WHAT MIGHT BE THE CASE AFTER A CHANGE IN VIEW

ANTHONY S. GILLIES

1. Introduction

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. It might be that p, in view of what I take myself to know, just in case I have not already ruled out that p just in case I do not already accept that not-p. Put another way, if and only if my epistemic state does not commit me to believing that it is not raining can we count It might be raining as among my rational epistemic commitments—things my current epistemic state commits me to accepting.[1]

Suppose that rationality demands of us, at least when we are reflective about our everyday beliefs, to have judgments of what (in view of what we take ourselves to know) might and might not be the case and that those judgments go something like the way I sketched above. And suppose that rational changes of belief are information preserving—that is, rational changes of belief are always, in a precise sense, minimal changes of belief. The surprising fact is that these assumptions are, on pain of triviality, inconsistent.[2] And so it looks like we are faced with a dilemma: we must choose, apparently anyway, between being ideally reflective and

Thanks to Kai von Fintel, Jim Joyce, Isaac Levi, David Makinson, Wlodek Rabinowicz, Hans Rott, and Krister Segerberg 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 Logic and Semantics Group at The University of Texas at Austin.

[1]Two quick terminological stipulations here. First, the relevant sense of might is both epistemic and solipsistic: intuitively, it is an expression of relative possibility—of what, in view of what I currently take myself to know, 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.

[2]See Levi (1988); Fuhrmann (1989); Rott (1989). I will follow Hansson (1999) and call this triviality result the "Fuhrmann Impossibility Theorem".

being doxastically conservative in belief revision. Matters, I will argue, justify us in being a bit more optimistic than this dilemma might otherwise suggest.

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. 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. The idea is that 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).[3] For our present purposes, let us shelve discussion of the subtleties of epistemic change in the context of such beliefs.

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. Such a restriction seems innocuous if we are thinking of the problem facing us as one of specifying the constraints on scientific theory change; it seems plausible to suppose that theories only change in response to objective, non-modal data. The best defense for both of these background assumptions is that they represent the "best case" scenario for extending models of epistemic change to the modal case: ignore defeasible beliefs and suppose that although agents are committed to things like It might be that p, their proper epistemic revisions are in response to non-modal inputs. Ideally we want to explore epistemic change without these restrictions. The problem however, can be seen as posing a threat to the cogency of the project with our simplifying restrictions. It is, in that sense, a more powerful dilemma with our restrictions than it would be without them.

So the "easy" turns out to be rather hard. But there is more to it than that. It is natural to think that one of the most basic requirements on rational changes in view is that they be conservative—information is not gratuitous and rational agents ought to seek to minimize information loss when revising beliefs. Equivalently: their commitments, after a change in view, should be minimally different from their commitments prior to that change. But such a constraint can only be as good as the relation of commitment that it is based upon. The path toward non-trivial revision models, in the context of reflective modalities like might, lies in getting straight about the structure and dynamics of epistemic commitment.

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,

[3]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.

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—between epistemic states and sets 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)[4], 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 $\phi \in$ CPL, then $\Diamond\phi, \neg\Diamond\phi \in L^+$.

Disjunction ($\vee$) and the material conditional ($\rightarrow$) can be introduced in the usual way in the classical fragment; the unary modality $\Box$, intuitively expressing the epistemic must, abbreviates $\neg\Diamond\neg$ in the modal fragment.[5]

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 it is always presented in the belief dynamics literature—is to view the problem "AGM-wise".[6] AGM epistemic states, at least for our purposes, can just be thought of as theories, i.e. sets of sentences closed under classical consequence. Officially, an AGM belief set is just a Cn-closed subset of L^+, where Cn is the classical consequence operator. This makes for an entirely trivial epistemic commitment relation: an agent in state K is committed to ϕ iff ϕ follows from K via classical entailment; but since K is already closed under Cn, this just becomes iff $\phi \in K$. Of course, since Cn only cares about CPL, our chosen closure property only constrains belief sets in their objective component.

[4]Let CPL be generated from a fixed set of atomic formulas $\{p, q, \ldots\}$ plus conjunction ($\wedge$) and negation ($\neg$) in the usual way.

[5]So the revision models of interest here will be triples $\langle \Sigma, \mathcal{R}, f \rangle$, where Σ is the set of epistemic states, $\mathcal{R} \subseteq \Sigma \times L^+$, and $f : \Sigma \times \text{CPL} \rightarrow \Sigma$.

[6]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.

But in the case of epistemic modals, we also want belief sets to be introspective, encoding not only the objective beliefs of an agent 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 be the case.

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

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

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

Part of the project of investigating epistemic change is to think of what the constraints are in moving from one equilibrated epistemic state to another, what makes some such transitions rational and others irrational. When we limit our attention to the non-modal case, closure is a natural way to ensure that the states are properly equilibrated. And when we are thinking about agents who are reflective about the objective beliefs they have, the least we can do is close belief sets under Poss as well.

Thus far we have said a bit about what it takes to be an equilibrated epistemic state, but the Fuhrmann result is fundamentally a problem about changes in view. Let us suppose that the only impetus for epistemic change is expressible in non-modal language. If a revision model satisfies some seemingly innocuous assumptions, then on pain of triviality the whole project of extending revision models to allow commitments with respect to what might be the case is impossible.

An AGM revision model is a set $\mathbb{K}$ of Poss-closed belief sets and a revision function $\star$ taking belief sets in $\mathbb{K}$ and sentences of CPL to belief sets in $\mathbb{K}$.[7] Intuitively, the trouble is this: (1) we want rational epistemic agents to be reflective about what (non-modal) facts they do and dont accept; (2) we want to allow that it can be rational to have an incomplete picture of the world—rationality surely does not preclude uncertainty; (3) revision should be successful, consistent, and preservative. But we cannot have all of this.

In a formal dress:

AGM Nontriviality ($\mathrm{NT}_{\mathrm{AGM}}$)**:** There is a belief set $K \in \mathbb{K}$ and $\phi \in$ CPL such that $\phi \notin K$ and $\neg\phi \notin K$

AGM Success ($\mathrm{S}_{\mathrm{AGM}}$)**:** For any $\phi \in$ CPL: $\phi \in K \star \phi$.

AGM Consistency ($\mathrm{C}_{\mathrm{AGM}}$)**:** If $\neg\phi \notin \mathrm{Cn}(\emptyset)$, then $K \star \phi$ is consistent.[8]

AGM Preservation ($\mathrm{P}_{\mathrm{AGM}}$)**:** If $\neg\phi \notin K$, then $K \subseteq K \star \phi$.

The labels are mnemonic. The first, Nontriviality, is the requirement that rationality not entail having full information. The other three are part of the AGM "Basic Postulates" on revision functions: ($\mathrm{S}_{\mathrm{AGM}}$) requiring that posterior states carry commitments to the information inducing the change, ($\mathrm{C}_{\mathrm{AGM}}$) insisting that revision should be consistency preserving and, where this does not conflict with ($\mathrm{S}_{\mathrm{AGM}}$), consistency restoring. And ($\mathrm{P}_{\mathrm{AGM}}$) is meant to codify our intuitions about information preservation—keep believing as much as you can after a change of view

[7]Officially: a triple $\langle \mathbb{K}, \mathrm{Cn}^+, \star \rangle$, where $\phi \in \mathrm{Cn}^+(X)$ iff either $\phi \in X$ or $\phi \in \mathrm{Cn}(X)$. Of course, since the K's in $\mathbb{K}$ are Cn-closed, this just amounts to membership: being in epistemic state K commits you to ϕ just in case $\phi \in K$.

[8]A set is consistent iff for no formula $\phi \in L^+$ is it the case that both ϕ and $\neg\phi$ are elements of it.

(Harman, 1984; Gärdenfors, 1986; Gärdenfors, 1988). For if you are not committed in a prior state to $\neg\phi$, revising with respect to ϕ should cause no real trouble—in that case the set of your prior commitments ought to be included in the set of your posterior commitments. Taken together, these three assumptions capture what Gärdenfors takes to be the core of rational changes in view; thus he writes: "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).

The Fuhrmann 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.[9]

Proposition 3.1 (Fuhrmann (1989))**.** *There is no nontrivial revision model $\langle \mathbb{K}, \star \rangle$ that satisfies AGM Success, AGM Consistency, and AGM Preservation.*

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

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, lets 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. We will want to focus attention on those which track the motivating idea behind closing the AGM belief sets under Poss above. With these in hand, it is but a short step to laying down the needed assumptions for the Fuhrmann result.

Intuitively, we can think of an epistemic state as a pair: a set of possibile 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

[9]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 $\phi \in$ CPL such that $\phi \in K_2$ and $\phi \notin K_1$. Since $\mathrm{Poss}(K_1) \subseteq K_1$, $\Diamond\neg\phi \in K_1$, and so $\Diamond\neg\phi \in K_2$. But $\phi \in K_2$ and $\mathrm{Poss}(K_2) \subseteq K_2$, and so $\Box\phi \in K_2$. So K_2 is inconsistent.

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.

Partition an agents commitments in a given state into those expressible in CPL and those that are only expressible in L^+. With respect to CPL, we want rational commitment to be identified with classical entailment. 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 CPL. If I believe that it is raining, then I do not consider it a serious possibility that it is not raining; if I don't believe that it is raining, then I consider it a real possibility that it is not raining. Might, in other words, 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 4.2. Let $[\![\cdot]\!]$ be the classical interpretation function over CPL. A relation $\models \subseteq I \times L^+$ is *basically reflective* iff, for any $\phi \in$ CPL:

(1) $s \subseteq [\![\phi]\!]$ iff $s \models \phi$;
(2) if $s \models \phi$, then $s \models \Box\phi$;
(3) if $s \not\models \neg\phi$, then $s \models \Diamond\phi$;
(4) if $[\![\alpha]\!] = [\![\beta]\!]$, then $s \models \phi$ iff $s \models \phi[\alpha/\beta]$, where $[\alpha/\beta]$ is just like ϕ with zero or more occurrences of α in ϕ replaced by β.

A basic state s is consistent with respect to $\models$ iff for no $\phi \in L^+$ is it the case that both ϕ and $\neg\phi$ are in $\{\psi \in L^+ : s \models \psi\}$. In this case we write $s \neq \bot$.

There is an obvious and natural interest in relations of epistemic commitment that are basically reflective: they 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, ϕ be any formula in CPL, and ψ be any formula in L^+. Let K_s be the smallest set such that:

(1) $s \subseteq [\![\phi]\!]$ iff $\phi \in K_s$;
(2) if $\phi \in K_s$, then $\Box\phi \in K_s$;
(3) if $\neg\phi \notin K_s$, then $\Diamond\phi \in K_s$;
(4) $[\![\alpha]\!] = [\![\beta]\!]$, 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: $s \models^+ \phi$ iff $\phi \in K_s$.

The K_s sets are a convenient intermediary to defining this basic commitment relation taking basic states to sets of beliefs (formulas in L^+). Clause (1) ensures that that relation respect 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. The K_s's, in fact, are just stable autoepistemic theories (Stalnaker, 1993; Moore, 1985), and so the basic commitment relation tracks a very basic kind of autoepistemic consequence.

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 *completed 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 $s \neq \bot$, then $s \models \Box\phi$ iff $s \models \phi$.*
(2) *For any $\phi \in$ CPL, either $s \models \Diamond\phi$ or $s \models \neg\Diamond\phi$.*
(3) *Let s, s' be any consistent states. If $\{\phi \in \text{CPL} : s \models \phi\} = \{\phi \in \text{CPL} : s' \models \phi\}$, then $\{\phi \in L^+ : s \models \phi\} = \{\phi \in L^+ : s' \models \phi\}$.*

Proof.
(1) For the left-to-right direction, suppose (toward a reductio) that $s \models \Box\phi$ but $s \not\models \phi$. Since $s \not\models \phi$ it follows that $s \models \Diamond\neg\phi$, and so $\{\psi : s \models \psi\}$ is inconsistent, and so $s = \bot$ after all—completing the reductio. For the other direction, suppose $s \models \phi$. And so clearly $s \models \Box\phi$, as required.

(2) Consider any $\phi \in$ CPL. Clearly, either $s \subseteq [\![\neg\phi]\!]$ or $s \not\subseteq [\![\neg\phi]\!]$. Suppose the former. Then $s \models \neg\phi$, and so $s \models \neg\Diamond\neg\neg\phi$—i.e., $s \models \neg\Diamond\phi$. So suppose the latter. If $s \not\subseteq [\![\neg\phi]\!]$, then $s \not\models \neg\phi$, and so $s \models \Diamond\neg\neg\phi$—i.e., $s \models \Diamond\phi$.

(3) Here we show that $\{\phi \in L^+ : s \models \phi\} \subseteq \{\phi \in L^+ : s' \models \phi\}$ (the other direction being symmetric). The modal fragment only contains formulas of the form $\Diamond\phi$ and $\neg\Diamond\phi$, and so we have two cases to consider. For the first case, suppose that $s \models \Diamond\phi$. Since $s \neq \bot$, $s \not\models \neg\Diamond\phi$, i.e., $s \not\models \Box\neg\phi$. And so, since $\models$ is basically reflective, $s \not\models \neg\phi$. But s and s' support all and only the same CPL formulas, so $s' \not\models \neg\phi$. Hence $s' \models \Diamond\neg\neg\phi$—i.e., $s' \models \Diamond\phi$. Now, for the second case, suppose $s' \not\models \neg\Diamond\phi$, that is, that $s' \not\models \Box\neg\phi$. So, given that $\models$ is basically reflective, $s' \not\models \neg\phi$. Thus, since s and s' support all and only the same objective formulas, $s \not\models \neg\phi$. It then follows that $s \models \Diamond\phi$. And since $s \neq \bot$, we have that $s \not\models \neg\Diamond\phi$, as required. $\square$

All of this suggests that Definition 4.2 does what we pre-theoretically wanted: we want to insist that agents are reflective about what plain facts (expressible in CPL) they are and are not committed to. And this makes the class of revision models with basically reflective commitment relations of interest. Now to the problem posed by epistemic modals. As before, 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 CPL to states. This restriction is meant to pick out the "easy case" for extending belief revision models to modal languages—allowing, indeed insisting, that agents have modal commitments but that the revision process is driven by non-modal inputs. When we are thinking of our problem as a problem about theory change this restriction corresponds to the natural one of supposing that scientific theories change as a result of observable facts and events. The surprising fact is that this "easy case" could not be harder.

Consider the class of revision models (for L^+) that are built from basically reflective relations between states and commitments. And suppose we only consider in which the revision function takes states in I and formulas of CPL to states. Let $\mathbb{M}$ pick out this class of revision models.[10] We already know that the revision models in $\mathbb{M}$ will be reflective in the way we are after. And, just to tidy up some of the core requirements for revision models, let (for any state s) $B_s = \{\phi : s \models \phi\}$ be the set of commitments (with respect to the relation $\models$) of an agent in s. Then we have:

Non-Triviality (NT): For some $s \in I$, $\phi \in$ CPL: $s \not\models \phi$ and $s \not\models \neg\phi$
Success (S): For any $s \in I$, $\phi \in$ CPL: $s \circ \phi \models \phi$
Consistency (C): For any $s \in I$, $\phi \in$ CPL: if $[\![\phi]\!] \neq \emptyset$, then $s \circ \phi \neq \bot$
Preservation (P): For any $s \in I$, $\phi \in$ CPL: if $s \not\models \neg\phi$ then $B_s \subseteq B_{s\circ\phi}$

It doesn't take much squinting to see that these are the just the same assumptions as the AGM-wise assumptions above: rationality doesn't rule out uncertainty, and revision should be successful, consistency preserving and restoring (when this doesnt conflict with being successful), and preservative. Unpacking (P), it requires that when $s \not\models \neg\phi$, $\{\psi : s \models \psi\} \subseteq \{\psi : s \circ \phi \models \psi\}$. That's just the AGM constraint: consistent addition should result in no information loss between prior and posterior state. But—and we will return to this point below—this more model-theoretic rendition of Preservation does make rather explicit that it is a constraint that lives and dies with the property of the consequence relation that it is built upon. More epistemologically put: doxastic conservatism of the kind embodied by Preservation is only as good as the relation from states to sets of epistemic commitments that it implicates. Since these last three criteria track the AGM-wise core of the basic postulates, say that a model M in $\mathbb{M}$ is basic iff it satisfies (S), (C), and (P). Fuhrmann's Impossibility Theorem can now be stated rather generally: 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 ϕ (in CPL) such that $s \not\models \phi$ and $s \not\models \neg\phi$. Since $\models$ is basically reflective, it follows that $s \models \Diamond\neg\phi$. Consider $s \circ \phi$. By (S), $s \circ \phi \models \phi$. Since $s \not\models \neg\phi$, by (P) we have that $\{\psi : s \models \psi\} \subseteq \{\psi : s \circ \phi \models \psi\}$. Hence, since $s \models \Diamond\neg\phi$, $s \circ \phi \models \Diamond\neg\phi$. But since $s \circ \phi \models \phi$, we know that $s \circ \phi \models \Box\phi$. Thus $\{\psi : s \circ \phi \models \psi\}$ is inconsistent, and so $s \circ \phi = \bot$. By (C) it follows that $[\![\neg\phi]\!] = W$. Since $\models$ is supraclassical it then follows that $s \models \neg\phi$. Contradiction. □

So what we thought should be easy has turned out to be rather hard. Triviality results like this need escape routes. There are some obvious non-starters. We do not want to jettison (NT), embracing the wildly implausible thesis that rationality rules out uncertainty. It is hard to muster any enthusiasm for denying (C). And, though some sense can be made of belief revision models in which (S) does not hold, the point of the assumption here is more or less definitional—we are interested in the rational dynamics of a state s when an agent in s takes on a commitment with respect to ϕ.

---

[10]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 $\circ : I \times$ CPL $\to I$.

Apart from the non-starters, it looks like our options are rather constrained: either it is the current incarnation of the doxastic conservative's preservation condition, (P), or else it is the kind of systematic autoepistemic commitment to epistemic modals, codified by our basically reflective relations, which has to go. And each strategy has found some favor. Let's take them in reverse order. Levi, for one, denies that we ever have beliefs of the form It might be that p. Strictly speaking, on this view, 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 agents corpus (AGM-style belief set), whereas the latter are confined to the agents meta-corpus (an AGM-style belief set which is closed under Poss). Beliefs and judgments of serious possibility are not of the same ilk.

When the Fuhrmann result is expressed in terms of AGM belief sets, it is easy enough to see how Levi's way out does indeed provide a way out—he denies closing belief sets under Poss 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.[11]

But the version of the trouble I am most concerned with is independent of any important distinction between an agent's beliefs and her other epistemic commitments. Instead we have an epistemic state (a set of worlds) and a relation of epistemic commitment (a consequence relation) relating states to subsets of L^+. Perhaps only some of an agents commitments can be properly called "beliefs", and maybe those are just those that Levi would be happy attributing to an agents corpus. But it does not really matter. Our assumptions—(NT), (S), (C), (P)—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 4.2 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", we have a problem not so easily dispatched.[12]

---

[11]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).

[12]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

Fuhrmann wants no part of denying that agents' corpora can contain proper beliefs which bear epistemic modals. So suppose we resolve the problem in favor of rational commitment being systematically autoepistemic at the expense of preservation conditions like (P). Then, Fuhrmann argues, we have a choice to make. Suppose we think that revisions go by way of (some version or other of) the Levi Identity—that the minimal revision of a state so as to include ϕ is the same as first the minimal weakening of that state to give up of the commitment to $\neg\phi$ followed by expanding it with respect to ϕ. What about this notion of "weakening"? The rational weakening of a state with respect to ϕ should be the minimal change to that state such that in the resulting state the agent is no longer committed to ϕ.[13] As a limiting case, if a state does not commit an agent to ϕ then weakening that state with respect to ϕ should produce no change at all. This near-platitudinous constraint on weakenings is a vacuity constraint. What Fuhrmann proves, in effect, is that the Levi Identity, Consistency, and such a vacuity constraint cannot all be non-trivially met. He then argues that a preservation condition—in fact, ($\mathrm{P}_{\mathrm{AGM}}$)—is entailed by such a vacuity constraint plus the Levi Identity. Thus, one cannot abandon that preservation condition without giving up either the vacuity constraint on weakenings or the Levi Identity. It is, according to Fuhrmann, the Levi Identity which is the culprit here. There are two senses of "weakening". One, a "mind-opening" sense, is tied-up with the Levi Identity—weakening a state with respect to $\neg\phi$ in order to take on the commitment with respect to ϕ. 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—in particular, weakenings for which the Levi Identity is true do not satisfy the vacuity constraint. That is a pre-theoretic mark against them, and hence a reason to suppose that the Levi Identity should be abandoned. Put another way, Fuhrmann's view 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.

5. Preservation and persistence

So Fuhrmann argues that it is the Levi Identity which is to blame for the triviality result. But the version of the trouble which concerns me most, Proposition 4.2, makes no mention at all of the Levi Identity or weakenings. In fact, I think the discussion about the Levi Identity is a red herring. I want to argue that the trouble surrounds our preservation condition (P), and it surrounds it in an essential way.

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.

[13]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.

We will have cause to reject it independent of any considerations about the Levi Identity. Once this is demonstrated, we can see where Fuhrmanns diagnosis goes wrong.

Pre-theoretically, our preservation conditions—both ($\mathrm{P}_{\mathrm{AGM}}$) and its more model-theoretic counterpart (P)—are 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. If we insist on a transparent property of vacuous revisions, then (P) will be satisfied by revision models in which the relation of commitment is, in the sense defined below, *persistent*.[14] Making this precise:

Definition 5.1 (Persistence)**.** Fix a consequence relation $\models \subseteq I \times L$ for a language L. A formula $\phi \in L$ is *persistent with respect to* $\models$ iff for any $s, s' \in I$: if $s \models \phi$ and $s' \subseteq s$, then $s' \models \phi$. The relation $\models$ is *persistent* iff all formulas in L are persistent with respect to it.

The interesting cases of revision are those in which we have to relinquish some of our commitments as well as take on some new ones. But there is an easier case we sometimes are fortunate enough to be in: 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 5.2 (Vacuous Revisions)**.** Let $M = \langle I, \models, \circ \rangle$ be a revision model for L^+. M satisfies *vacuous revisions* (VR) iff: for any $s \in I$ and $\phi \in \textsc{cpl}$, if $s \not\models \neg\phi$ then $s \circ \phi = s \cap [\![\phi]\!]$.

Proposition 5.1. *Let $M = \langle I, \models, \circ \rangle$ be any revision model for L^+. If $\models$ is persistent and M satisfies* (VR)*, then M satisfies* (P)*.*

Proof. Suppose the hypothesis, and consider any $\phi \in \textsc{cpl}$ such that $s \not\models \phi$. By (VR), $s \circ \phi = s \cap [\![\phi]\!]$. Now consider any $\psi \in L^+$ such that $s \models \psi$. We have to show that $s \circ \phi \models \psi$. But since $s \circ \phi = s \cap [\![\phi]\!]$, it follows that $s \circ \phi \subseteq s$. Whence it follows by the persistence of $\models$ that $s \circ \phi \models \psi$. Thus, $\{\psi : s \models \psi\} \subseteq \{\psi : s \circ \phi \models \psi\}$, as required. $\square$

So, in the presence of a near-platitudinous constraint on vacuous 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

[14]See Veltman (1985, 1996). Persitence and non-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 non-monotonic in the normal sense.

to $\neg\Diamond\neg p$). So we have reason to drop (P) apart from Fuhrmanns Impossibility Theorem. But, given that (VR) and persistence entail (P), we cannot jettison (P) without also abandoning either (VR) or the persistence of our chosen consequence relation.

This may seem an uncomfortable choice since (VR) seems utterly obvious and persistence generally makes for well-behaved notions of consequence. We have reason, however, to expect that rational epistemic commitment cannot be modeled by a persistent consequence relation. The argument parallels the one above: I can be committed to $\Diamond\neg p$ in a state which commits me to neither p nor $\neg p$—say the state $s = \{w_1, w_2\}$ where p is true at w_1 but not at w_2. But, were I instead in the state $s' = \{w_1\}$ which is properly contained in s I clearly ought not be committed to $\Diamond\neg p$ for it is just not so that in the narrower state s' it might be that not-p. But this is squarely at odds with the mandates placed on us by a persistent consequence relation.[15] This phenomenon surrounding might does a tremendous amount of work when it comes to the Fuhrmann result—we will see it twice more.

The diagnosis is now coming into view: we want to drop (P) from the list of restrictions placed on revision models since it carries absurd predictions about what epistemic commitments agents ought to have. But (VR) plus a persistent consequence relation forces (P) upon us. It is hard to muster much enthusiasm for rejecting (VR). Luckily, we do not have to: persistent consequence relations are flatly implausible as models of rational epistemic commitment in the context of epistemic modals—and for just the reasons that (P) proved implausible. So persistence is the real problem, and we are better off without it.

How does this diagnosis relate to Fuhrmann's? For Fuhrmann the result turns crucially on epistemic weakenings and the Levi Identity. In particular, his argument is that (P) is entailed by the Levi Identity plus a constraint on vacuous weakenings; (P) is the culprit, he thinks, and so he takes aim at the Levi Identity. Fuhrmann and I agree on the culprit, but disagree on the mastermind, for in the version I have put before us the Levi Identity makes no appearance at all. To make a comparison easier, we need a suitable version of the Levi Identity and at least a thin notion of epistemic weakenings in terms of our states. Once we have done this it will turn out that Fuhrmann's diagnosis tacitly assumes that the background consequence relation is persistent—the Levi Identity and equivocations on 'weakening' are red herrings.

Definition 5.3 (Levi Identity)**.** Let $M = \langle I, \models, \circ \rangle$ be a revision model for L^+. Let $\downarrow$ be a state *downdating* function, taking states and formulas in CPL to states such that $s \subseteq s \downarrow \phi$. M satisfies the Levi Identity (LI) iff: for any $s \in I$ and $\phi \in$ CPL, $s \circ \phi = (s \downarrow \neg\phi) \cap [\![\phi]\!]$.

Definition 5.4 (Vacuous Weakening)**.** Let $M = \langle I, \models, \circ \rangle$ be a revision model for L^+. M satsifies *vacuity of weakenings* (VW) iff for any $s \in I$ and $\phi \in$ CPL: if $s \not\models \phi$ then $s \downarrow \phi = s$.

(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

[15] 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).

one might just as well go the other direction, starting with a revision function and defining weakenings in terms of it. This direction for the link between revision and weakenings is known, in the belief revision lore, as the Harper Identity:

Definition 5.5 (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 (LI) iff: for any $s \in I$ and $\phi \in$ CPL, $s \downarrow \phi = s \cup (s \circ \neg\phi)$.

There is a close, and well-known, relationship between Fuhrmann's cast of characters and mine: assuming (LI), (VW) implies (VR); and assuming (HI), (VR) implies (VW).[16] So there is a broad and natural sense in which our constraints on vacuous weakenings and vacuous revisions come to just the same thing. For completeness, I will reproduce the facts here:

Proposition 5.2. *Let $M = \langle I, \models, \circ \rangle$ be a revision model for L^+. If M satisfies* (LI) *then it satisfies* (VW) *only if it satisfies* (VR)*. And if M satisfies* (HI) *then it satisfies* (VR) *only if it satisfies* (VW)*.*

Proof. Assume M satisfies (LI) and (VW). Suppose $s \not\models \neg\phi$, and consider $s \circ \phi$. By (LI), $s \circ \phi = (s \downarrow \neg\phi) \cap [\![\phi]\!]$. But, by (VW), $(s \downarrow \neg\phi) = s$, and so $s \circ \phi = s \cap [\![\phi]\!]$.

Now assume that M satisfies (HI) and (VR). Suppose $s \not\models \neg\phi$, and consider $s \downarrow \neg\phi$. By (HI), $s \downarrow \neg\phi = s \cup (s \circ \phi)$. But, by (VR), $s \circ \phi = s \cap [\![\phi]\!]$, and so $s \downarrow \neg\phi = s$. □

Given this broad equivalence, Fuhrmann's way out will hold up to scrutiny only if—as he claims—the Levi Identity and the appropriate vacuity constraint on weakenings entail the relevant preservation condition. And, indeed, as properties of AGM-style belief sets the entailment claimed exists.[17] The trouble is that once we have brought the properties of the relation of epistemic consequence out of the background and into the fore, it is clear that the entailment just does not go through: (LI) and (VW) just do not entail (P). What is needed, in addition to the Levi Identity and a vacuity constrain on weakenings is the assumption that the relation of epistemic commitment specified in the model is persistent.

Corollary 5.3. *Let $M = \langle I, \models, \circ \rangle$ be any revision model for L^+. If $\models$ is persistent and M satisfies* (LI) *and* (VW)*, then M satisfies* (P)*.*

Proof. This follows immediately from Propositions 5.1 and 5.2. □

That the assumption that the consequence relation is persistent is critical is easy to see by verifying that the relation $\models^+$ from Example 1 is not persistent and with this concept of commitment (LI) and (VW) just do not entail (P), by reasoning which is by now familiar.[18] For 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 (VW),

---

[16]See, e.g., Gärdenfors (1988); Grove (1988).

[17]The relevant version of the Levi Identity is just this: $K \star \phi = (K \sim \neg\phi) + \phi$, where $K + \phi = \mathrm{Cn}(K \cup \{\phi\})$; the relevant vacuity constraint on weakenings (contractions) is just that if $\phi \notin K$ then $K \sim \phi = K$. These two conditions are indeed enough to get us ($\mathrm{P_{AGM}}$). Proof: Suppose $\neg\phi \notin K$, and consider $K \star \phi$. By the Levi Identity $K \star \phi = (K \sim \neg\phi) + \phi$ and the vacuity constraint gives us that $K \sim \neg\phi = K$. Thus $K \star \phi = K + \phi$, and this will clearly contain all the prior beliefs in K since Cn is monotonic.

[18]Given a state s, we first formed a stable autoepistemic theory K_s, and then defined $\models^+$ as: $s \models^+ \phi$ iff $\phi \in K_s$.

$s \circ p = s \cap [\![p]\!] = \{w_1\}$. Here we have a case where $s \models^+ \Diamond\neg p$, but $s \circ p \not\models^+ \Diamond\neg p$, violating (P).[19]

Fuhrmann and I agree that (P) is to blame in the impossibility result; we disagree about the ultimate source of this culpability. He thinks it lies in the Levi Identity and two notions of epistemic weakening. But, I have argued, this is a mistake. Properties of the consequence relation assumed in our models may well have non-trivial effects for those models. This is a reason to bring those relations out of the background and into plain view. And, in the Fuhrmann Impossibility Theorem, doing this lets us make some progress. We can see that 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

A diagnosis of the difficulty in the Fuhrmann Impossibility Theorem does not solve the problem of epistemic change in the context of might. To turn our diagnosis into a solution to the problem, we need to produce a non-trivial revision model which has the properties we want.[20] And this begins with the construction of a reasonable, and more subtle, concept of epistemic commitment. The most obvious model for epistemic commitment is the "propositional containment" analysis. The idea is familiar: a state (a set of possibilities) commits an agent to ϕ just in case that state is included in the set of ϕ-worlds. This is, of course, exactly the model for CPL consequence. So the natural path would seem to be to generalize the notion of a "ϕ-world" so that we can pick out, in addition to ϕ-worlds, the worlds where ϕ might be the case—the set of $\Diamond\phi$-worlds. And, having done that, we would need only to inspect whether a given state is included in the set of $\Diamond\phi$-worlds to see if that state indeed carries that commitment.

But as soon as we start this process of looking for $\Diamond\phi$-worlds that can be coupled with the propositional containment analysis of epistemic commitment we run into some apparent trouble. There seems to be no set answering to the description "set of worlds where ϕ might be the case" for a state to be included in. For suppose otherwise. Consider a scenario in which my state contains just two possibilities,

[19]It is not obvious, at first blush, how it is that persistence or monotonicity has been smuggled in when we are viewing the problem AGM-wise. In a way, that is just my point in advocating looking at the problem in the way I have been—for the two frameworks are obviously equivalent and yet in one guise the properties of the consequence relation are obviously important to getting the diagnosis right. But, in any case, the assumption was there in the AGM-wise version all along: even though Poss is not a monotone operator, Cn^+ clearly is a monotone operator over Poss-closed belief sets. And this is the background notion of commitment that is operative—it is what is implicated, e.g., by ($\mathrm{NT_{AGM}}$), ($\mathrm{S_{AGM}}$), ($\mathrm{C_{AGM}}$), and ($\mathrm{P_{AGM}}$).

[20]Strictly speaking, this will really only give us the *beginnings* of a solution since we are still only considering revision with respect to non-modal inputs. Eventually, for a full solution, this assumption will need relaxing as well.

w_1 and w_2, and that the first is a raining-world and the second is not.[21] Then, intuitively speaking, I am committed in view of my other beliefs to It might be raining. So, given the propositional containment story about commitment, both of these worlds must be in the set of might-raining-worlds. In particular, w_2 must be. But then if I learn—in a simple-minded update sense of 'learn'—that it is not raining, my new state containing only w_2, intuitively, does not commit me to It might be raining.[22] But in that case w_2 had better not be in the set of might-raining-worlds, which is rather unfortunate since we already decided that w_2 *is* a might-raining-world. Contradiction.

Now, what this simple argument might *not* 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.[23]

The intuition behind the propositional containment analysis of epistemic commitment is something like the following. A state commits you to ϕ just in case learning that ϕ in that state would not make any difference to you with respect to ϕ's being the case—the information that ϕ conveys is already present in the state. And the propositional containment analysis captures this intuition well enough so long as ϕ contains no modalities. But it is certainly not the only way of making this intuition precise. An equally appealing way would be to treat epistemic commitment as a fixed-point: s commits you to ϕ iff adding the information conveyed by ϕ to s forces no change whatever to s—i.e., just in case the information conveyed by ϕ is already present in s. In the classical case, these come to the same thing since "adding the information conveyed by ϕ" amounts to just intersecting a state with the set of ϕ-worlds.

I want to begin with this fixed-point intuition, and try to get straight about what sort of commitment relation we want when it comes to might. But the classical way of making it precise will be of no more help to us than is the propositional containment story for it cashes out the key bit in terms of intersecting with a set of ϕ-worlds. And if ϕ is a modal, then we are right back where we started. So we need a more general way of thinking about adding information conveyed by bits of our regimented language. I will offer a way below—not entirely new, but new to this way of thinking about problems in the dynamics of belief—and then use it to bootstrap our way into a non-persistent concept of epistemic commitment.

Up to this point I have been content to follow Levi and Fuhrmann in their focus on epistemic modalities which are restricted to depth at most one. But this

---

[21]If you prefer a more sophisticated propositional containment story according to which it is the "optimal set" of possibilities that must be included in the set of ϕ-worlds, then just assume that your optimal set includes just w_1 and w_2. This simple argument is unaffected by this twist.

[22]Simple-minded updates are meant to correspond to the simplest kind of epistemic change in which new information is added without any safeguards to prevent collapse into inconsistency. AGM expansion is one such operator, as is set intersection of a prior state with a non-modal content $[\![\phi]\!]$, as is the update function I will define below. This sense of 'update' is quite different from, and should not be confused with, the sense which is at issue in Katsuno and Mendelzon (1992).

[23]This is analogous to Gibbard's view about indicative conditionals (Gibbard, 1981). See also Gillies (2004a).

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 CPL and such that if $\phi, \psi \in L^\diamond$ then $\neg\phi, (\phi \wedge \psi), \diamond\phi \in L^\diamond$.

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. I have also followed the traditional belief revision lore up to this point in assuming that updating a state (which corresponds to belief-set expansion in the AGM framework) induces only the trivial dynamic of set-intersection: s updated with ϕ is just $s \cap [\![\phi]\!]$. But, in general, this will not do. First, if ϕ contains an epistemic modality, say if $\phi = \diamond p$, then there just is no set answering to the description $[\![\diamond p]\!]$. Second, if we hope to bootstrap our way into a suitable concept of epistemic commitment by first getting straight about updates, then we need to extend the notion of update to cover the modalities in a satisfactory way. And, of course, we want the new notion of update to agree with the old where it can.

The general idea is alarmingly simple. Let's 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).[24] Generalizing things a bit:

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

(1) $s \uparrow p = \{w \in s : w \in [\![p]\!]\}$
(2) $s \uparrow \neg\phi = s \setminus (s \uparrow \phi)$
(3) $s \uparrow (\phi \wedge \psi) = (s \uparrow \phi) \uparrow \psi$
(4) $s \uparrow \diamond\phi = \{w \in s : s \uparrow \phi \neq \emptyset\}$

Updating a state with an atom p just eliminates from that state all not-p possibilities. Negation is just set subtraction: updating with $\neg\phi$ eliminates from a state just those possibilities which would survive an update with ϕ. 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 $[\![\cdot]\!]$ has, and so when we use it to define a commitment relation we will get something much better.

[24]This "test" behavior of modals is a motivating intuition for Update Semantics. See Veltman (1996); van der Does, *et al.* (1997); Groenendijk, Stokhof, and Veltman (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).

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 $w_1 \in [\![p]\!]$ but $w_2 \notin [\![p]\!]$. 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, $s \uparrow \Diamond p \neq \bigcup_{w \in s} \{w\} \uparrow \Diamond p$. □

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, $[\![\cdot]\!]$ 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 $\Diamond\phi$-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 $[\![\cdot]\!]$ is insensitive to. Since we are thinking of commitment as a fixed-point of the update function, the non-distributivity of the latter will have dynamic effects on former.

Definition 6.4 (Commitment)**.** Let ϕ, ψ be any formulas in $L^{\Diamond}$. Then:

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

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.[25]

Definition 6.5. A relation $\models \subseteq I \times L^{\Diamond}$ is *completely reflective* iff for any $s \in I$:

(1) for any $\phi \in$ CPL, $s \models \phi$ iff $s \subseteq [\![\phi]\!]$;
(2) for any $\phi \in L^{\Diamond}$, if $s \models \phi$ then $s \models \Box\phi$;
(3) for any $\phi \in L^{\Diamond}$, if $s \not\models \phi$ then $s \models \Diamond\neg\phi$.

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

Proof. First, note that for any $\phi \in$ CPL, $s \uparrow \phi = s \cap [\![\phi]\!]$. Condition (1) is then easy to establish: for, $s \Vdash \phi$ iff $s \uparrow \phi = s$. Since $s \uparrow \phi = s \cap [\![\phi]\!]$ in this case, we have that $s \uparrow \phi = s$ iff $s \cap [\![\phi]\!] = s$, i.e. iff $s \subseteq [\![\phi]\!]$.

(2) Suppose $s \Vdash \phi$. Thus $s \uparrow \phi = s$. But note that $s \Vdash \Box\phi$ iff $s \uparrow \Box\phi = s$—i.e., iff $\{w \in s : s \uparrow \phi = s\} = s$, and this iff $s \uparrow \phi = s$, as required.

(3) Suppose $s \nVdash \Diamond\neg\phi$. (We need to see that $s \Vdash \phi$.) So $s \uparrow \Diamond\neg\phi \neq s$. That is, $\{w \in s : s \uparrow \neg\phi \neq \emptyset\} \neq s$, which implies that $s \uparrow \neg\phi = \emptyset$. Whence it follows that $(s \setminus (s \uparrow \phi)) = \emptyset$, and hence $s \uparrow \phi = s$. And so $s \Vdash \phi$. □

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

[25] 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.

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 $w \in [\![p]\!]$ and $w_2 \notin [\![p]\!]$. 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 $s' \nVdash \Diamond p$ 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 $\phi \in L^{\Diamond}$ we have that $\phi \nVdash \phi$. For let $\phi = \Diamond p \wedge \neg p$ and consider an s like that in the proof above. But, you might say, surely this is a mistake!

Surely 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 which 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: $\phi \in L^{\Diamond}$ is *consistent* iff there is a state $s \in I$ such that $s \uparrow \phi \neq \emptyset$.[26] 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 opened 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 isnt raining at all. Assuming that we are content enough to treat sequences of sentences as cases of intersentential conjunction, then we have just the sort of situation we are after. 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 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 $\phi \in L^{\Diamond}$ is *cohesive* iff there is a non-empty state s such that $s \Vdash \phi$, i.e. such that $s \uparrow \phi = \phi$.[27] And so something like $\Diamond p \wedge \neg p$ isnt really the kind of thing that an agent can have as a commitment.

[26]This captures the classical concept of consistency for the non-modal fragment as a special case: for if $\phi \in$ CPL then there is an s such that $s \uparrow \phi \neq \emptyset$ iff $[\![\phi]\!] \neq \emptyset$.

[27]This is the same property that Groenendijk, Stokhof, and Veltman (1996) call "coherence". It is easy to see that cohesiveness implies consistency but not vice versa.

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 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. A family of non-trivial models

So far I have argued against the doxastic conservative's (P) by arguing against 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.

What we need is a revision function (or a family of such functions) taking states in I to states that, when paired with our more dynamic commitment relation $\Vdash$, satisfies versions of (NT), (S), (VR), and (C).[28] There is a range of such 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 W is an *s-implausibility ordering* iff for any $w \in s$:

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

Let $\min(\phi, \leq_s)$ be the set of worlds $w \in [\![\phi]\!]$ such that no $w' \in [\![\phi]\!]$ is strictly less s-implausible than w. A revision function $\circ : I \times \textsc{cpl} \to I$ is *broadly conditional* iff for any s and $\phi \in \textsc{cpl}$ there is an s-implausibility ordering $\leq_s$ such that $s \circ \phi = \min(\phi, \leq_s)$.

It is easy to see that revision models based on such broadly conditional revision functions and our dynamic consequence relation satisfy (S), (VR), and (C); provided the spave 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 $\phi \in \textsc{cpl}$ and $s \in I$.*

(1) (S) $s \circ \phi \Vdash \phi$.
(2) (VR) *If* $s \not\Vdash \neg\phi$ *then* $s \circ \phi = s \cap [\![\phi]\!]$.
(3) (C) *If* $[\![\phi]\!] \neq \emptyset$ *then* $s \circ \phi \neq \emptyset$.

[28]Since the models will based upon $\Vdash$ they will, of course, fail to satisfy (P).

(4) (NT) *If W is non-trivial, i.e., if there are two distinct possibilities, then M is non-trivial.*

Proof. For (1), note that $s \circ \phi = \min(\phi, \leq_s)$ and so $s \circ \phi \subseteq [\![\phi]\!]$. Since $\phi \in$ CPL we know that for any s' whatever $s' \uparrow \phi = s' \cap [\![\phi]\!]$, and so clearly $(s \circ \phi) \uparrow \phi = s \circ \phi$, and thus $s \circ \phi \Vdash \phi$.

(2) Suppose $s \nVdash \neg\phi$, for an arbitrary $\phi \in$ CPL. Since $\phi \in$ CPL, this implies that $s \cap [\![\phi]\!] \neq \emptyset$. Now, $s \circ \phi = \min(\phi, \leq_s)$ for some s-implausibility ordering $\leq_s$. We first show that $\min(s, \leq_s) \subseteq s \cap [\![\phi]\!]$. Assume $w \notin s \cap [\![\phi]\!]$. If $w \notin [\![\phi]\!]$, then it follows straightaway that $w \notin \min(\phi, \leq_s)$. So suppose $w \in [\![\phi]\!]$ but $w \notin s$. We need to show that there is a $w' \in [\![\phi]\!]$ such that $w' <_s w$. Since $s \cap [\![\phi]\!] \neq \emptyset$, let $w' \in s \cap [\![\phi]\!]$. By construction of $\leq_s$, it then follows that $w' <_s w$, as required. To see that $s \cap [\![\phi]\!] \subseteq \min(\phi, \leq_s)$ consider an arbitrary $w \in s \cap [\![\phi]\!]$. Again, the construction of $\leq_s$ gives us that w is minimal in $\leq_s$ (since $w \in s$), and so must be in $\min(\phi, \leq_s)$.

(3) Suppose $[\![\phi]\!] \neq \emptyset$. Since, for $\phi \in$ CPL, $s \circ \phi = \min(\phi, \leq_s)$ it is immediate that $s \circ \phi \neq \emptyset$. (And, it doesn't take much to see that, if a state $s \neq \emptyset$, then $\{\phi : s \Vdash \phi\}$ is consistent.)

(4) If there are two possibilities, w_1 and w_2, then there is a non-modal ϕ about which they disagree. Thus, the limiting case when $s = W$ suffices for (NT). □

(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. 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. 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, I want to suggest, is best got at by way of the story I have been telling.

The triviality result we have been considering forces us to re-think some issues about the relationship between the doxastic conservative's preservation condition (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 trivial and limiting cases of revision (revising a state with ϕ when that state carries no prior commitment to $\neg\phi$) reduces to simple-minded updates of the sort codified by AGM-style belief set expansion and set intersection of a state with $[\![\phi]\!]$. Such an assumption straightaway yields a preservation condition, and so has got to go when we have such 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 ϕ as the *revision* by $\Box\phi$, and define the contraction with respect to ϕ as the *revision* by $\Diamond\neg\phi$.

This lesson is not altogether the right one. First, it is just not so that if the limiting case of revision reduces to simple-minded updating we get (P). We saw above that (VR) is not only a near platitude, 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 $\Diamond\phi$. And when we do we will surely want to have that a state revised by $\Diamond\phi$ amounts to a weakening of that state with respect to $\neg\phi$. 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 (VR) to get it.

Now to the second alternative lesson.[29] 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.

[29]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 closure under Poss.

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. 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 the preservation condition is obviously the right move to make, and it was just a mistake to ever think it applied more widely than that. The trouble is, I think, that restricting (P) to CPL in this way drains it of much of its philosophical significance. It was meant, at least by many of its proponents, to capture and codify intuitions that doxastic conservatism is a deeply rational mandate.[30] Without such a substantive interpretation of (P), for instance, Harman's arguments against foundationalist belief revision lose their normative force (Harman, 1984). I am no fan of (P), but I think those who are fans of it 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. That seems a high price to pay. 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.

[30]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., P. Gärdenfors, and D. Makinson (1985). "On the logic of theory change: Partial meet functions for contraction and revision," *Journal of Symbolic Logic* **50**: 510530.

van Benthem, J. (1996). *Exploring Logical Dynamics*. Stanford: CSLI Press.

van der Does, J., W. Groeneveld, and F. Veltman (1997). "An update on 'might'," *Journal of Logic, Language and Information* **6**: 361–381.

Fuhrmann, A. "Reflective modalities and theory change," *Synthese* **81**: 115–134.

Gärdenfors, P. (1986). "Belief revisions and the Ramsey Test for conditionals," *Philosophical Review* **95**: 81–93.

Gärdenfors, P. (1988). *Knowledge in Flux*. Cambridge: MIT Press.

Gibbard, A. (1981). "Two recent theories of conditionals," in W. L. Harper, R. Stalnaker, and G. Pearce (eds.), *Ifs* (Dordrecht: Reidel).

Gillies, A. S. (2001). "A new solution to Moore's paradox," *Philosophical Studies* **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* (Amsterdam: ILLC).

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., M. Stokhof, and F. Veltman (1996). "Coreference and modality," in S. Lappin (ed.), *Handbook of Contemporary Semantic Theory* (Oxford: Blackwell), pp.179–213.

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

Hansson, S. O. (1999). *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Doredrecht: Kluwer.

Harman, G. (1984). "Positive versus negative undermining in belief revision," *Noûs* **18**: 39–49.

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

Katsuno, H. and A. O. Mendelzon (1992). "On the difference between updating a knowledge base and revising it," in: P. Gärdenfors (ed.), *Belief Revision* (Cambridge: Cambridge University Press), pp.183–203.

Kratzer, A. (1977). "What 'must' and 'can' must and can mean," *Linguistics and Philosophy* **1**: 337–355.

Levi, I. (1979). "Serious possibility," in: E. Saarinen, R. Hilpinen, I. Niiniluoto, and M. Provence Hintikka (eds.), *Essays in Honour of Jaakko Hintikka* (Dordrecht: D. Reidel), pp.219–236. Reprinted in: I. Levi, *Decisions and Revisions* (Cambridge: Cambridge University Press), 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*. New York: Cambridge University Press.

Moore, R. C. (1985). "Semantical considerations on nonmonotonic logic," *Artificial Intelligence* **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*. New York: Oxford University Press.

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* (Dordrecht: Kluwer Academic Publishers), pp.105–134.

Stalnaker, R. (1993). “A note on nonmonotonic modal logic,” *Artificial Intelligence* **64**: 183–196.

Veltman, F. (1985). *Logics for Conditionals*. Ph.D. dissertation, Department of Philosophy: University of Amsterdam.

Veltman, F. (1996). “Defaults in update semantics,” *Journal of Philosophical Logic* **25**: 221–261.

Department of Philosophy, University of Michigan, Ann Arbor
E-mail address: thony@umich.edu