NOÛS 38:4 (2004) 585–616

*Epistemic Conditionals and Conditional Epistemics**

ANTHONY S. GILLIES
University of Michigan, Ann Arbor

1. Introduction: Two Problems of "Ordinary" Indicatives

There has been a murder in the mansion. There are, and we all know that this is so, only three suspects: the gardener, the driver, and the butler. We believe, and we have our reasons for so believing, that if the gardener did not do it, then the driver did. We share belief in this particular conditional, and this conditional is of a particular kind. It is what Quine calls an "ordinary indicative" conditional (so-called because such conditionals are typically expressed in the indicative mood). We could, just as well, call them ordinary *epistemic* conditionals since, intuitively anyway, they express some conditional connection about epistemic possibilities. Either way, these "ordinary" conditionals have proved to be rather extraordinary.[1]

There are two kinds of obstacles to a satisfying theory about indicatives, and these two classes of obstacles mirror a distinction within the class of epistemic conditionals. A plausible pragmatic constraint on indicatives is that they carry a presupposition that their antecedents might be true. Let us call indicatives which meet this presupposition—i.e., indicatives whose antecedents are consistent with our picture of the world—*open indicatives* or *open epistemic conditionals*. But the pragmatic constraint on indicatives can be cancelled; some indicatives have antecedents the belief in which would require us to revise our epistemic states. The well-known example due to Ernest Adams (1975) illustrates this fact: I believe that Oswald killed Kennedy, but I also believe the indicative conditional If Oswald didn't kill Kennedy, then someone else did.[2] Let us call indicatives of this latter sort *belief-contravening indicatives* or *belief-contravening epistemic conditionals*.

The first kind of obstacle to a comprehensive understanding of indicatives is specific to the belief-contravening indicatives. Frank Ramsey (1931) suggested, in a famous footnote, that the basic structure to reasoning about indicatives goes by way of a seemingly simple test: "If two people are arguing 'If p will q?' and are both in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q ..." (p. 247n1). Suppose that in some sense or other, analyzing and accepting an indicative does require this hypothetical endorsement of the conditional's antecedent. Now, if belief in that antecedent is incompatible with what we take ourselves to know (as is the case with belief-contravening epistemic conditionals), then to keep such conditionals from being trivial our story about indicatives requires saying what the landscape of rational belief change is. In particular, it requires a theory of rational *revision* of epistemic states. This is the *Problem of Epistemic Change*, a serious epistemological problem very much deserving of our interest and attention.[3] We will here, however, turn our attention away from this general problem in order to focus more clearly on the structure of epistemic conditionals. Accordingly we will be shelving all discussion of belief-contravening indicatives.

Open epistemic conditionals are indicatives belief in the antecedents of which is, in some sense, compatible with our epistemic state. On the surface, giving an analysis of the rational kinematics of these seems a much simpler challenge to meet. But open indicatives are not without their own serious worries. This is the second kind of obstacle which threatens our understanding of epistemic conditionals, and what will be my main focus here.

The main puzzle I want to address is a problem which I call the *Reduction Problem*. This is the problem of giving an account of the semantic structure of the ordinary locution If p, then q without reducing it to the truth-functional analysis of material implication according to which it means either not-p or q.[4] Understood in this way, the Reduction Problem looks like a problem for semantics. However, in Section 2 I will argue that it is also a problem for epistemology, and perhaps threatens to do more damage in this new guise. It will turn out to pose difficulties for the epistemic commitments of rational agents who have beliefs which are expressed by indicative conditionals. Plausible assumptions on epistemic commitment seem to lead inexorably to the result that believing indicatives is no different from believing material conditionals. So, for those of us who want to resist this conclusion, we must find a way out.

To motivate my own favored escape route, in Section 3 I look at some of the other options in logical space. None of these solutions—solutions favored, for example, by Frank Jackson (1979), Robert Stalnaker (1975), and Vann McGee (1985) as ways out of the Reduction Problem *qua* semantic problem—will do as solutions to the Reduction Problem *qua* epistemic problem. The central problem is that we need, I will argue, a picture of epistemic commitment which is sensitive to informational context. With the philosophical foundations in place, Sections 4 and 5 contain my positive theory about conditional epistemics. To anticipate, a dynamic view of rational epistemic commitment which takes information seriously gets us everything we want in a theory of open indicatives: something like the Ramsey Test undergirds how we think about them, intuitively rational inference patterns are sanctioned by the theory as rational, and belief in indicatives is distinct from believing material conditionals.

2. Reduction Problem

In talking about what my epistemic commitments are given some specified information—as opposed to what my actual beliefs are—we are talking about what an ideally rational agent with that information would conclude. What is rationally warranted, moreover, varies with the complexity of the epistemic environment. Our world is a complicated one: we have to be able to form beliefs defeasibly, and this is the rational thing to do in a world like ours. But assessing the structure of rational epistemic commitments in simpler environments has its philosophical value, too. The structure of epistemic commitment which is based on and derived from open indicatives, it seems to me, is just such a case.

Open indicatives are just those indicatives belief in the antecedents of which is compatible with what we take ourselves to know. So, in assessing epistemic commitments with respect to open indicatives, let us take the most idealized stance possible: agents never revise their picture of the world, and all reasons are non-defeasible. It will turn out that even in this ideal case scenario our assumptions on epistemic commitment will get us into trouble—we will have cause to restructure our rough-and-ready sense of epistemic commitment, even in the simplest case.

The puzzle is this. There are a number of plausible assumptions about what it is rational to believe based on epistemic conditionals—assumptions which we are all only too happy to agree to—which jointly entail that epistemic conditionals are indistinguishable from the material conditional analysis familiar from our days in introductory logic courses.

To state the problem precisely, suppose we have a background propositional language $\mathcal{L}_{\Rightarrow}$ with the set of propositional atoms $\mathcal{A} = \{p, q, \ldots\}$, and closed under the truth-functional connectives for negation ($\neg$), conjunction ($\wedge$), and material implication ($\rightarrow$). Suppose further that $\mathcal{L}_{\Rightarrow}$ is closed under the indicative if-then ($\Rightarrow$). On occassion we will have use for the fragment of $\mathcal{L}_{\Rightarrow}$ which contains no formulas with any occurrences of the epistemic conditional $\Rightarrow$. Call this fragment $\mathcal{L}_0$.[5] Then we can state four assumptions on the concept of rational epistemic commitment.

First, suppose that $\mathcal{L}_{\Rightarrow}$ is adequate to express the epistemic commitments that we are interested in. Next, let us assume that we have a sort of deduction theorem for epistemic commitment: ϕ is sufficient to commit an agent to ψ iff that agent is committed outright to the epistemic conditional $\phi \Rightarrow \psi$. Third, the set of rational epistemic commitments of an agent (at a given time) is closed under three intuitively rational inference rules: modus ponens, exportation, and what Stalnaker calls the "direct argument" for the indicative. Take each of these in turn. If I am committed to ϕ and I am committed to the indicative conditional $\phi \Rightarrow \psi$, then I am on that basis committed to ψ—rational epistemic commitment is closed under modus ponens. Likewise for exportation—belief in (i.e., acceptance of) $(\phi \wedge \psi) \Rightarrow \chi$ is enough to commit an agent to the right-nested if-then $\phi \Rightarrow (\psi \Rightarrow \chi)$—and the direct argument—belief in (i.e., acceptance of) the disjunction $\phi \vee \psi$ is enough to commit an agent to the indicative $\neg\phi \Rightarrow \psi$ (at least for ϕ and ψ not containing occurrences of '$\Rightarrow$').

The fourth and final assumption on rational epistemic commitment constrains the properties of the relation "...epistemically commits an agent to...". Straightaway we know that however this relation behaves, it must be that holding a set of beliefs commits us to holding each of the individual beliefs that makes up that set. Moreover, rational epistemic commitment is meant to pick out something like a "justified in the limit" notion of rationality. One way of codifying this is to say that an agent is automatically committed to every tautology. Now, we are confining ourselves to open indicatives, and correspondingly to the most idealized epistemic environment. So we are pretending that the only reasons agents ever have are conclusive, non-defeasible reasons. And thus agents never encounter defeaters of any sort. This means that if a set of beliefs commits an agent to ϕ and then the agent picks up some new beliefs, then the agent is still committed to ϕ. These three properties correspond to three natural properties of consequence relations. Let $\Gamma \models \phi$ symbolize the fact that ϕ is a consequence of Γ. The fourth assumption is just that epistemic commitment is a classical consequence relation, in the sense of Definition 2.1.

Definition 2.1 (Classical Consequence). A relation $\models_{CL}$ on sets of formulas and formulas is a *classical consequence relation* for a language $\mathcal{L}$ iff (subscript omitted when unambiguous):

(1) $\emptyset \models \phi$ for any tautology ϕ. In this case we write $\models \phi$;
(2) $\Gamma \models \gamma$ for each $\gamma \in \Gamma$ (inclusion);
(3) If $\Gamma \models \phi$ and $\Gamma \subseteq \Delta$, then $\Delta \models \phi$ (monotonicity).

Here, then, is our set of assumptions on our idealized concept of rational epistemic commitment.

Assumption 1. Epistemic commitments can be any set of formulas of $\mathcal{L}_{\Rightarrow}$, including open indicatives, and all epistemic commitments are expressible in $\mathcal{L}_{\Rightarrow}$.
Assumption 2. The information that ϕ is enough to epistemically commit an agent to ψ iff that agent is committed to the indicative $\phi \Rightarrow \psi$ outright.
Assumption 3. The set of rational epistemic commitments is closed under modus ponens (MP), exportation (EXP), and the direct argument (DA) for $\Rightarrow$.[6]

Assumption 4. "…epistemically commits an agent to…" is a classical consequence relation for $\mathcal{L}_{\Rightarrow}$.

But we cannot have all of this. If indicatives are at least as strong as the material conditional, and weaker than strict implication, then there is a standard result which shows that in the presence of (MP) and (EXP), indicatives are indistinguishable from material conditionals.[7] The structure of epistemic conditionals, apparently anyway, guarantees that the commitments they carry are all and only the commitments carried by material conditionals.

Proposition 2.1 (Stalnaker 1975; Gibbard 1981; McGee 1985; Veltman 1985). *Let ϕ, ϕ' be any formulas in $\mathcal{L}_{\Rightarrow}$ and suppose $\models$ is a relation of rational epistemic commitment for $\mathcal{L}_{\Rightarrow}$ which meets Assumptions 1–4. Then if ϕ and ϕ' are alike except that some occurrences of '$\Rightarrow$' in ϕ are replaced by '$\rightarrow$' in ϕ', $\phi \models \phi'$ and $\phi' \models \phi$.*

So far, Proposition 2.1 is only a surprising, and perhaps disappointing, result. No deep philosophical puzzle yet. But there is a serious puzzle to be had—the rational kinematics of believing indicatives must be different from that of believing material conditionals. The next example aims to establish that. So—this is where the Reduction Problem earns its title—at least one of our assumptions needs to be rejected; but each assumption is intuitively something we want from our theory of epistemic conditionals.

Now, to the example I have in mind:

Example 1. Suppose, as before, that there has been a murder in the mansion. The mansion staff can be partitioned into the *grounds staff* (people who work outside the mansion) and the *house staff* (people who work inside the mansion proper). Again, the three suspects are the driver (grounds staff), the butler (house staff), and the gardener (grounds staff). You are the lead investigator, and your young assistant (who, in fact, is but an apprentice) has been collecting clues and reports back to you. She has collected this one clue: the butler has an airtight alibi. Your assistant thus decrees

(1) Therefore: if a member of the grounds staff did it, then it was the driver.

Being a seasoned inspector, you disagree:

(2) It's not so that if a member of the grounds staff did it, then it was the driver.
After all, it might still be the gardener who did it.

The truth-functional analysis of conditional statements (i.e., the material conditional) treats an indicative If *p*, then *q* as equivalent to the disjunction not-*p* or *q*. This cannot be right. If the truth-functional analysis were right, then your denial in (2) would commit you to accepting that a member of the grounds staff did it and it was definitely not the driver. This is too strong a commitment, and would render your reply unwarranted. But, given your information, (2) is not only the reasonable thing to *say*, it is the rational thing to *believe*—and believing that a member of the grounds staff did it and it was definitely not the driver would be decidedly irrational. So the material conditional analysis cannot be right. Being epistemically committed to an indicative If *p*, then *q* is not the same thing as being epistemically committed to the disjunction not-*p* or *q*.

So the Reduction Problem really is a *problem*. Moreover, it is a serious *epistemological* problem. Stalnaker, Gibbard, McGee, and Veltman each prove a version of Proposition 2.1 as a fact about the semantics of indicatives, and address the Reduction Problem as a problem somewhere near the semantics-pragmatics interface.[8] But, as I have stated it here, this is a result about the epistemics of indicatives, and I like to think of the Reduction Problem as a problem at the semantics-epistemics interface. There is, I think, typically a routine way to turn most problems of semantics into problems for epistemology; this is because the rough-and-ready concept picked out by the epistemic locution "...epistemically commits us to..." is the same as that picked out by the semantic locution "...intuitively entails...". The assumptions about semantic entailments needed to prove Proposition 2.1 as a semantic fact in the way that Stalnaker, Gibbard, McGee, and Veltman do correspond in a straightforward way to our four assumptions about rational epistemic commitment. But there is philosophical ground to be gained by thinking of the Reduction Problem as an epistemic problem. I will argue in the next section that some of the escape routes to the Problem that might seem plausible when we are thinking of it as a semantics puzzle are simply not live options when we are thinking of it as an epistemic problem.

3. Escape Routes

There are three general sorts of reaction to the Reduction Problem that have been proposed.[9] One can just insist that it is not a problem at all, and explain away apparent data by appeal to certain pragmatic features of sincere utterance. This is the route endorsed by holders of the "strict equivalence thesis" according to which the indicative is not distinct from the material conditional. Frank Jackson (1979) and David Lewis (1976) belong to this camp. A related way out—Stalnaker's (1975) way out—is to deny that the desired inferences for the indicative are semantically valid but insist that they are pragmatically reasonable. A third option is to just reject some of the desired inference rules as anything like intuitively valid. This is McGee's (1985) basic strategy. None of these routes, or so I will argue, is open when we consider the Reduction Problem as a problem in the epistemics of indicatives.

The first two kinds of options both involve some kind of retreat to the pragmatics of indicatives. These kinds of moves seem plausible enough when we are thinking of the puzzle as a problem at the semantics-pragmatics interface. But this move to pragmatics seems to be wide of the mark if the problem we are interested in is an epistemological one. These kinds of solutions distinguish between what is valid (a semantic notion) and what is felicitous (a pragmatic notion); ϕ might entail ψ even though an utterance of ψ is not always felicitous in the presence of ϕ. This sort of distinction is, no doubt, crucial for understanding language use. But our concern is an epistemic one—we are concerned only with what the rational epistemic commitments of agents are, quite apart from what they may or may not say about those commitments. If holding ϕ commits us to ψ, then if we accept that ϕ it is clear that ψ is one of our epistemic commitments regardless of the pragmatic facts about potential utterances of ψ. So positive pragmatic status (whether it be warranted assertability or pragmatic reasonableness) is not necessary for epistemic commitment. So explaining away troublesome data for the strict equivalence thesis by appeals to negative pragmatic status will not do.

Nor is it sufficient for epistemic commitment. Both Stalnakerian pragmatic reasonableness and warranted assertability are characteristically non-monotonic affairs, allowing cancellations and the like. But that then places such a metric squarely outside of the scope of the relations we have opted to focus attention on—we are ignoring defeasibility of this kind pretending that information simply aggregates in a monotonic way. So explaining away the intuition that certain inference patterns seem rational by appeals to positive pragmatic status will not do.[10]

That leaves us with the last escape route. McGee sees the Reduction Problem as pitting modus ponens against the law of exportation; we cannot, he argues, have both. Exportation (on ordinary indicatives) seems unassailable. If I believe

(3) If the butler is out and the gardener is out, then the mansion is empty

then I am *ipso facto* committed to

(4) If the butler is out, then
if the gardener is out then the mansion is empty.

It is rational to hold conditionals like (4) on the basis of conditionals like (3), and our story about indicatives needs to reflect that.

While exportation seems unassailable, McGee has several counterexamples to modus ponens. Suppose, again, that there has been a murder in the mansion. The mansion staff is made up of the *house staff* (butler, maids, etc.) and the *grounds staff* (driver, gardener, security, etc.). The only three suspects are the butler, the gardener, and the driver.

Example 2. Suppose that the gardener almost certainly did it, it is possible but unlikely that the butler did it, and the driver almost certainly did not do it.[11] We can partially express this ordering on the likely culprit by the complex conditional:

(5) If a member of the grounds staff is the culprit,
then if it is not the gardener who is guilty, the driver is.

And, in fact, since the gardener almost certainly did it, we have every reason to think the antecedent of (5) is true:

(6) A member of the grounds staff is the culprit.

(5) and (6) would seem to entail, by modus ponens, that we also hold that

(7) If it is not the gardener who is guilty, the driver is.

But this is not at all what our information supports.

McGee's view is that the Reduction Problem poses a dilemma (he grants that indicatives and the material conditional must be distinct): we have to choose between exportation and modus ponens. Exportation seems unassailable, and his counterexamples—cases like Example 2—are meant to cast serious doubt on the status of modus ponens with respect to rational epistemic commitments. All of McGee's counterexamples have this essential form. Modus ponens on a complex right-nested open epistemic conditional can take us from premises we endorse to a conclusion (which is a simple open epistemic conditional) we do not.

We should pause for just a moment to recall two facts. First, we are here only interested in *open* epistemic conditionals. Second, the Reduction Problem is a problem for this class of conditionals. So if rejecting modus ponens is a viable escape route to the Reduction Problem and rejecting modus ponens is meant to be motivated by McGee-like counterexamples, then the only relevant readings of the conditionals in the counterexamples had better be of the non-revisionary, open kind.

That being said, when first presented with McGee-like counterexamples to modus ponens, one feels like the victim of some kind of trick.[12] The more clearly in view you keep your commitment to (6), the more the embedded conditional (7) seems a good thing to believe; but when we let (6) slip out of sight just a bit then the embedded conditional seems a crazy thing to believe.

Intuitively, something like this seems to be going on. It is as if we are being asked to entertain certain premises, and decide whether or not we endorse them. But then, when we turn to considering the conclusion, we are asked to pretend we have not endorsed the premises at all—we are asked, in particular, to let our commitment to (6) slip out of view. Our intuitions about our epistemic commitments to (7) are being harvested in an information vacuum. And this sort of slippage is illegitimate. Information matters, and in particular informational context is important in assessing rational epistemic commitments.

Now, it is one thing to point to a metaphor like "information vacuum" and quite another to give an argument that that is what McGee's counterexample turns on. An example with the same content as his above, but one which naturally blocks the information slippage that I allege is at the heart of McGee-like counterexamples, it seems to me, would be sufficient. More precisely, McGee's counterexample requires that we accept three things: (1) a particular right-nested complex if-then, call it P; (2) the antecedent of P; and (3) the denial of the consequent of P. What I need is a right-nested if-then Q such that: (1) Q is intuitively equivalent to P and has the same surface syntax (plus or minus a bit); (2) we accept the antecedent of Q; and (3) the construction of Q is such that it blocks the alleged slippage between premise-beliefs and the conclusion (and so that we also accept the consequent of Q).

Example 3. Suppose we have exactly the same set-up as in the previous case, Example 2. But instead of (5), we have what looks like an equivalent complex conditional:

(8) If a member of the grounds staff is the culprit,
then if he is not the gardener, he is the driver.

Furthermore, we still have reason to think the antecedent of (8) obtains. Our information still supports (6):

A member of the grounds staff is the culprit.

Now consider the consequent of (8):

(9) If he is not the gardener, he is the driver.

Whether or not we are committed to *this* embedded conditional turns on what the referent of he is supposed to be on the most preferred reading. Taken in isolation, there is no fact of the matter: he is serving to refer back to some already-introduced referent in the discourse. If we pretend there is no preceding discourse, then the question is open. But the example clearly provides us with a preceding discourse—most salient, the antecedent (6): A member of the grounds staff is the culprit. In this linguistic context, he in (9) most plausibly refers to the culprit who is a member of the grounds staff—or, more briefly, the grounds-staff-culprit. But now consider (9) with this referent for he:

(10) If the grounds-staff-culprit is not the gardener,
the grounds-staff-culprit is the driver.

And this *is* supported by the information in the example. The only grounds staff suspects are the gardener and the driver, so given that one of them did it (which is what (6) guarantees), and that person was not the gardener, it must have been the driver.

The difference between Example 2 and Example 3 lies solely in what appears to be stylistic variation—the contents of the relevant complex conditionals seem to be the same. But Example 3 has some strategically placed anaphora. In order to evaluate our commitment to (9) we cannot lose sight of the premise-beliefs (8) and (6). And the fact that we endorsed these premise-beliefs makes a difference to how we assess our epistemic commitment with respect to (9).

This is not some odd feature of complex conditionals; it is the epistemic counterpart of a well-known fact in the formal semantics of natural language: interpretation is not only context-dependent, it is also highly context-affecting—anaphoric reference like that in Example 3 is a paradigmatic example of the phenomenon (Groenendijk and Stokhof, 1991). And, in inferences in natural language, anaphora in conclusions can reach back into the premises for their referents. The same is true for rational epistemic commitment. If I accept as true

(11) A woman walks in Vondel Park and talks in Vondel Park.

then I am committed to

(12) So, she walks in Vondel Park.

(11) commits me to (12), and how this works hinges on the fact that the referent for she can get passed from premises to conclusions.[13] Appreciating this, we can see what has gone wrong in McGee-like counterexamples. Information is introduced, but (by fiat) made inaccessible to our assessment of rational epistemic commitment that follows. And that is illegitimate.

It is interesting that Stalnaker's way with the direct argument can be used in much the same way to make McGee-like counterexamples to modus ponens seem less compelling. The direct argument is, while not semantically valid in Stalnaker's treatment, what he calls a *reasonable inference*. Reasonable inference is the pragmatic cousin of semantic validity/entailment. The basic idea is that what is reasonable depends on a *context set*: the sum total of the background information that has been introduced up to that point in a conversation. But we are not allowed to ignore any information in the context set; once information gets introduced, it is part of the background and has a bearing on what inferences are reasonable. The direct argument says that from the disjunction p or q it is reasonable in this technical sense to infer the epistemic conditional If not-p then q.

Suppose we have exactly the same set-up as in Example 2: there has been a murder in the mansion, and we have three suspects.

Example 4. Since there are only two grounds staff suspects at all (the gardener and the driver), the antecedent of (5) is equivalent to a disjunction: either the gardener did it, or the driver did it. So, we have that our information supports the complex conditional

(13) If either the gardener or the driver is the culprit,
then if it is not the gardener who is the culprit, it is the driver.

Again, we believe the antecedent, since we think it likely that the gardener did it:

(14) Either the gardener or the driver is the culprit.

By the direct argument, from (14) we have

(15) If it is not the gardener who is the culprit, it is the driver.

And this just is the consequent of (13).

If in moving from (14) to (15) we are not allowed to ignore the information that (14) introduces into the context set, then the direct argument in this case will commit us to exactly what modus ponens would have. The point for my purposes is just the same as that illustrated in Example 3: in assessing our rational epistemic commitments it is illegitimate to ask that what has been introduced be treated as though it hasn't. And this is what McGee-like counterexamples with respect to open epistemic conditionals ask us to do.

4. Conditional Epistemics

We have good reason to insist that an account of epistemic conditionals and rational epistemic commitment—even in the easy case we are concerned with here where information simply aggregates—must be sensitive to informational context. This sensitivity goes two ways: assessing commitments is sensitive to what information an agent has, and accepting new beliefs changes that information landscape. What I want to do now is turn to giving a precise theory of epistemics—including the epistemics of open indicative conditionals—which trades on taking informational context seriously.

The theory of epistemics I want to give here is related in spirit and in substance to two other approaches. In spirit, it is very much related to a proposal made by Gärdenfors (1984, 1988) for a reconstruction of propositional logic. His idea is to use the rational constraints on information aggregation as a basis for logic, eliminating the need for propositions defined as sets of possible worlds. The project is, much like the one here, one of turning a traditionally semantic problem into an epistemic one. While my story and his share this programmatic bit in common, in formal substance and philosophical point they differ more than they agree.

The theory of conditional epistemics that I want to give here (again, in the simple case where information growth is monotonic) resembles in its formal apparatus the semantic theories in Veltman (1985, 1996).[14] But the philosophical point is considerably different: whereas his interest is primarily semantic, mine is decidedly epistemic.

As before, we shall concentrate on a language $\mathcal{L}_{\Rightarrow}$ of propositional logic. Officially, $\mathcal{L}_{\Rightarrow}$ is the smallest set that contains any propositional atoms that we will need, $\mathcal{A} = \{p, q, \ldots\}$, and is closed under negation ($\neg$), conjunction ($\wedge$), and the epistemic conditional ($\Rightarrow$). $\mathcal{L}_0$ is just the classical fragment of $\mathcal{L}_{\Rightarrow}$.

Theories in philosophical logic often aim at characterizing entailment relations of one sort or another. The data against which these theories are judged are our intuitive judgments of "consequence" or "entailment" for whatever the problem domain is. Typically, entailment is just propositional containment—P entails Q iff the proposition expressed by P is included in the proposition expressed by Q. The bulk of the work, then, goes into pinning down truth-conditions for the P's and Q's as an analysis of propositions. The linchpin in all of this is that what a formula means is just what its truth-conditions are. But why should we buy into this dogma? I say we should not. For present purposes, though, I can settle for the weaker claim that we should not *presume* that meaning is exhausted by truth-conditions. There is, in fact, a *prima facie* case to be made against the received view. If meaning is exhausted by truth-conditions, then (borrowing an example which is due to Barbara Partee)

(16) I have lost ten marbles and found all but one. It might be under the couch.

should have the same meaning as

(17) I have lost ten marbles and found nine of them. #It might be under the couch.

since the first sentence in each case manifestly has the same truth-conditions. But the second sentence in (17) is semantically marked, and not in (16). So, apparently anyway, there is more to meaning than truth-conditions—semantics is richer than propositional content. Hence we should not flatly assume that all interesting semantic structure is truth-conditional structure. Recent work in the formal semantics of natural language has tried to do without the dogma that meaning is exhausted by truth-conditions. Veltman explains the motivating intuition behind dynamic semantics in this way. In dynamic semantics,

> the slogan 'You know the meaning of a sentence if you know conditions under which it is true' is replaced by this one: 'You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it'. (Veltman, 1996, p. 221)

Accordingly, the focus is not on giving a recursive specification of propositional *content*, but on giving a recursive specification of information *change*.

It does not take much squinting to see that Partee's marble sentences may just as well be about epistemic commitment: (16) is an intuitively coherent thing to believe but (17) is not. Similarly, instead of focusing on when indicative conditionals are true, I want to focus instead on how they change the information states of agents who accept them. And since, as I have argued, "...intuitively entails..." picks out the same concept as "...epistemically commits an agent to...", there is a clear sense in which semantic models of epistemic change, when we are restricting ourselves to the case of simple-minded updates in which information is non-defeasible and simply aggregates, just *are* dynamic semantic theories. So, giving a theory about the rational kinematics of formulas in $\mathcal{L}_{\Rightarrow}$ with respect to simple-minded epistemic updating, can be thought of as giving a theory of the *meanings* of formulas of $\mathcal{L}_{\Rightarrow}$. And, echoing the slogan I introduced at the beginning of this paper, this firmly places the Reduction Problem at the semantics–epistemics interface.

Now to the theory of conditional epistemics. As we pick up new information, we have to update our picture of the world, bringing our set of epistemic possibilities into line with what we have learned. This is true even in our idealized environment in which all information is non-defeasible. For present purposes, the information that agents have for epistemic concerns—the agents' *information states* or *acceptance bases*—can just be identified with a set of possible worlds.[15] For our purposes, it is enough that possible worlds (more economically: *possibilities*) decide the truth value of the atomic facts we care about. That is, worlds are functions from atoms to truth-values.[16]

Definition 4.1 (Possibilities and Acceptance Bases). Fix a set $\mathcal{A}$ of atomic formulas. w is a *possibility* iff $w : \mathcal{A} \to \{0, 1\}$. W is the set of such w's. s is an *acceptance base* iff $s \subseteq W$. I is the set of such s's.

There are two limiting cases of acceptance bases. The first is that W itself is an acceptance base. In this case, the agent has not yet accepted anything—no possibilities have yet been ruled out. At the other extreme we have the *absurd* acceptance base—the empty set of possibilities $\emptyset$. This is the result of "learning" inconsistent information; the agent has ruled out too much. A happier, though rarer, state of affairs is when an agent's acceptance base is a singleton. Any epistemic state with a singleton acceptance base is a state of *perfect information*. It is easy to verify that $\langle I, \subseteq\rangle$ is a partial order.

Base updating is a process which takes us from one information state to the next, given some new input.

Definition 4.2 (Conditional Base Updating). Consider any $w \in W$, $s \in I$, $p \in \mathcal{A}$, and $\phi, \psi \in \mathcal{L}_{\Rightarrow}$. The *Conditional Base Updating* function $\uparrow_{CBU}$: $I \longrightarrow I$ is defined by the following recursion (subscript omitted when unambiguous):

(1) $s \uparrow p = \{w \in s : w(p) = 1\}$;
(2) $s \uparrow \neg\phi = s \setminus (s \uparrow \phi)$;
(3) $s \uparrow (\phi \wedge \psi) = (s \uparrow \phi) \uparrow \psi$;
(4) $s \uparrow (\phi \Rightarrow \psi) = \{w \in s\text{: } (s \uparrow \phi) \uparrow \psi = s \uparrow \phi\}$.

Updating an information state with an atom p eliminates all possibilities from that state in which p is false. Negation is set subtraction. To update an information state with a conjunction, an agent first updates with the first conjunct, and then updates the resulting state with the second conjunct: conjunction is functional composition. So far, epistemic updating proceeds in the expected way: learning new facts about the world leads us to eliminate possibilities, winnowing away at our uncertainty.

The case for epistemic conditionals is different. The intuition behind the Ramsey Test is something like the following. I should believe an indicative If p, q just in case learning p given my present information would be enough to commit me to q. Epistemic conditionals seem to tell us more about the structure of our information about the world than they do (directly, anyway) about the world. And this is the intuition codified above in Definition 4.2. Encountering an epistemic conditional invites an agent to perform a test on her acceptance base. If the result of updating her acceptance base s with ϕ and then ψ would add no more information than merely updating s with ϕ, then updating s with the epistemic conditional $\phi \Rightarrow \psi$ returns the original acceptance base s. Otherwise, it returns the absurd acceptance base. Equivalently: if hypothetically updating s with ϕ already commits her to ψ, then her epistemic state indeed passes the test posed by the epistemic conditional $\phi \Rightarrow \psi$. Otherwise—i.e., if hypothetically updating with ϕ does *not* rule out $\neg\,\psi$—then it fails the test.

An agent does not learn anything about the world when she encounters an indicative conditional; indicatives only give her direct information about her (pre-established) acceptance base, and so about her pre-established epistemic state. An immediate objection springs to mind: surely this 'test' behavior of indicatives is counterintuitive; we clearly *do* learn something about the world when we encounter indicatives. Two replies. First, the objection gets all its force by running together the conclusive and the defeasible. Suppose that given my present information I cannot discriminate between *p, q* and *r*: I do not know which possibility is the actual world, but given my information either *p* or *q* or *r* is true (but I do not know which). Now, you are informed about the relevant facts more than I am. If you say to me If not-*p*, then r, haven't you told me something significant about the world? Should I not remove all not-*p*-and-not-*r* possibilities from my acceptance base? Not if the epistemic conditional acts as a test in the way I am suggesting. What is going on in this case, it seems to me, is that by saying If not-*p*, then r you have given me information *directly* about the structure of your information state and only *indirectly* about the world. You have given me a defeasible reason to think the world is a certain way, namely the way you think it is. But we agreed to limit our attention to the simple case where all reasons are conclusive and agents engage only in simple-minded updates; so the would-be objection just does not apply.

Second, if our attitudes toward a certain epistemic conditional differ—you are committed to it, and I am not—then there must be some plain facts in the world about which we differ, and some minimal set of such facts such that resolving our differences there would *ipso facto* resolve our difference on this conditional. So, take the case above. I am not committed to the indicative If not-*p* then *r* and you are. Once our disagreement surfaces, you can silence my protests by informing me about the plain facts which underwrite your acceptance of the conditional. And this is a perfectly normal conversational path. So, although I cannot straightaway learn the conditional, I *can* learn plain facts on which its acceptance supervenes.

Now that I have said how acceptance bases can be updated with conditional information, I have to say how it is that agents can be committed to epistemic conditionals. But first, I want to just note one property of this updating function on acceptance bases. It is an *eliminative* function in that, for any s and any ϕ whatever, $(s \uparrow \phi) \subseteq s$.[17] This, to be sure, is a *very* simplistic view of the kinematics of information. But that is precisely what we want: the easy case in which agents have no defeaters, no backtracking, and no revision.

Intuitively, an epistemic state commits you to ϕ just in case the information that ϕ carries is already contained in the information encoded in that epistemic state. One way of making this intuition precise is to think of epistemic commitment as a sort of fixed-point construction: s commits you to ϕ, ϕ is supported by your acceptance base s, just in case updating s with ϕ produces no change at all to s—just in case the information that ϕ carries is already present in s.

Definition 4.3 (Support, Entailment, Equivalence). Let s be any acceptance base, and ϕ, ψ be any formulas of $\mathcal{L}_{\Rightarrow}$. We define the Conditional Epistemic Updating support and commitment relation, $\models_{CEU}$, as follows (subscript omitted when unambiguous):

(1) *s supports ϕ*, $s \models \phi$, iff $s \uparrow \phi = s$;
(2) *ϕ commits an agent to ψ* (ϕ entails ψ), $\phi \models \psi$, iff for all s: $s \uparrow \phi \models \psi$;
(3) ϕ and ψ are *epistemically equivalent*, $\phi \cong \psi$, iff for all s: $s \uparrow \phi = s \uparrow \psi$.

An agent in s is committed to all and only those ϕ supported by s. A necessary proposition, notice, induces no change in an agent's information state—every tautology is everywhere supported. And so every agent is committed (in whatever state) to every tautology.[18]

The information expressed by epistemic conditionals is information about the structure of our acceptance bases. And this information expresses *global* modal properties of those bases. To see what this amounts to, consider the epistemic modal *might* ($\Diamond$). Intuitively, It might be that p, in its epistemic sense, is supported just in case in some possibility not yet ruled out, p. If an agent with acceptance base s has not already ruled out that ϕ, then updating with *might* ϕ should return s; otherwise—i.e., in case the agent *has* already ruled out the ϕ, then updating with *might* ϕ leads to the absurd epistemic state. The epistemic might has a natural dual: epistemic *must* ($\Box$). The sense in which the indicative expresses global modal properties of acceptance bases is reflected in the fact that we could have defined it using the unary operator for epistemic must.

To see that this is the case, we first need to introduce the material conditional ($\rightarrow$) in the usual way. Then, starting from our language with indicatives, I will show how we can introduce the unary epistemic modalities as an abbreviation. Finally, I will show that these abbreviations respect the update properties of these modals.[19]

First, then: the material conditional ($\rightarrow$) can be introduced in the standard way.

Definition 4.4. Where ϕ, ψ are any sentences of $\mathcal{L}_{\Rightarrow}$, we have the following abbreviation:

$$(\phi \rightarrow \psi) =_{df} \neg(\phi \wedge \neg\psi).$$

Similarly, we can introduce as abbreviations the epistemic *might* and its dual *must*.

Definition 4.5. Where ϕ is any formula in $\mathcal{L}_{\Rightarrow}$, and $\mathbf{1}$ is any tautology, introduce the unary modalities via the following abbreviations:

(1) $\Box\phi =_{df} \mathbf{1} \Rightarrow \phi$;
(2) $\Diamond\phi =_{df} \neg\Box\neg\phi$.

These abbreviations, moreover, preserve the update properties that we want. It *must* be that ϕ, in state s, just in case updating s with ϕ would not change s at all. Similarly, it *might* be that ϕ, in s, just in case updating s with ϕ would not leave s empty.

Proposition 4.1. *Let* ϕ *be any formula in* $\mathcal{L}_{\Rightarrow}$ *and s be any acceptance base. Then*:

(1) $s \uparrow \Box\, \phi = \{w \in s\colon s \uparrow \phi = s\}$.
(2) $s \uparrow \Diamond\, \phi = \{w \in s\colon s \uparrow \phi \neq \emptyset\}$.

Notice, also, that we could have defined the kinematic profile of $\Rightarrow$ as being a test on what must be the case, given the antecedent:

Proposition 4.2. *Let s be any acceptance base and* ϕ, ψ *be any sentences in* $\mathcal{L}_{\Rightarrow}$. *Then*:

(1) $s \uparrow (\phi \Rightarrow \psi) = \{w \in s\colon (s \uparrow \phi) \models \Box\, \psi\}$.
(2) $(\phi \Rightarrow \psi) \cong \Box(\phi \rightarrow \psi)$.

Calling indicatives "epistemic conditionals" reflects just this structure: they are global tests on what must epistemically be the case, given the antecedent.

Veltman (1985) investigates the relationship between indicatives and epistemic modals. A key advantage of his data semantics approach—which is present in my theory of conditional epistemic updating—is that this tight connection explains some otherwise puzzling semantic/epistemic data about the negations of indicatives.[20] Consider, once again, Example 1. There has been a murder in the mansion, and you are the lead investigator. Your young assistant learns that the butler has an alibi, and concludes hastily from this:

Therefore: if a member of the grounds staff did it, then it was the driver.

You rightly disagree, pointing to the fact that it might still be the gardener who is guilty.

One moral to be drawn is that in believing the negation of an indicative $\phi \Rightarrow \psi$ you are not committed to endorsing the conjunction of the antecedent with the negation of the consequent: you need not accept ϕ while $\neg\psi$. This requirement would be too strong, and would render your reply in (2) unwarranted. To deny the epistemic conditional $\phi \Rightarrow \psi$, it is sufficient that given your information state, being informed about ϕ does not *ipso facto* rule out $\neg\psi$. It *might* be that ϕ and $\neg\psi$ anyway. Looking at epistemic conditionals as tests on what epistemically must be the case, given the antecedent, explains this in a straightforward way as the next fact illustrates.

Proposition 4.3. *Let ϕ, ψ be any formulas of $\mathcal{L}_{\Rightarrow}$. Then*: $\neg(\phi \Rightarrow \psi) \cong \Diamond(\phi \wedge \neg\psi)$.

The base update function $\uparrow$, restricted to formulas which contain no occurrences of the epistemic conditional, does not differ from thinking of updating as going by way of the classical possible worlds semantics for propositional logic:

Proposition 4.4 (Veltman 1996). *Let ϕ and ψ be any formulas of $\mathcal{L}_0$, and let $[\![\phi]\!]$ abbreviate $W \uparrow \phi$. Then $[\![\cdot]\!]$ has just the properties of the classical propositional interpretation function*:

(1) $[\![p]\!] = \{w \in W : w(p) = 1\}$, *for any* $p \in \mathcal{A}$
(2) $[\![\neg\phi]\!] = W \backslash [\![\phi]\!]$.
(3) $[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$.

But in the full language for epistemic commitments $\mathcal{L}_{\Rightarrow}$, conditional epistemic updating induces a different structure on the set of information states. In particular, simple-minded updates with formulas containing open epistemic conditionals do not allow for arbitrary unions of arguments—base update does not distribute over *I*.

Definition 4.6. A function on sets f distributes over its domain X iff for every $x \in X$:

$$f(x) = \bigcup_{a \in x} f(\{a\}).$$

Proposition 4.5. *Let $[\![\cdot]\!]$ be as defined above. Then*:

(1) $[\![\cdot]\!]$ *distributes over the set I of acceptance bases.*
(2) $\uparrow$ *does not distribute over the set I of acceptance bases.*

5. Anti-Reductionism

We should take stock of our main puzzle in the epistemics of open indicatives: the Reduction Problem. The puzzle is that if rational epistemic commitment satisfies four intuitively plausible assumptions, then the epistemics of the indicative boil down to just the epistemics of the material conditional. But it seems clear that the two kinds of conditional are just different epistemic beasts. And so at least one of our four assumptions on rational epistemic commitment must be jettisoned.

But which one? Not Assumption 1, for that was just meant to restrict our attention to a simple formal language to help us attack our problem. Not Assumption 2; the whole point of epistemic conditionals is that rational epistemic commitment ties them to a sort of deduction theorem. And Assumption 4, apparently anyway, is meant to limit our attention to the "easy case" where information simply aggregates and all reasons are conclusive and non-defeasible. So Assumption 3 looks to be the culprit. Stalnaker and McGee's preferred ways out do, in fact, place the blame here. But neither of these solutions, I have argued, seems plausible. Stalnaker's pragmatic move just won't help us out of our epistemological problem, and McGee's counterexample to modus ponens for open indicatives rests on a kind of informational equivocation that is illegitimate.

All of this led us to thinking of epistemic commitment and simple-minded updates for our conditional language in a way that is sensitive to informational context and how we update that context in precise ways. The issue now is whether all of this talk of informational context is a difference that makes a difference in the Reduction Problem. It is. Epistemic updating and our dynamic notion of rational epistemic commitment, $\models_{CEU}$, both of which grew out of the intuition that information is not gratuitous, provide another way out of the Reduction Problem *qua* epistemological problem.

It is clear, by construction, that Assumption 1 is met in CEU. Our toy language is adequate for expressing epistemic commitments, and any formula of our language is a potential commitment. Assumption 2 is the requirement that rational epistemic commitment interact with indicatives to give us a sort of deduction theorem. I am committed outright to the open epistemic conditional $\phi \Rightarrow \psi$ just in case if I learn ϕ that would commit me to ψ. The next theorem establishes that our story of conditional epistemic updating meets this requirement.

Proposition 5.1 (Deduction Theorem). *Let s be any acceptance base, and ϕ, ψ be any sentences of $\mathcal{L}_{\Rightarrow}$. Then $s \uparrow \phi \models \psi$ iff $s \models \phi \Rightarrow \psi$.*

The next assumption (Assumption 3) is the requirement that the set of rational epistemic commitments of an agent at a time is closed under modus ponens, exportation, and the direct argument for the epistemic conditional.[21] Again, we can show that CEU countenances each such inference pattern as rational.

Proposition 5.2 (MP, EXP, and DA). *Let s be any acceptance base, and let ϕ, ψ, χ be any formulas in $\mathcal{L}_{\Rightarrow}$. Then*:

(1) $(s \uparrow (\phi \Rightarrow \psi)) \uparrow \phi \models \psi$ (MP).
(2) $s \uparrow ((\phi \wedge \psi) \Rightarrow \chi) \models (\phi \Rightarrow (\psi \Rightarrow \chi))$ (EXP).
(3) $s \uparrow (\alpha \vee \beta) \models \neg\alpha \Rightarrow \beta$, *for* $\alpha, \beta \in \mathcal{L}_0$ (DA).

All of this, and yet we can still resist reduction: epistemic conditionals and the material conditional are epistemologically distinct connectives. The material conditional $\phi \rightarrow \psi$ induces just the change in acceptance base s that we would expect: it is just the union of the update $s \uparrow \neg\phi$ with that of $s \uparrow \psi$. And, by Proposition 4.4, if ϕ and ψ contain no occurrences of $\Rightarrow$, this is the same as $[\![\neg\phi]\!] \cup [\![\psi]\!]$, which is the standard truth-conditional interpretation. So, we are justified in taking $\rightarrow$ to be the material conditional. Now it just remains to be seen that it induces a different update on epistemic states (via acceptance bases) than does the epistemic conditional $\Rightarrow$.

Proposition 5.3 (Distinctness). $(\phi \Rightarrow \psi) \not\cong (\phi \rightarrow \psi)$.

So the indicative is not epistemically equivalent to the material conditional. But isn't this just what Proposition 2.1 said we cannot have? Close: for full-blown reduction we also need to suppose *that epistemic commitment is classical*. CEU support (and so commitment), however, fails to have all the properties of classical entailment. Classical consequence relations, for our purposes, are assumed to satisfy three constraints: they must support all truth-functional tautologies, they must satisfy inclusion, and they must satisfy monotonicity. CEU-support certainly requires that all tautologies are everywhere supported. However, the other two properties of classical consequence relations fail to hold for $\models_{CEU}$. It is interesting that monotonicity and inclusion fail for much the same reason—namely, that $\models_{CEU}$ is a dynamic notion of epistemic commitment.

First, consider monotonicity. Classical consequence is monotonic: if $\Gamma \models_{CL} \phi$ then $\Delta \models_{CL} \phi$ for each $\Delta \supseteq \Gamma$. Monotonicity can be thought of as a special case of the more general property of *persistence* (Veltman, 1985, 1996). Intuitively, a consequence (support) relation is persistent if adding more information never leads us to "taking back" what has already been concluded. In the classical case, "more information" is represented as a bigger premise set on the left of the consequence relation (and so monotonicity coincides with persistence in the classical case).

In CEU, "more information" can naturally be thought of as a shrinking set of possibilities to the left of $\models_{CEU}$: more information is a more complete picture of the world. CEU-support is not persistent. More complete information can take us from a state in which a sentence is supported to one in which the same sentence is not supported. In a more formal dress:

Definition 5.1 (Persistence). A support relation $\models \subseteq I \times \mathcal{L}_{\Rightarrow}$ is *persistent* iff for any $s \in I$ and $\phi \in \mathcal{L}_{\Rightarrow}$: if $s \models \phi$ then $s' \models \phi$, for every $s' \subseteq s$.

Proposition 5.4 (Non-persistence). *CEU-support*, $\models_{CEU}$, *is not persistent*.

Here is a simple counter-example to persistence involving epistemic conditionals. Suppose you have just two possibilities in your acceptance base—the first a world in which it is both sunny and warm, and the second a world in which it is not sunny but it is warm. If your acceptances are arranged in this way, then you do not accept the conditional If it is warm, then it is sunny since learning that it is warm would not rule out the not-sunny world. In fact, your epistemic state supports the negated conditional: It is not the case that if it is warm then it is sunny. But now, let us say that you learn, in a simple-minded update sense of 'learn', that it is sunny. In the resulting acceptance base—which is a proper subset of your earlier base—the negated conditional is no longer supported: you no longer accept that It is not the case that if it is warm then it is sunny since every warm world, according to the facts you have, is in fact a sunny one.

A suitable translation of inclusion also fails for CEU-commitment. What is telling is that it *ought* to fail in the context of open epistemic conditionals. Inclusion is the requirement that holding the commitments Γ commits an agent to each member γ in Γ. And this seems reasonable enough, indeed obvious enough, if commitments are restricted to plain facts expressible in classical propositional logic. But once we take seriously the idea that epistemic commitment might have dynamic properties, then this is much less clear. A natural translation of the Inclusion property is to say that a commitment relation is *reflexive*—indeed, Inclusion is a generalization of Reflexivity:

Definition 5.2 (Reflexivity). A support relation $\models$ is *reflexive* iff for any $\phi \in \mathcal{L}_{\Rightarrow}, \phi \models \phi$.

Proposition 5.5. *CEU-support, $\models_{CEU}$, is not reflexive.*

Reflexivity fails for reasons of non-persistence. Consider the formula $(\Diamond p \wedge \neg p)$. Updating a state s with this will leave that state either empty or with only not-p possibilities. In the latter case, the resulting state will not support the conjunction $(\Diamond p \wedge \neg p)$ since it will not support that it might be that p. The epistemic modal $\Diamond p$ is not persistent. Of course, there are examples which illustrate this with respect to open indicatives as well: consider the conjunction $\neg(p \Rightarrow q) \wedge q$ in a state like the one described above (where p is It is warm and q is It is sunny).

And, of course, non-persistence has another consequence: conjunction is not (in general) commutative by the lights of CEU—that is, it is not the case that $\phi \wedge \psi$ is always epistemically equivalent to $\psi \wedge \phi$. For consider the case sketched above: $(\Diamond p \wedge \neg p)$ induces quite a different update profile than does $(\neg p \wedge \Diamond p)$. Updating any state with a conjunction is just functional composition. So, updating a state which contains just a p-world and a $\neg p$-world with $\Diamond p$ followed by $\neg p$ leaves such a state containing just the $\neg p$-world. However, reversing the order of the updates makes a difference: updating a state with $\neg p$ eliminates the p-world. If we then update *this* state by $\Diamond p$, the result is the absurd state—quite a different result.[22] Now, if a conjunction involves no modalities at all—that is, if it is a conjunction expressible in the classical fragment alone—then Proposition 4.4 guarantees that such a conjunction will commute. But no such property should be expected in the full language. The epistemic might is meant to act as a check on what has not yet been ruled out. As what has been ruled out changes (even in a straightforward, monotonic way) we should expect this to be reflected in what modals are supported. Given the interdefinability of might with the epistemic conditional plus negation, similar remarks apply to commitments like It is not so that if..., then....[23]

Some potential commitments inevitably "crash" in that the information they carry can never be successfully incorporated to any acceptance base—update always reduces to absurdity. Let us call formulas expressing such commitments *inconsistent*. Conjunctions like $(\neg p \wedge \Diamond p)$ are inconsistent in this sense. We saw above that its commutation does not have this property—there are acceptance bases which do not go empty upon an update with $(\Diamond p \wedge \neg p)$. That being said, such a conjunction does seem intuitively rather different from more mundane ones like It is raining and it is cold. Intuitively speaking, to successfully update with a conjunction like $(\Diamond p \wedge \neg p)$, i.e. to update a non-empty acceptance base with it without reducing that base to $\emptyset$, one has to be in a state in which it might be that p (whence the first conjunct poses a test we pass) and then one has to learn that in fact it is not so that p. Other conjunctions, unadorned with any modalities, need not *require* such a mid-conjunction addition of information. What this means is that even though there are acceptance bases which can be successfully updated with conjunctions like $(\Diamond p \wedge \neg p)$ without going empty (and hence such conjunctions are consistent), there is no *single* non-empty base which *supports* such a conjunction—updating with such a conjunction requires a change in view, albeit a monotonic shrinking of one's acceptance base.[24] And so there is no non-empty state which can be the fixed-point of the update with $(\Diamond p \wedge \neg p)$. It is, in the end, this fact which sheds light on why reflexivity fails.

All of this, keep in mind, even though we are constraining our attention to the easy case in which all information simply aggregates—all information, we have been assuming, is non-defeasible in the sense that agents never have default rules and their associated expectations. And my picture of epistemic updating and rational epistemic commitment in CEU has been faithful to that. What Proposition 5.4 illustrates is that defeasibility (non-monotonicity) is not the only source of non-persistence in the set of our rational epistemic commitments. Epistemic modalities, like the ones involved in epistemic conditionals, are another source. So while it makes perfect sense to limit our attention to the "easy case" of information kinematics when looking at open indicatives, it is not reasonable on that basis to insist that rational epistemic commitment be a classical consequence relation. It must be more dynamic than that. Assumption 4, I contend, is the culprit in the Reduction Problem.

6. Conclusion: Just how strict are these conditionals?

I want to conclude the story I have been telling about open epistemic conditionals by looking a bit at just how strict these conditionals are required to be. It is clear, I hope, that my indicatives are *variably strict* conditionals in the sense that, depending on the space of possibilities left uneliminated by an agent in a certain information state, the epistemic kinematics of these conditionals tests for a strict connection between the antecedent and its consequent. In fact, the epistemics of conditionals is variable twice over: they are *variably* variably strict conditionals.

In one sense open indicatives in CEU are variably strict because the acceptance bases shrink as agents monotonically gain information. And since the conditionals' kinematics prescribes that they act as tests upon those bases, how strict the conditional is depends on the size of the base: the fewer worlds in s—the stronger the agent's present epistemic commitments are with respect to plain facts—then *ceteris paribus* the weaker the connection between antecedent and consequent needed for the state to pass the test posed by the indicative.

The other sense in which indicatives are variably strict has to do with the *growth* of acceptance bases. Intuitively, an acceptance base s' which properly extends s contains *less* information than does s. There are two ways that these bases might grow. The first way is connected to belief-contravening indicatives, and the Problem of Epistemic Change. If an agent accepts that p, then all of the possibilities in her acceptance base are p-worlds. But she could learn that she was mistaken; in fact not-p. The contours of rational belief change in such cases demand that her acceptance base in her revised state not be included (properly or otherwise) in her old acceptance base. But this way that bases can grow is orthogonal to our main focus here on open indicatives.

The other mechanism by which acceptance bases can grow is relevant to the case of open epistemic conditionals. All along we have tacitly been thinking of the space of worlds W as a finite space of worlds, containing just the worlds which are relevant in covering the space of possibilities with respect to our contextually fixed set $\mathcal{A}$ of atoms which are relevant. Of course, there are more possible worlds than this. But not all of them are always relevant. To use a phrase introduced by David Lewis (1979, 1996), we *properly ignore* most of the possibilities most of the time. But sometimes a world which was properly ignored becomes salient, and so the space of possibilities shifts. Such shifts make our acceptance bases bigger (since such bases are measures of uncertainty *with respect to* the space of possibilities), and so the conditional connection required between an antecedent and a consequent of an indicative in order for an acceptance base to pass the test it poses is *ceteris paribus* stricter than before the shift.

Lewis (1996) gives a series of rules for the proper ignoring of worlds. Since his aim is an analysis of of the locution "S knows that p" and my aim is rather the epistemics of open indicatives, not all of his rules are important for our present aims. In fact, it is one rule in particular that is relevant here: the Rule of Attention—a possibility which is not ignored is not *properly* ignored. This rule, notice, is an indicative rule. Bringing up heretofore ignored possibilities makes them relevant, and so shifts the space of W. Thus shifted, indicatives which were once supported may become unsupported.[25]

Examples which seem to suggest that indicatives do not support "antecedent strengthening", it seems to me, can be thought of rather as showing how previously ignored possibilities can be introduced into the epistemic context. Suppose I am throwing a party. No official invitations, just word of mouth. W, let us suppose, just covers the atoms B, C, D and F for Billy's coming, Chris's coming, Danny's coming, and it being a fun party. I believe that

(18) If Billy comes to the party, it will be loads of fun.

So in my acceptance base all the Billy-worlds are also loads-of-fun-worlds. Then you remind me:

(19) But if Billy and Alex come, it will be dull. They only talk with each other.

You are quite right! I was *presupposing* that Alex would not come! But now that you mention the possibility, I cannot go on presuming it will not happen unless I have some special reason for *ruling out* that Alex will come. If I do not have such a reason, then I must accommodate this possibility. The space of possibilities W, now, has to have grown, and with it my acceptance base. It is with respect to this larger space of possibilities that the conditional in (19) is supported, while (18) no longer is. We can continue the story: a bit more thought on my part and I realize that Eric might also come. And Eric, it is common knowledge, cancels out Alex's coming so that the three of them all make it a great party. And so on. Such contextual shifts in the space W of possibilities then varies the *relative size* of my acceptance base to the space of possibilities. And it is the relative size of an acceptance base which governs how tight a connection must exist between the antecedent and the consequent for a base to pass the test posed by an indicative.[26]

This way of looking at the failure of antecedent strengthening on indicatives explains why order matters when considering juxtaposing pairs of indicative conditionals. If our discourse unfolds from (18) to (19) then everything, both conversationally and epistemically speaking, is unproblematic. But, if the order is reversed, then we get the feeling that something is amiss. If it is the case that if Billy and Alex come, it will be dull—and you inform me of this fact—then I cannot go on believing that if Billy comes it will be loads of fun unless and until I can eliminate the Alex-comes possibilities. Once the space of possibilities is shifted outward to include Billy-and-Alex worlds, and if it is the case that all such possibilities are dull-party possibilities, then my information just does not support the simple conditional claim.

Appendix A. Proofs of main propositions

PROPOSITION 4.1 Let ϕ be any formula in $\mathcal{L}_{\Rightarrow}$ and s be any acceptance base. Then:

(1) $s \uparrow \Box\phi = \{w \in s : s \uparrow \phi = s\}$.
(2) $s \uparrow \Diamond\phi = \{w \in s : s \uparrow \phi \neq \emptyset\}$.

Proof. (1) Note that, if **1** is any truth-functional tautology then for any s whatever, $s \uparrow \mathbf{1} = s$. Then, applying Definitions 4.5 and 4.2, we have that

$$\begin{aligned} s \uparrow \Box\phi &= s \uparrow (\mathbf{1} \Rightarrow \phi) \\ &= \{w \in s : (s \uparrow \mathbf{1}) \uparrow \phi = s\} \\ &= \{w \in s : s \uparrow \phi = s\} \end{aligned}$$

which proves the theorem.

(2) By Definitions 4.5 and 4.2, we have that

$$\begin{aligned} s \uparrow \Diamond\phi &= s \uparrow \neg\Box\neg\phi \\ &= s \backslash s \uparrow \Box\neg\phi \end{aligned}$$

Now, by clause (1) of the proposition we know that

$s \uparrow \Box\neg\phi = \{w \in s : s \uparrow \neg\phi = s\} = \{w \in s : s \setminus s \uparrow \phi = s\}$. So, $s \uparrow \Box\neg\phi = s$ if $s \setminus (s \uparrow \phi) = s$, i.e. if $s \uparrow \phi = \emptyset$. And $s \uparrow \Box\neg\phi = \emptyset$ if $s \setminus (s \uparrow \phi) = \emptyset$, i.e., if $s \uparrow \phi \neq \emptyset$. Thus we have that

$$s \uparrow \Diamond\phi = s \backslash s = \emptyset \text{ if } s \uparrow \phi = \emptyset$$

and

$$s \uparrow \Diamond\phi = s \backslash \emptyset = s \text{ if } s \uparrow \phi \neq \emptyset.$$

That is, $s \uparrow \Diamond\phi = \{w \in s : s \uparrow \phi \neq \emptyset\}$, as required. □

PROPOSITION 4.2 Let s be any acceptance base and ϕ, ψ be any sentences in $\mathcal{L}_{\Rightarrow}$. Then:

(1) $s \uparrow (\phi \Rightarrow \psi) = \{w \in s : (s \uparrow \phi) \models \Box\psi\}$.
(2) $(\phi \Rightarrow \psi) \cong \Box(\phi \rightarrow \psi)$.

Proof. (1) This is immediate from Proposition 4.1 and Definition 4.3.

(2) We need to see that, for any s, $s \uparrow (\phi \Rightarrow \psi) = s \uparrow \Box(\phi \rightarrow \psi)$. Consider any s; $s \uparrow \Box(\phi \rightarrow \psi) = s$ if $s \uparrow (\phi \rightarrow \psi) = s$ and $s \uparrow \Box(\phi \rightarrow \psi) = \emptyset$ otherwise. So suppose that $s \uparrow (\phi \rightarrow \psi) = s$. (We show that, in this case, $s \uparrow (\phi \Rightarrow \psi) = s$.) Now, applying our

definitions for defined connectives, $s \uparrow (\phi \rightarrow \psi) = s$ iff $s \uparrow \neg(\phi \wedge \neg\psi) = s$. And, by several applications of Defintion 4.2,

$$\begin{aligned} s \uparrow \neg(\phi \wedge \neg\psi) = s \quad & \text{iff} \quad s \backslash s \uparrow (\phi \wedge \neg\psi) = s \\ & \text{iff} \quad s \uparrow (\phi \wedge \neg\psi) = \emptyset \\ & \text{iff} \quad (s \uparrow \phi) \uparrow \neg\psi = \emptyset \\ & \text{iff} \quad (s \uparrow \phi) \backslash (s \uparrow \phi) \uparrow \psi = \emptyset \\ & \text{iff} \quad (s \uparrow \phi) \uparrow \psi = s \uparrow \phi \end{aligned}$$

But then it is clear that $s \uparrow (\phi \Rightarrow \psi) = s$, as required. If, on the other hand, $s \uparrow (\phi \rightarrow \psi) \neq s$, then similar reasoning establishes that $s \uparrow (\phi \Rightarrow \psi) = \emptyset$, and this is precisely the same as $s \uparrow \Box(\phi \rightarrow \psi)$, which completes the proof. □

PROPOSITION 4.3 Let ϕ, ψ be any formulas of $\mathcal{L}_{\Rightarrow}$. Then:

$$\neg(\phi \Rightarrow \psi) \cong \Diamond(\phi \wedge \neg\psi).$$

Proof. Immediate from Proposition 4.2. □

PROPOSITION 4.5 Let $[\![\cdot]\!]$ be as defined above. Then:

(1) $[\![\cdot]\!]$ distributes over the set I of acceptance bases.
(2) $\uparrow$ does not distribute over the set I of acceptance bases.

Proof. (1) Note that $[\![\cdot]\!]$ distributes over I iff it distributes over the limiting case of an acceptance base W. So we need to see that $[\![\cdot]\!]$ distributes over W. The following lemma is established by a routine induction on $\phi \in \mathcal{L}_0$:

LEMMA: For any $w \in W$: $\{w\} \uparrow \phi = \{w\}$ iff $w \in [\![\phi]\!]$ (else $\{w\} \uparrow \phi = \emptyset$).

The proof of the theorem then proceeds by an induction on $\phi \in \mathcal{L}_0$. For the atomic case, let p be any atom. From Proposition 4.4 we have that $[\![p]\!] = W \uparrow p = \{w \in W : w(p) = 1\}$. Clearly, this is the same as $\bigcup_{w \in W} (\{w\} \uparrow p)$. So assume, for the inductive hypothesis, that

$$[\![\phi]\!] = \bigcup_{w \in W} (\{w\} \uparrow \phi).$$

Consider $[\![\neg\phi]\!]$:

$$\begin{aligned} [\![\neg\phi]\!] &= W \backslash [\![\phi]\!] \text{ (Prop. 4.4)} \\ &= W \backslash \bigcup_{w \in W} (\{w\} \uparrow \phi) \text{ (\textit{Inductive Hyp.})} \\ &= \bigcap_{w \in W} (W \backslash \{w\} \uparrow \phi) \text{ (set theory).} \end{aligned}$$

It follows from the lemma that $\bigcap_{w\in W}(W\backslash\{w\}\uparrow\phi)$ will be just the set of w's in W such that $\{w\}\uparrow\phi=\emptyset$. That is,

$$\bigcap_{w\in W}(W\backslash\{w\}\uparrow\phi)=\{w\in W:\{w\}\uparrow\phi=\emptyset\}.$$

But

$$\{w\in W:\{w\}\uparrow\phi=\emptyset\}=\bigcup_{w\in W}(\{w\}\backslash\{w\}\uparrow\phi).$$

And, by the definition of $\uparrow$, this latter is just $\bigcup_{w\in W}(\{w\}\uparrow\neg\phi)$, as required.
Consider $[\![\phi\wedge\psi]\!]$:

$$\begin{aligned}[\![\phi\wedge\psi]\!]&=[\![\phi]\!]\cap[\![\psi]\!]\ \text{(Prop. 4.4)}\\&=(\bigcup_{w\in W}\{w\}\uparrow\phi)\cap(\bigcup_{w\in W}\{w\}\uparrow\psi)\ \text{(Inductive Hyp.)}\\&=\bigcup_{w\in W}((\{w\}\uparrow\phi)\cap(\{w\}\uparrow\psi))\ \text{(set theory).}\end{aligned}$$

It follows from the lemma that $\bigcup_{w\in W}((\{w\}\uparrow\phi)\cap(\{w\}\uparrow\psi))$ will be just the w's in W such that $\{w\}\uparrow\phi=\{w\}$ and $\{w\}\uparrow\psi=\{w\}$. But notice that if w is such that $\{w\}\uparrow\phi=\{w\}\uparrow\psi=\{w\}$, then it must be that $(\{w\}\uparrow\phi)\uparrow\psi=\{w\}$. Thus it follows that

$$\bigcup_{w\in W}((\{w\}\uparrow\phi)\cap(\{w\}\uparrow\psi))=\bigcup_{w\in W}((\{w\}\uparrow\phi)\uparrow\psi).$$

And, by the definition of $\uparrow$, this latter is just $\bigcup_{w\in W}(\{w\}\uparrow(\phi\wedge\psi))$, as required.

(2) Let $s=\{w_1, w_2\}$ where $w_1(p)=1$, $w_1(q)=0$, and $w_2(p)=w_2(q)=1$. Now, consider the update of s with the formula $\neg(p\Rightarrow q)$:

$$\begin{aligned}s\uparrow\neg(p\Rightarrow q)&=s\backslash s\uparrow(p\Rightarrow q)\ \text{(Def. 4.2)}\\&=s\backslash\emptyset\\&=s\end{aligned}$$

But:

$$\begin{aligned}\{w_1\}\uparrow(\neg(p\Rightarrow q))&=\{w_1\}\backslash(\{w_1\}\uparrow(p\Rightarrow q))\ \text{(Def. 4.2)}\\&=\{w_1\}\backslash\emptyset\ \text{(Def. 4.2)}\\&=\{w_1\}\\\{w_2\}\uparrow(\neg(p\Rightarrow q))&=\{w_2\}\backslash(\{w_2\}\uparrow(p\Rightarrow q))\ \text{(Def. 4.2)}\\&=\{w_2\}\backslash\{w_2\}\\&=\emptyset\end{aligned}$$

And so $\bigcup_{w\in s}\{w\}\uparrow(\neg(p\Rightarrow q))=\{w_1\}\neq s=s\uparrow(\neg(p\Rightarrow q))$. □

PROPOSITION 5.1 Let s be any acceptance base, and ϕ, ψ be any sentences of $\mathcal{L}_{\Rightarrow}$. Then $s \uparrow \phi \models \psi$ iff $s \models \phi \Rightarrow \psi$.

Proof.

$$\begin{aligned} s \uparrow \phi \models \psi \text{ iff } \quad & (s \uparrow \phi) \uparrow \psi = s \uparrow \phi \text{ (Def. 4.3)} \\ \text{iff } \quad & s \uparrow (\phi \Rightarrow \psi) = s \text{ (Def. 4.2)} \\ \text{iff } \quad & s \models \phi \Rightarrow \psi \text{ (Def. 4.3).} \end{aligned}$$

PROPOSITION 5.2 Let s be any acceptance base, and let ϕ, ψ, χ be any formulas in $\mathcal{L}_{\Rightarrow}$. Then:

(1) $(s \uparrow (\phi \Rightarrow \psi)) \uparrow \phi \models \psi$ (MP).
(2) $s \uparrow ((\phi \wedge \psi) \Rightarrow \chi) \models (\phi \Rightarrow (\psi \Rightarrow \chi))$ (EXP).
(3) $s \uparrow (\alpha \vee \beta) \models \neg\alpha \Rightarrow \beta$, for $\alpha, \beta \in \mathcal{L}_0$ (DA).

Proof. (1) In order to prove this clause of the theorem, note that it is sufficient to prove the following

LEMMA: $((s \uparrow (\phi \Rightarrow \psi)) \uparrow \phi) \uparrow \psi = (s \uparrow (\phi \Rightarrow \psi)) \uparrow \phi$

Proof of Lemma. By Definition 4.2, either $s \uparrow ((\phi \Rightarrow \psi) = \emptyset$ or $s \uparrow ((\phi \Rightarrow \psi) = s$. If the former, then $\emptyset \uparrow \phi = (\emptyset \uparrow \phi) \uparrow \psi = \emptyset$. Hence, trivially, in this case the claim holds. So suppose the latter—i.e., that $s \uparrow (\phi \Rightarrow \psi) = s$. Since $s \uparrow (\phi \Rightarrow \psi) = s$, it is enough to show that $(s \uparrow \phi) \uparrow \psi = s \uparrow \phi$. By Definition 4.2, $s \uparrow (\phi \Rightarrow \psi) = s$ iff $(s \uparrow \phi) \uparrow \psi = s \uparrow \phi$, as required to prove the lemma.

(2) Let $s' = s \uparrow ((\phi \wedge \psi) \Rightarrow \chi)$. We must show that $s' \models \phi \Rightarrow (\psi \Rightarrow \chi)$. By Definition 4.2, either $s' = \emptyset$ or $s' = s$. If $s' = \emptyset$, then trivially $s' \models \phi \Rightarrow (\psi \Rightarrow \chi)$. So suppose that $s' = s$. That is, suppose that $s \uparrow ((\phi \wedge \psi) \Rightarrow \chi) = s$. Then we have the following:

$$\begin{aligned} s \uparrow ((\phi \wedge \psi) \Rightarrow \chi) = s \quad \text{iff} \quad & (s \uparrow (\phi \wedge \psi)) \uparrow \chi = s \uparrow (\phi \wedge \psi) \text{ (Def.4.2)} \\ \text{iff} \quad & s \uparrow (\phi \wedge \psi) \models \chi \text{ (Def.4.3)} \\ \text{iff} \quad & (s \uparrow \phi) \uparrow \psi \models \chi \text{ (Def.4.2)} \\ \text{iff} \quad & s \models \phi \Rightarrow (\psi \Rightarrow \chi) \text{ (Prop.5.1, twice)} \\ \text{iff} \quad & s' \models \phi \Rightarrow (\psi \Rightarrow \chi) \text{ (since } s' = s). \end{aligned}$$

(3) Let $s' = s \uparrow (\alpha \vee \beta)$. Hence, $s' = (s \uparrow \alpha) \cup (s \uparrow \beta)$. Since $\alpha_1 \beta \in \mathcal{L}_0$, by Proposition 4.4, $s' = (s \cap [\![\alpha]\!]) \cup (s \cap [\![\beta]\!])$. Now, $s' \models \neg\alpha \Rightarrow \beta$ iff $s' \uparrow \neg\alpha \models \beta$, by Proposition 5.1. And so iff $(s' \uparrow \neg\alpha) \uparrow \beta = s' \uparrow \neg\alpha$. But $s' \uparrow \neg\alpha \subseteq s \cap [\![\beta]\!]$ and (since $\uparrow$ is eliminative) we then have that $(s' \uparrow \neg\alpha) \uparrow \beta = s' \uparrow \neg\alpha$, as required. □

PROPOSITION 5.3 $(\phi \Rightarrow \psi) \not\cong (\phi \rightarrow \psi)$.

Proof. It is enough to show that there is an acceptance base $s \in I$, and formulas $\phi, \psi \in \mathcal{L}_{\Rightarrow}$ such that $s \uparrow (\phi \rightarrow \psi) \neq s \uparrow (\phi \Rightarrow \psi)$. Let $s = \{w_1, w_2, w_3, w_4\}$ where:

$w_1(p) = w_1(q) = 1$; $w_2(p) = 1$, $w_2(q) = 0$; $w_3(p) = 0$, $w_3(q) = 1$; and $w_4(p) = w_4(q) = 0$. Then we have

$$\begin{aligned} s \uparrow (p \rightarrow q) &= (s \uparrow \neg p) \cup (s \uparrow q) \\ &= \{w_3, w_4\} \cup \{w_1, w_3\} \text{ (Def.4.2)} \\ &= \{w_1, w_3, w_4\} \end{aligned}$$

But, $s \uparrow (p \Rightarrow q) = \{w \in s : (s \uparrow p) \uparrow q = s \uparrow p\}$ (Def. 4.2), and

$$\begin{aligned} s \uparrow p &= \{w_1, w_2\} \quad \text{while} \\ (s \uparrow p) \uparrow q &= \{w_1, w_2\} \uparrow q \text{ (Def.4.2)} \\ &= \{w_1\} \end{aligned}$$

So, $s \uparrow (p \Rightarrow q) = \emptyset \neq \{w_1, w_3, w_4\} = s \uparrow (p \rightarrow q)$, as required. □

Proposition 5.4 CEU-support, $\models_{CEU}$, is not persistent.

Proof. Let $s = \{w_1, w_2\}$ where: $w_1(p) = 1$, $w_1(q) = 0$; and $w_2(p) = w_2(q) = 1$. Consider the simple-minded update of s with $\neg(p \Rightarrow q)$:

$$\begin{aligned} s \uparrow (\neg(p \Rightarrow q)) &= s \backslash (s \uparrow (p \Rightarrow q)) \\ &= s \backslash \emptyset = s \text{ (Def 4.2 twice).} \end{aligned}$$

And so we have that $s \models \neg(p \Rightarrow q)$. However, $s' = \{w_2\} \subseteq s$ and

$$\begin{aligned} \{w_2\} \uparrow (\neg(p \Rightarrow q)) &= \{w_2\} \backslash \{w_2\} \uparrow (p \Rightarrow q) \\ &= \{w_2\} \backslash \{w_2\} = \emptyset \end{aligned}$$

Since $\{w_2\} \uparrow (\neg(p \Rightarrow q)) \neq \{w_2\}$, we have that $s' \not\models \neg(p \Rightarrow q)$ for some $s' \subseteq s$, as required.

Proposition 5.5 CEU-support, $\models_{CEU}$, is not reflexive.

Proof. Let $s = \{w_1, w_2\}$ where: $w_1(p) = 1$, $w_1(q) = 0$; and $w_2(p) = w_2(q) = 1$. Consider the update of s with $\neg(p \Rightarrow q) \wedge q$:

$$\begin{aligned} s \uparrow (\neg(p \Rightarrow q) \wedge q) &= (s \uparrow \neg(p \Rightarrow q)) \uparrow q \\ &= (s \backslash (s \uparrow p \Rightarrow q)) \uparrow q \\ &= s \uparrow q \\ &= \{w_2\} \end{aligned}$$

Let $s' = \{w_2\}$. Now, $\models_{CEU}$ is reflexive only if $s' \models \neg(p \Rightarrow q) \wedge q$. But, by Definition 4.3, $s' \models \neg(p \Rightarrow q) \wedge q$ iff $s' \uparrow (\neg(p \Rightarrow q) \wedge q) = s'$. And this only if $s' \uparrow \neg(p \Rightarrow q) = s'$. But $s' \uparrow \neg(p \Rightarrow q) = \emptyset$. Whence it follows that $s' \not\models \neg(p \Rightarrow q) \wedge q$, as desired. □

Notes

* Thanks are due to Nicholas Asher, David Chalmers, Kai von Fintel, Shaughan Lavine, Dick Oehrle, John Pollock, Mary Rigdon, David Sosa, Jason Stanley, Frank Veltman, and three anonymous referees for their careful, persistent and helpful comments. Thanks also to audiences at The University of Arizona, the Department of Logic and Philosophy of Science at the University of California at Irvine, The University of Michigan, The University of Texas at Austin, and at Harvard University.

[1] There is a worry—a serious worry—over exactly how to specify the class of "indicative" or "epistemic" conditionals. If we are using surface grammar as our guide, then we may be painting ourselves into a bit of a corner since the indicative/subjunctive distinction among conditionals is not entirely a happy one: the subjunctive mood in English seems to be rather impoverished—so much so that there is controversy about whether there is *any* subjunctive mood in English at all. Matters only get more complicated when we turn to trying to mark the difference between indicatives and subjunctives (between epistemics and counterfactuals). However, I hope that what I have to say here is independent of these worries. I have in mind a particular subset of the class of non-counterfactual conditonals which typically get expressed in the indicative mood—the same class of conditionals others before me (Gibbard, Edgington, Jackson, Lewis, Stalnaker, Veltman and countless others) have been more or less content to label "indicative" and "epistemic". This class of conditionals has some other serious puzzles associated with it, and I hope to tackle some of those.

[2] That this conditional is an epistemic one, and not a "hidden" counterfactual, can be seen by contrasting it with If Oswald hadn't killed Kennedy, then someone else would have.

[3] See, for example, the motivation, theories, and references in Gärdenfors (1988), Hansson (1999), and Gillies (2004).

[4] By calling this the Reduction *Problem* I am anticipating just a bit an argument in Section 2 that the material conditional analysis of the ordinary "if-then" is, at least as far as rational belief is concerned, inadequate: the lights of that analysis leaves us with irrational epistemic commitments.

[5] $\mathcal{L}_0$ is just the language of classical propositional logic.

[6] The clearest intuitions about DA are cases in which from $\phi \vee \psi$ we conclude $\neg\phi \Rightarrow \psi$ *when* ϕ, ψ *contain no occurrences of* '$\Rightarrow$'. In order to make our puzzle as strong as possibile, this will always be the version of DA up for consideration in this paper.

[7] The direct argument plus our assumptions about rational epistemic commitment poses the same problem so long as we have MP. Similarly, Hanson (1991) argues that indicatives must be truth-functional if they satisfy what I am calling here Assumption 2. But his arguments only work in the context of modus ponens and a classical entailment relation.

[8] As far as I can tell, each proves the result independently of the others.

[9] I don't pretend that this constitutes an exhaustive survey of the ink spilled on indicative conditionals. But what ink has been spilled can for our present purposes be lumped and grouped in the way I describe.

[10] To be clear, I am not suggesting that Jackson or Lewis or Stalnaker claims to be solving an epistemological problem with a pragmatic move. The epistemological problem is mine; I am suggesting that their preferred ways out of the Reduction Problem *qua* semantics problem are not open to those us interested in solving the Reduction Problem *qua* epistemics problem.

[11] For example, suppose it is the case that we think only someone with both easy enough means and strong enough motive did it. The gardener ranks relatively high along both dimensions, the butler quite high on motive but not means, and the driver moderately high on means but not motive. Now suppose we have evidence which indicates that the means actually used moderately indicts some member or other of the grounds staff. Then we might be almost certain the gardener did it, still think it seriously possible that the butler did it (our evidence which

seems to point in the direction of the grounds staff isn't that strong and might be misleading), but think it pretty remote that the driver did it.

[12] In informal surveys, philosophers seem to be split more or less evenly on their first intuitive reaction to McGee-like counterexamples: roughly half admit feeling like something has been pulled over on them, and roughly half are quick to say that (finally!) someone has a clear example that MP is not a rational inference rule.

[13] For more on this sort of phenomenon, and a semantics for (first-order extensional) predicate logic that accounts for it, see Groenendijk and Stokhof (1991).

[14] Resembles, but does not mirror. The core intuition behind the account of indicatives in "data semantics" (Veltman, 1985) and the account I offer here is, I think, the same. But the resulting logic there is quite different from that here (different argument forms are validated). I have adopted and adapted what I take the core intuition to be at play in "update semantics" (Veltman, 1996) about the various non-modal formulas and unary epistemic modals, and tried to make that fit with something like the Ramsey Test for epistemic conditionals. But, again, the details differ.

[15] This assumption is justified by our agreement to limit ourselves to the easy case of information updating where all information is non-defeasible. There is no in-principle obstacle to relaxing the assumption—in Gillies (2004), e.g., epistemic states are taken to be pairs: a set of possible worlds plus some ordering information over the space of possibilities. The term "acceptance bases" is meant to remind us that the possibilities in a base are possibilities not ruled out by what we take ourselves to know, irrespective of defeasible rules.

[16] I assume that the functions are non-empty, of course.

[17] Whence it follows that the absurd state is a fixed-point of every update: $\emptyset \uparrow \phi = \emptyset$.

[18] Our concept of commitment can be generalized in the expected way: $\phi_1, \ldots, \phi_n \models \psi$ iff for any s: $(s \uparrow \phi_1) \uparrow \ldots \uparrow \phi_n \models \psi$. This consequence relation is an "Update-to-Test" consequence relation (van Benthem, 1996). There are, of course, other possibilities for defining a consequence relation for our language. For instance, one might define $\Vdash$ as follows: $\phi_1, \ldots, \phi_n \Vdash \psi$ iff for any s: if $s \models \phi_1$ and $\ldots s \models \phi_n$ then $s \Vdash \psi$. Such a consequence relation, a "Test-to-Test" consequence relation, is ill-suited for our present purposes. What went wrong in the McGee-like cases, as we saw, was that information was not allowed to pass from premises to conclusion. The intuition we have been pursuing has it that ϕ commits an agent to ψ just when upon learning that ϕ one is thereby in a state in which ψ is supported.

[19] Proofs of the theorems have been relegated to the appendix.

[20] Data semantics (DS) does have some difficulty with the interaction between the indicative and *might*. For example, $\phi \Rightarrow \Diamond\psi \models_{DS} \phi \Rightarrow \psi$, which seems counterintuitive. Not so in CEU: $\phi \Rightarrow \diamond\psi \not\models_{CEU} \phi \Rightarrow \psi$— If p, then might q does not commit an agent to If p, then q.

[21] The direct argument, as before, is limited to formulas in $\mathcal{L}_0$.

[22] The failure of commutivity in the presence of might is not new here; see, for example, Veltman (1996) and van Benthem (1996). The differential treatment of conjunctions involving epistemic modals is put to work in Gillies (2001) to solve Moore's paradox. And, it should be noted, that although not in general commutative, conjunction in CEU is associative.

[23] It is worth pointing out that although conjunction is not (in general) commutative by the lights of CEU, *disjunction* is (or would be, if we introduced it in the usual way). Question: How can this be, given the definability of disjunction in terms of conjunction and negation? Answer: Negation in dynamic semantics acts as a barrier to dynamic effects.

[24] Groenendijk, Stokhof, and Veltman (1996) call this property "coherence", but I have (elsewhere) called it "cohesiveness" since it requires that a piece of information hang together in a single state.

[25] Of course the shifting of the space of not properly ignored possibilities has similar effects in the classical fragment as well.

[26] A similar accommodation-based story can be, and is, told about the relationship between 'would'-counterfactuals and their 'might'-counterfactual cousins (Gillies, 2003). The difference between that treatment and the one here is a matter of complexity of the semantics:

there I roll accommodation into the semantics proper, whereas here I am treating it as post-semantic. This simplification is to keep what I take the main story surrounding the Reduction Problem clearly in view.

References

Adams, E. W. (1975). *The Logic of Conditionals. Dordrecht: Reidel.*

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

Edgington, D. (1995). "Conditionals," *Mind* **104**: 235–329.

Gärdenfors, P. (1984). "The dynamics of belief as a basis for logic," *British Journal for the Philosophy of Science* **35**: 1–10.

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

Gibbard, A. (1981). "Two recent theories of conditionals," in W. L. Harper, , and (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 Dekker (ed.), *Proceedings of the Fourteenth Amsterdam Colloquium (Amsterdam: ILLC).*

Gillies, A. S. (2004). "New foundations for epistemic change," *Synthese,* **138**: 1–48.

Grice, H. P. (1975). "Logic and conversation," in P. Cole and (eds.), *Syntax and Semantics (New York: Academic Press).*

Groenendijk, J. and M. Stokhof (1991). "Dynamic predicate logic," *Linguistics and Philosophy* **14**: 39–100.

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

Hanson, W. (1991). "Indicative conditionals are truth-functional," *Mind* **100**: 53–72.

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

Jackson, F. (1979). "On assertion and indicative conditionals," *Philosophical Review* **85**(4): 565–589.

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

Lewis, D. (1976). "Probabilities of conditionals and conditional probability," *Philosophical Review* **85**(3): 297–315.

Lewis, D. (1979). "Scorekeeping in a language game," *Journal of Philosophical Logic* **8**: 339–359.

Lewis, D. (1996). "Elusive knowledge," *Australasian Journal of Philosophy* **74**(4): 549–567.

McDermott, M. (1996). "On the truth conditions of certain 'if'-sentences," *Philosophical Review* **105**(1): 1–37.

McGee, V. (1985). "A counterexample to modus ponens," *Journal of Philosophy* **82**(9): 462–471.

Morreau, M. (1992). "Epistemic semantics for counterfactuals," *Journal of Philosophical Logic* **21**(1): 33–62.

Nute, D. (1980). *Topics in Conditional Logic. Dordrecht: Reidel.*

Ramsey, F. P. (1931). "General propositions and causality," in *Foundations of Mathematics and Other Essays (New York: Routledge and Kegan Paul).*

Stalnaker, R. (1968). "A theory of conditionals," in N. Rescher (ed.), *Studies in Logical Theory (Oxford: Basil Blackwell), pp. 98–112.*

Stalnaker, R. (1975). "Indicative conditionals," *Philosophia* **5**(3): 269–286.

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.