg$ ) in the usual way.

<sup>7</sup> So the revision models of interest here will be triples  $\langle \Sigma, f, \mathcal{R} \rangle$ , where  $\Sigma$  is the set of epistemic states,  $\mathcal{R} \subseteq \Sigma \times L^+$ , and  $f: \Sigma \times \text{CPL} \rightarrow \Sigma$ . In a syntactic model – like the AGM model, where epistemic states are just identified with the (full) set of beliefs of the agent – we can drop the reference to  $\mathcal{R}$  since it just amounts to set-membership.

<sup>8</sup> The AGM theory being the well-known benchmark for theories of epistemic change Alchourrón et al. (1985); Gärdenfors (1988). This is the way that Levi (1988), Fuhrmann (1989), Rott (1989), and Hansson (1999) all discuss the triviality result.

<sup>9</sup> A belief set is consistent iff for no formula  $\varphi \in L^+$  is it the case that both  $\varphi$  and  $\neg\varphi$  are elements of it.

<sup>10</sup> We can already see that there is trouble on the horizon: belief sets that are closed under Poss are *saturated*, i.e., they have no consistent proper supersets that extend them in objective beliefs and are closed under Poss. Proof: Suppose  $K_1 \subset K_2$ , where both are closed under Poss. Then there is a  $\varphi \in \text{CPL}$  such that  $\varphi \in K_2$  and  $\varphi \notin K_1$ . Since  $\text{Poss}(K_1) \subseteq K_1$ ,  $\diamond\neg\varphi \in K_1$ , and so  $\diamond\neg\varphi \in K_2$ . But  $\varphi \in K_2$  and  $\text{Poss}(K_2) \subseteq K_2$ , and so  $\square\varphi \in K_2$ . So  $K_2$  is inconsistent.

<sup>11</sup> Autoepistemic theories were first investigated by Stalnaker, and then by Moore, in the early 1980s (Stalnaker, 1993; Moore, 1985).

<sup>12</sup> That is,  $\mathbb{M}$  is the class of models  $M$  such that  $M$  is a revision model for  $L^+$  and  $M = \langle I, \models, \circ \rangle$  where  $\models$  is basically reflective and  $\circ : I \times \text{CPL} \rightarrow I$ .

<sup>13</sup> See Levi (1988).

<sup>14</sup> This way with the trouble seems to confine us to a picture of revision that in principle cannot be extended to allow for revising in response to modal information. That seems a mistake, and so we ought to be skeptical about it – if I antecedently think my favorite café is closed but you tell me that it might be open (and I know that you have just come in from walking down the very street where that café is), I may very well want to revise my picture of things accordingly. This looks like a case where the revision is with respect to an epistemic modal, and a case where post-revision I ought to be committed to that new bit of modal information. This reason for finding Levi’s strategy wanting on this point is rather similar to the Levi–Gärdenfors divergence about nested conditionals (Levi, 1988; Gärdenfors, 1988).

<sup>15</sup> There are other, related, escape routes one might try at this point. One might, for instance, argue that our basically reflective relations don’t get us into any trouble at all if we just recognize that modal commitments are time-sensitive and that the proper representation is to put the time into the commitment: an agent in a state  $s$  is committed to  $\diamond_t p$ . Then, at a posterior  $t'$ , when her state  $s'$  commits her to  $\neg p$ , she is committed to  $\neg_{\diamond_t} p$ ; but she can be faithful to (P) by retaining her commitment to  $\diamond_t p$ . And this is perfectly consistent since it is obvious that  $\diamond_t p$  and  $\neg_{\diamond_t} p$  are not incompatible. But this is unsatisfying. First, it appeals to expressive resources in the object language that, following the tradition in the modeling of belief dynamics, I have been assuming we don’t have. Second, it proves too much. If we allow such expressive resources for modals, then it is hard to resist the move for non-modal information. Put the time into the commitment: agents have beliefs like  $p_t$  and  $q_{t^*}$ , representing that they believe  $p$  at  $t$  and  $q$  at  $t^*$ . But then we can avoid the call for non-trivial revision in the non-modal fragment altogether. An agent can be certain that  $p$  at  $t$ , and then certain that  $\neg p$  at a posterior  $t'$  without having to revise at all. This is perfectly consistent since it is obvious that  $p_t$  and  $\neg p_{t'}$  are not incompatible.

<sup>16</sup> This particular strain of orthodoxy is due to Fuhrmann.

<sup>17</sup> In the AGM framework, a weakening is a contraction function; in a possible worlds framework, “contraction” is a pretty awful description of what happens — it makes more sense to call such operations *downdates*. The neutral term “weakening” covers both.

<sup>18</sup> The relevant version of the Levi Identity is just this:  $K \star \varphi = (K \sim \neg\varphi) + \varphi$ , where  $K + \varphi = \text{Cn}(K \cup \{\varphi\})$ ; the relevant vacuity constraint on weakenings (contractions) is just that if  $\varphi \notin K$  then  $K \sim \varphi = K$ . These two conditions are indeed enough to get us

(P<sub>AGM</sub>). Proof: Suppose  $\neg\varphi \notin K$ , and consider  $K \star \varphi$ . By the Levi Identity  $K \star \varphi = (K \sim \neg\varphi) + \varphi$  and the vacuity constraint gives us that  $K \sim \neg\varphi = K$ . Thus  $K \star \varphi = K + \varphi$ , and this will clearly contain all the prior beliefs in  $K$  since Cn is monotonic.

<sup>19</sup> Strictly speaking, a model satisfies (LI) *with respect to* a choice for  $\downarrow$ .

<sup>20</sup> See, e.g., Gärdenfors (1988); Grove (1988).

<sup>21</sup> See note 18. I've just put that diagnosis in a possible worlds framework.

<sup>22</sup> Given a state  $s$ , we first formed a stable autoepistemic theory  $K_s$ , and then defined  $\models^+$  as:  $s \models^+ \varphi$  iff  $\varphi \in K_s$ .

<sup>23</sup> See Veltman (1985, 1996). Persistence and monotonicity are related: if  $\models$  is persistent, and we define a notion entailment in terms of it in the normal way, then that entailment relation will be monotonic in the normal sense.

<sup>24</sup> Non-persistent consequence relations for epistemic modals (and epistemic conditionals) is a familiar theme in dynamic semantics: see Veltman (1985, 1996); van der Does et al. (1997); Gillies (2004a).

<sup>25</sup> This is analogous to Gibbard's view about indicative conditionals (Gibbard, 1981). See also Gillies (2004a).

<sup>26</sup> AGM expansion is one such operator, as is set intersection of a prior state with a non-modal content  $\llbracket \varphi \rrbracket$ , as is the update function I will define below.

<sup>27</sup> This "test" behavior of modals is a motivating intuition for Update Semantics. See Veltman (1996); van der Does et al. (1997); Groenendijk et al. (1996); van Benthem (1996). This picture of modals is generalized (though not for the case of epistemic modals) with an accommodation mechanism in Gillies (2003).

<sup>28</sup> This is a so-called "update-to-test" consequence relation. It is worth pointing out that this one, however, does not have all of the structural properties as similar dynamic consequence relations do, in particular, those in van der Does et al. (1997), since I allow modals to occur within the scope of conjunctions, negations, and other modals.

<sup>29</sup> This captures the classical concept of consistency for the non-modal fragment as a special case: for if  $\varphi \in \text{CPL}$  then there is an  $s$  such that  $s \uparrow \varphi \neq \emptyset$  iff  $\llbracket \varphi \rrbracket \neq \emptyset$ .

<sup>30</sup> This is the same property that Groenendijk et al. (1996) call "coherence". It is easy to see that cohesiveness implies consistency but not *vice versa*.

<sup>31</sup> This has been advanced by, among others, David Makinson. Levi advocates something like this in various passages in his Levi (1988). His official line there, however, is denying (R<sub>AGM</sub>).

<sup>32</sup> To be sure, this is not Levi's view. But it does, I think, fairly characterize the conservative intuitions of the likes of Harman, Gärdenfors, and others.

## REFERENCES

- Alchourrón, C., Gärdenfors, P. and Makinson, D. (1985): On the logic of theory change: Partial meet functions for contraction and revision, *J. Symbolic Logic* **50**, 510–530.
- van Benthem, J. (1996): *Exploring Logical Dynamics*, Stanford.
- van der Does, J., Groeneveld, W. and Veltman, F. (1997): An update on 'might', *J. Logic, Lang. Inf.* **6**, 361–381.
- Fuhrmann, A. (1989): Reflective modalities and theory change, *Synthese* **81**, 115–134.
- Gärdenfors, P. (1986): Belief revisions and the Ramsey Test for conditionals, *Philos. Rev.* **95**, 81–93.
- Gärdenfors, P. (1988): *Knowledge in Flux*, MIT, Cambridge.

- Gibbard, A. (1981): Two recent theories of conditionals, in W. L. Harper, R. Stalnaker and G. Pearce (eds.), *Ifs*, Reidel, Dordrecht.
- Gillies, A. S. (2001): A new solution to Moore's paradox, *Philos. Stud.* **105**, 237–250.
- Gillies, A. S. (2003): Modal scorekeeping and 'might'-counterfactuals, in P. Dekker and R. van Rooy (eds.), *Proceedings of the Fourteenth Amsterdam Colloquium*, ILLC, Amsterdam.
- Gillies, A. S. (2004a): Epistemic conditionals and conditional epistemics, *Noûs* **38**, 585–616.
- Gillies, A. S. (2004b): New foundations for epistemic change, *Synthese* **138**, 1–48.
- Groenendijk, J., Stokhof, M. and Veltman, F. (1996): Coreference and modality, in S. Lappin (ed.), *Handbook of Contemporary Semantic Theory*, Blackwell, Oxford, pp. 179–213.
- Grove, A. (1988): Two modellings for theory change, *J. Philos. Logic* **17**, 157–170.
- Hansson, S. O. (1999): *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Kluwer, Dordrecht.
- Harman, G. (1984): Positive versus negative undermining in belief revision, *Noûs* **18**, 39–49.
- Katsuno, H. and Mendelzon, A. O. (1991): Propositional knowledge base revision and minimal change, *Artif. Intell.* **52**, 263–294.
- Katsuno, H. and Mendelzon, A. O. (1992): On the difference between updating a knowledge base and revising it, in P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 183–203.
- Kratzer, A. (1977): What 'must' and 'can' must and can mean, *Linguist. Philos.* **1**, 337–355.
- Levi, I. (1979): Serious possibility, in E. Saarinen, R. Hilpinen, I. Niiniluoto and M. Provenge Hintikka (eds.), *Essays in Honour of Jaakko Hintikka*, D. Reidel, Dordrecht, pp. 219–236. Reprinted in I. Levi, *Decisions and Revisions*, Cambridge University Press, Cambridge, pp. 147–161.
- Levi, I. (1988): Iteration of conditionals and the Ramsey Test, *Synthese* **76**, 49–81.
- Levi, I. (1996): *For the Sake of Argument*, Cambridge University Press, New York.
- Moore, R. C. (1985): Semantical considerations on nonmonotonic logic, *Artif. Intell.* **25**, 75–94.
- Rott, H. (1989): Conditionals and theory change: revisions, expansions, and additions, *Synthese* **81**, 91–113.
- Rott, H. (2001): *Change, Choice and Inference*, Oxford University Press, New York.
- Spohn, W. (1988): Ordinal conditional functions: A dynamic theory of epistemic states, in W. L. Harper and B. Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics II*, Kluwer, Dordrecht, pp. 105–134.
- Stalnaker, R. (1993): A note on nonmonotonic modal logic, *Artif. Intell.* **64**, 183–196.
- Veltman, F. (1985): Logics for Conditionals. PhD dissertation, Department of Philosophy: University of Amsterdam.
- Veltman, F. (1996): Defaults in update semantics, *J. Philos. Logic* **25**, 221–261.

*Department of Philosophy,  
University of Michigan,  
Ann Arbor, MI, USA  
e-mail: thony@umich.edu*