EMBEDDING 'IF AND ONLY IF'

NOVEMBER 4, 2007

Abstract

Many hold that indicative conditionals do not express propositions. Of the different ways to motivate this view, considerations to do with embedding have perhaps been the most influential. The line goes something like this: if conditionals express propositions then they should embed freely under sentential operators. But we have two reasons to say that they do not freely embed. First, the linguistic data suggest that conditionals do not embed meaningfully in many cases, and second, embedding conditionals is crucial to the proofs of triviality results like Lewis'.

We challenge this way of supporting the view that conditionals do not express propositions. The crux of our argument is that *if and only if* embeds much more freely than *if/then*, but is subject to a triviality result like the one for *if/then*. These facts suggest that one side of the dual-support story described in the previous paragraph is mistaken. If the moral of Lewis' triviality result is that *if/then* does not express a proposition, then a triviality result for *if and only if* should mean the same for *if and only if*. But then why does *if and only if* embed successfully in the very contexts where *if/then* does not embed?

1. INTRODUCTION

It is a common view about indicative conditionals that they do not express propositions, but instead express the speaker's high conditional degree of belief in the consequent given the antecedent. When I say that if Sergio makes coffee then he'll spill some, I merely express an epistemic attitude; specifically, that my probability that Sergio will spill some coffee, given that he is brewing a pot, is high. But I do not assert a proposition to that effect, and correlatively, I do not utter anything that can properly be evaluated for truth or falsity.

Of the different ways to motivate expressivism about conditionals, considerations to do with embedding have perhaps been the most influential. The line goes something like this: if conditionals express propositions, then we should expect them to embed freely under sentential operators. This is an unwelcome consequence for two reasons. First, the linguistic data suggest that conditionals do not embed freely. Second, embedding conditionals is crucial to the proofs of triviality results like Lewis' (1976). But expressivism does not assign propositions to conditionals, so it is not committed to free embedding and its unattractive consequences, lending the view dual support.[1]

We want to challenge this way of supporting the expressivist view. The crux of our challenge will be that *if and only if* embeds much more freely than *if/then*, but is subject to a triviality result like the one for *if/then*. This fact suggests that one side of the dual-support story described in the previous paragraph is mistaken. If the moral of Lewis' result is that *if/then* does not express a proposition, then a triviality result for *if and only if* should mean the same for *if and only if*. But then why does *if and only if* embed successfully in the very contexts where *if/then* does not embed?

[1]Edgington is one example of an author who endorses both sides of this dual-support line; she explicitly endorses Lewis-type results to support the view that conditionals do not express propositions (Edgington, 1995, pp. 273-6), and also says that "reflections about [the interpretability of] compounds support the conclusion that conditionals don't have truth conditions." (Edgington, 2000, p. 115) She prefers the conclusion that conditionals are conditional assertions to the view that they express conditional degrees of belief, but our arguments apply to the conditional assertion view as well. For definiteness, we will focus on expressivism.

Bennett (2003) is an expressivist who likewise appeals to Lewis-style triviality results (pp. 103-4), and cites the linguistic evidence that conditionals do not embed freely in a positive light (pp. 95-8). He is, however, not entirely explicit that the linguistic data support expressivism, as opposed to just being consonant with expressivism.

After reviewing the background behind expressivism in the following section, we begin the meat of our argument in section 3. There we argue that *if and only if* embeds much more freely than *if/then*; most importantly, it embeds in the very same sentences where *if/then* does not embed. Then, in section 4, we argue that a Lewis-style triviality result extends to *if and only if*. Having set out the materials of the argument, we conclude in section 5 with a discussion of its workings and potential implications.

2. Expressivism and Embedding

Expressivism is a conjunction of two theses, one semantic and the other pragmatic:

Non-Propositionality Conditionals do not express propositions.

Expression Conditionals express high conditional probability.

These claims are logically independent and, strictly speaking, it is the former that is supported by the linguistic data about embedding and the triviality results, since it is the semantic claim of Non-Propositionality that blocks free embedding. If conditionals express propositions, then those propositions ought to feed into the semantics of sentential operators the same as any other proposition, to determine a meaning for the whole sentence. If they do not, then the operator has no propositional material to operate on, and so no meaningful whole should result. This, claims Gibbard (1981), is what happens in examples like

(1) *If Kripke was there if Strawson was there, Anscombe was there.[2]

[2]Notice that interpretability is not helped if we move to the *if if A then B, then C* formulation, instead of the *if A if B, then C* formulation that Gibbard uses. We will follow Gibbard and others in focusing on the *if A if B, then C* formulation here and in subsequent examples; readers who are concerned that something may hang on this choice should check to their satisfaction that there is no significant difference in each instance.

Gibbard claims that, if he tells us this about the latest conference, we do not know what he has told us. It is hard to say what circumstances would make this utterance true or false, which is what Non-Propositionality would lead us to expect.[3]

Embedded conditionals do seem perfectly interpretable in some cases. Conjunctions provide easy examples, but nesting conditionals can be made to work too:

(2) If John's company will buy him a car if he does well this year, then they care about employee satisfaction.[4]

Expressivists account for such cases by saying that they have special readings "that do not come straight from the conventional meanings of the sentences." (Bennett, 2003, p. 96). In the case of (2), the story would be that we understand the embedded conditional as stating the factual basis that would make that conditional acceptable, namely that John's company has some policy or intention to reward good performance. When we hear (2) then, we understand it as saying something something to the effect of:

(3) If John's company has a policy of rewarding high performers, then they care about employee satisfaction.

But interpreting (2) as (3) is a matter of pragmatics. (2) cannot express the same proposition as (3) because (2) does not express a proposition at all.

Non-Propositionality, together with a little pragmatic footwork, predicts the data we find when we try to embed *if/then*. But Expression is indirectly supported too, since it is a natural way to understand the function of indicative

[3]For the sake of argument, we will be taking a good number of controversial claims for granted. For example, we will assume that embeddings like (1) do sound as bad as the expressivist claims, even though there are certainly those who think otherwise. We will also assume that the Ramsey Test (see below) is highly plausible, even though there are putative counterexamples; see, fore example, (Hájek and Hall, 1994) and (Nolan, 2003).

We also limit our claims about embedding to one of the prime cases the expressivist is interested in: the antecedents of conditionals. See (Edgington, 1995) for other putative cases of embedding failure. (We are, frankly, unimpressed by any of the other cases.)

[4]Thanks to Jennifer Nagel for the example.

conditionals if Non-Propositionality is true. This is especially so, given the plausibility of the Ramsey Test, according to which a conditional's acceptability always goes with its conditional probability:

Ramsey Test The acceptability of $A \to B$ for an agent always matches her estimate of $p(B|A)$.[5]

The Ramsey Test is acknowledged as highly plausible by many,[6] and is accounted for by Expression on the expressivist view.

The Ramsey Test is especially important because it drives the argument for Lewis' triviality result, which is supposed to offer further support for Non-Propositionality. If we reject Non-Propositionality and treat conditionals as stating propositions, then the natural way to account for the Ramsey Test is to suppose that the probability of a conditional is always the same as its corresponding conditional probability:

The Equation $p(A \to B) = p(B|A)$ for any A and B.

If we reject Non-Propositionality and accept The Equation, we can derive absurdity as follows. Together with the assumption of conditionalization,[7] The Equation gives us the result that for any A, B, and C,

$$p(A \to B|C) \quad = \quad p(B|A \wedge C). \tag{4}$$

If we let C be B this gives us

$$p(A \to B|B) \quad = \quad p(B|A \wedge B) = 1. \tag{5}$$

[5] Strictly speaking, the Ramsey Test says something weaker: that we evaluate $A \to B$ by supposing A, making the minimal necessary adjustments in our epistemic corpus, and seeing whether belief in B results. Our formulation adds the assumption that supposing and minimally adjusting works, or should work, by Bayesian conditionalization. This is a common assumption in the literature, and is innocuous enough for now. Later, in section 5.3, we will see one way in which it is not so innocuous.

[6] Again, as we said earlier (fn. 3), for the sake of argument we are taking several controversial claims for granted. The Ramsey Test is one of them.

[7] Conditionalization says that the rational response to the evidence that C is to adopt one's old conditional credence in X as one's new unconditional credence in X, for all X.

Similarly, letting C be $\neg B$ gives us

$$p(A \to B|\neg B) \;=\; p(B|A \wedge \neg B) = 0. \tag{6}$$

Now apply the theorem of total probability to $p(A \to B)$ to get

$$\begin{aligned} p(A \to B) &= p(A \to B|B)p(B) + p(A \to B|\neg B)p(\neg B) \\ &= 1 \cdot p(B) + 0 \cdot p(\neg B) \\ &= p(B). \end{aligned} \tag{7}$$

Thus we are led from the rejection of Non-Propositionality and acceptance of The Equation to the absurd conclusion that the probability of a conditional is always the same as the probability of its consequent.

But if we accept Non-Propositionality, the argument can be blocked. Given Non-Propositionality, we may say that $(A \to B) \wedge C$ fails to express a proposition, so it does not make sense to assign it a probability. But the proof turns on assigning the probability $p(A \to B|C)$, which is defined by the ratio $p((A \to B) \wedge C)/p(C)$. Non-Propositionality frees us to say that the numerator here is senseless, and so are lines (4)-(7) of the proof for the same reason.

That is, in compact form, the dual-support story for expressivism. By embracing Non-Propositionality we free ourselves to reject, as meaningful, embeddings of conditionals, and this agrees with some of the linguistic data while simultaneously blocking the triviality proof. Since the Ramsey Test is highly plausible, it is natural to account for it in the context of Non-Propositionality by appealing to Expression. Non-Propositionality and Expression are thus jointly supported by the linguistic data and Lewis' triviality result, pushing us towards expressivism about conditionals.

3. Embedding 'If and Only If'

We claim that *if and only if* embeds more freely than *if/then*, even though it is subject to similar triviality results. Let's see the embedding data first. Consider the following parallel to Gibbard's (1):

(8) If Strawson was there if and only if Kripke was there, Anscombe was there.

(8) is fully interpretable. We can well imagine circumstances in which the antecedent of (8) is false and yet we would accuse Gibbard of speaking falsely were he to tell us (8). If Anscombe's plans did not involve Kripke or Strawson in any way, nor did anything affecting her decision, then (8) is false. It does not matter whether Kripke went if and only if Strawson did. If, for example, Anscombe had the unimpeded plan to attend just in case she did not sleep through her alarm that morning, then we would reject (8). The comparison with (1) is striking and the ease of comprehension clear.[8]

It is tempting to dismiss (8) on the grounds that the biconditional is a term of art, which our formal training has taught us to use as shorthand for the material biconditional. But whether or not *if and only if* is a part of colloquial language, it is a part of academic language, and the same sorts of arguments and considerations that bother *if/then* in English can be used to bother *if and only if* in the augmented English of academia. Furthermore, it is false that *if and only if* is academic English for the material biconditional. The material biconditional

(9) The Toronto Maple Leafs won the 1978 Stanley Cup $\equiv$ the price of milk was \$99 per liter that year.

is true because both sides are false, but the 'natural' biconditional

(10) The Toronto Maple Leafs won the 1978 Stanley Cup if and only if the price of milk was \$99 per liter that year.

is not at all acceptable. Indeed, the obvious conjecture, that a biconditional $A \leftrightarrow B$ is acceptable only if both $A \rightarrow B$ and $B \rightarrow A$ are acceptable, is very reliable. So

[8]Notice that the same judgments apply if we use *just in case* in place of *if and only if*.

it is no surprise that $\leftrightarrow$ cannot be analyzed as $\equiv$ (at least, no more successfully than $\rightarrow$ can be analyzed as $\supset$).

Expressivists do have a story to tell when *if/then* successfully embeds, and they may try to apply the same strategy to explain why (8) is acceptable. Recall, we supposedly find (2) acceptable because we substitute a factual basis for the embedded conditional *John's company will buy him a car if he does well*. Likewise, the expressivist will say, we find (8) acceptable because the locution *if and only if* suggests a factual basis that can be substituted into the antecedent. But why should we expect that *if and only if* would be more suggestive of a factual basis than *if/then*? It is also mysterious what the factual basis would even be in (8). We can imagine various things that might make Kripke's attendance and Strawson's attendance mutually dependent, and on which Anscombe's attendance could depend in turn. But which of these do we substitute? There is no obvious candidate. Do we then substitute the disjunction of all of them? If so, why don't we simply do this for all embedded conditionals? Whenever a conditional is embedded, we could interpret it by substituting for the embedded conditional the disjunction of all the propositions that would lead us to find that conditional acceptable. We are again at a loss to explain the difference between (1) and (8).

It is no fluke that (8) is better than (1). Here are a few more pairs to illustrate:

(11) *If Jimmy has a cat if he has a dog, then he has a parrot.

(12) If Jimmy has a cat if and only if he has a dog, then he has a parrot.

and

(13) *If Mary has tuberculosis if the test paper shows green, they are testing for the lethal kind.

(14) If Mary has tuberculosis if and only if the test paper shows green, they are testing for the lethal kind.

In each of these pairs, the latter example reads more acceptably than the first.

4. Triviality for 'If and Only If'

Given that *if and only if* embeds much more readily than *if/then*, the expressivist should be tempted to withhold Non-Propositionality for *if and only if*. And yet, doing so leads to triviality much the same as it did for *if/then*. We can get a triviality result for *if and only if* in a couple of ways. The first is to assume a conjunctive analysis of $A \leftrightarrow B$, according to which it is equivalent to $(A \rightarrow B) \wedge (B \rightarrow A)$. Then we use (4)-(6) from above to derive

$$\begin{aligned} p(A \leftrightarrow B|B) &= p((A \rightarrow B) \wedge (B \rightarrow A)|B) \\ &= p(B \rightarrow A|B) \\ &= p(A|B). \end{aligned} \tag{15}$$

Similarly, (5) and (6) also yield

$$\begin{aligned} p(A \leftrightarrow B|\neg B) &= p((A \rightarrow B) \wedge (B \rightarrow A)|\neg B) \\ &= 0. \end{aligned} \tag{16}$$

Applying the theorem of total probability to $A \leftrightarrow B$, we then have

$$\begin{aligned} p(A \leftrightarrow B) &= p(A \leftrightarrow B|B)p(B) + p(A \leftrightarrow B|\neg B)p(\neg B) \\ &= p(A|B)p(B) + 0 \cdot p(\neg B) \\ &= p(A \wedge B). \end{aligned} \tag{17}$$

So we have the absurd result that the probability of a biconditional is always the same as the probability of the corresponding conjunction.

The conjunctive analysis of $A \leftrightarrow B$ as $(A \rightarrow B) \wedge (B \rightarrow A)$ is fairly plausible but could be questioned. All we really need to derive an absurd result, however, is that the biconditional have nearly the same probability as the conjunction of the conditionals, and this is borne out by examples. But if that does not satisfy, we can avoid the detour via $\rightarrow$ and The Equation and work directly with an analogue of The Equation for $\leftrightarrow$.

There are some obvious candidate analogues of The Equation for $\leftrightarrow$. The acceptability of $A \leftrightarrow B$ naturally goes with the conditional probabilities $p(B|A)$ and $p(A|B)$, so you might think that the general rule is

$$p(A \leftrightarrow B) = \min[p(B|A), p(A|B)] \tag{18}$$

or maybe

$$p(A \leftrightarrow B) = p(B|A) \times p(A|B). \tag{19}$$

To avoid hanging the argument on the details, let's not assume either the minimum rule or the product rule, and instead go with the general assumption that

$$p(A \leftrightarrow B) = f(p(B|A), p(A|B)) \tag{20}$$

where f is a monotonically increasing function on both arguments, with $f(x,1) = x$ and $f(x,0) = 0$. Then we have:

$$\begin{aligned} p(A \leftrightarrow B|B) &= f(p(A|B \wedge B), p(B|A \wedge B)) \\ &= f(p(A|B), 1) \\ &= p(A|B) \end{aligned} \tag{21}$$

and

$$\begin{aligned} p(A \leftrightarrow B|\neg B) &= f(p(A|B \wedge \neg B), p(B|A \wedge \neg B)) \\ &= f(p(A|B \wedge \neg B), 0) \\ &= 0. \end{aligned} \tag{22}$$

Applying the theorem of total probability as per usual, we get

$$\begin{aligned} p(A \leftrightarrow B) &= p(A \leftrightarrow B|B)p(B) + p(A \leftrightarrow B|\neg B)p(\neg B) \\ &= p(A|B)p(B) + 0 \cdot p(\neg B) \\ &= p(A \wedge B) \end{aligned} \tag{23}$$

as before.

The assumption that f be monotonically increasing is very reasonable but clearly does no work in the proof. What matters is just that $f(x,1) = x$ and $f(x,0) = 0$. These facts hold of both the minimum and product rules, and are plausible general constraints. If $p(A|B)$ is certain, then the only obvious reason to doubt $A \leftrightarrow B$ is (roughly) that A might be true but B false. So the result should hang entirely on $p(B|A)$. And if A is certainly false if B is true, then $A \leftrightarrow B$ doesn't even have a shot, so it should have probability 0. There is room to quibble about the details here, but not much, since the monotonicity assumption will pick up some of the slack if the constraints are challenged. Thanks to that assumption, slight alterations of these constraints will leave the triviality result only slightly improved, and drastic alterations would be hard to motivate. Between this argument and the previous one, $\leftrightarrow$ looks to be squarely in the same boat as $\rightarrow$ with respect to triviality.

5. Discussion

The challenge we put to the expressivist's dual-support story is this: if the moral of triviality results is that conditionals do not embed because they do not express propositions, why does *if and only if* embed meaningfully despite having a triviality result? We saw that the expressivist has a pragmatic story to tell about cases like (2), but it failed to deliver the proper cut between cases such as (1) and (2), as well as between (1) and (8). We've also seen that the pragmatic story has a great deal of trouble explaining the sheer inability to make any natural sense of sentences like (1).

In fairness, opponents of Non-Propositionality may have a hard time accounting for some the same data. For them, the default expectation is that embeddings will be meaningful, so they owe an explanation why cases like (1) are so hard to make sense of. Without a fuller understanding of the terrain, the appropriate preliminary conclusion would seem to be that things are more or less a draw for now. Until a comprehensive story is available that handles both the linguistic data and the triviality problem, neither party has a clear advantage. Regrettably, we have no fully worked out story to offer. We do have several directions for further research, none of which seem to fully capture all the data we'd like to capture. We offer them here as suggestions to be developed.

5.1 Parsing

There are some famous cases in which a perfectly interpretable sentence appears to be uninterpretable. Most famous amongst these are the 'garden path sentences', like:

(24) The horse raced past the barn fell.

(24) sounds awfully bad, but is dramatically improved when understood as saying:

(25) The horse that was raced past the barn fell.

Perhaps there is something similar at work in cases like (1). The repetition of 'if', for example, may be confusing to parsing the sentence in a way that requires some pragmatic help for interpretation. We might take some solace in the fact that it is somewhat difficult to interpret similar sentences,[9] such as:

(26) ?Because Mary came because Shiela came, Steve came.

Compare (26) with the much more acceptable (27):

(27) Because Mary came because and only because Sheila came, Steve came.

Compare also, in this vein:

(28) ?When Mary went when Sheila went, Sally went. (or: *When when Sheila went, Mary went, Sally went.)

(29) When Mary went when only when Sheila went, Sally went.

[9]We are grateful to John Hawthorne here.

(28) is not easily interpreted. (29) is admittedly pretty awkward but to our ear much easier to interpret. Perhaps there is an explanation for the badness of sentences like (1) that is based on considerations to do with parsing rather than the semantics or syntax.

One might complain that this fails to explain the felicity of cases like the company-car conditional, (2). This is a fair complaint, but it is important to note that not all sentences that are syntactically like (24) are equally hard to interpret, e.g.:

(30) The car cleaned in the barn died.

(31) ?The man cleaned in the barn died.

(30) is much easier to interpret than (31) or (24) because 'the car' isn't an appropriate agentive subject for the putative verb phrase, 'cleaned in the barn', making a reading parallel to (25) the only available reading. It is not merely syntactic form that is responsible for (mis)leading the parser down a garden path – it is also the availability of the garden path reading.[10]

5.2 Semi-Factual Semantics

The parsing story takes successful embedding to be the default expectation, and explains away exceptions as parsing confusions. Because this approach leaves embeddings meaningful, if hard to understand, the expressivist's way out of triviality results is unavailable to it. How might we supplement the parsing story to escape triviality? Taking a cue from the expressivist, maybe what we need here is a semantics that is intermediate between a straight-forwardly factual semantics for conditionals and denying that they have any semantic value

[10] As Jessica Rett has pointed out to us, phonetic prominence might play a role in these cases. She notes that the following sounds much better than (1):

(32) ?If Kripke was there if, you know, Strawson was there, then Anscombe was there.

We are inclined to think that (32) is still hard to interpret, but we agree that adding material helps. This is consistent with a general parsing story.

at all. Suppose we could assign semantic values to conditionals that were not strictly factual. Then we might be able to make semantic sense of embeddings, but reject the probabilities assigned in Lewis-type proofs; either on the grounds that these non-factual contents are not eligible to bear probabilities, or on the grounds that they simply will not bear the particular probabilities the proof needs. We illustrate this approach here.

Inspiration for this idea comes from Gibbard's own work on expressivism in ethics. Expressivists about ethics think that ethical statements serve to express normative attitudes rather than to state normative facts. A notorious problem for ethical expressivists is that ethical sentences do embed quite freely, so they would seem to have a semantic value of some kind. Gibbard (1990) famously handles the problem by assigning non-standard propositions as the semantic values of ethical sentences — propositions whose contents are a mixture of factual and normative elements.

Start with standard possible-world semantics, on which the semantic value of a sentence is the set of possible worlds where the sentence is true. Then augment the machinery by adding to the set of possible worlds the set of *normative systems*, where a normative system is a complete specification of everything that is permitted, forbidden, and required; given a norm system n, every action in every possible scenario is n-permitted, n-forbidden, or n-required. We then pair each possible world with each possible norm-system to get the set of *fact-norm worlds*, $\langle w, n \rangle$. An ethical sentence can then be assigned the set of fact-norm worlds where it 'holds' as its semantic value. For example, the sentence

(33) You should give Johnny his wagon back.

is assigned the set of $\langle w, n \rangle$ pairs where you giving Johnny his wagon back is n-required at w. We now have a semantic value for ethical sentences that allows them to be meaningfully embedded under sentential operators, using the same rules of composition as we would in ordinary possible-worlds semantics.

To achieve a similar effect for conditionals, we need to find something like a set of fact-norm worlds to use as the semantic-value of a conditional. And we need it to be something that will not bear probabilities in the way they are used in Lewis' proof. Since an indicative conditional is, for the expressivist, supposed to say something like "my conditional probability $p(B|A)$ is high", an obvious thing to try is to assign to $A \rightarrow B$ the set of centered-worlds $w_{t,\alpha}$ where the speaker is α, the time of utterance is t, and α's degree of belief $p(B|A)$ is high at t in $w_{t,\alpha}$. But this makes the contents strictly factual, and of the wrong sort — conditionals are now reports about your own mental state, which is surely wrong.

More in line with Gibbard's approach would be to separate out the epistemic component that conditionals are supposed to express and use it in the way that Gibbard uses the normative component of ethical sentences: as the non-factual half of a mixed factual/non-factual semantics. The analogue of n here would be a system of epistemic norms, specifying what degree of belief one ought to have in any given situation. Let's use the probability function p to represent the system of epistemic norms on which $p(H|E)$ is the degree of belief you ought to have in H in a situation where your evidence is E. Our mixed worlds are then pairs of centered worlds and probability functions, $\langle w_{t,\alpha}, p \rangle$. A conditional $A \rightarrow B$ is true at $\langle w_{t,\alpha}, p \rangle$ just in case $p(B|A \wedge E)$ is high, where E is α's total evidence at t in $w_{t,\alpha}$. The intuitive idea is: a conditional conveys the agent's endorsement of a set of epistemic norms and her presumption of an evidential state that together mandate having a high $p(B|A)$.

This gives us a semantic value that is not strictly factual, but is it enough to block Lewis' proof? These contents may be able to bear probabilities; maybe $p(A \rightarrow B)$ can be made sense of as the plausibility that the correct norms, together with one's evidence, demand a high $p(B|A)$. But even if conditionals can bear probabilities on this semantics, they will not bear probabilities in a way that vindicates The Equation. Suppose that you have $p(B|A) = 1/2$. What should your $p(A \rightarrow B)$ be? That is, how confident should you be that your evidence and the correct epistemic perspective together recommend a high conditional $p(B|A)$? Not confident at all, under ordinary circumstances. If you are confident that you are rational, and confident that you have taken account of all your evidence appropriately, then you ought to be confident that your degree of belief is correct. And if your degree of belief is $p(B|A) = 1/2$, then you should think it very unlikely that it should be higher. So for you, $p(A \rightarrow B)$ might well be near 0, while $p(B|A)$ is $1/2$. Nothing about rationality or linguistic convention could mandate The Equation on this view, breaking the hold of Lewis' proof.[11]

This particular proposal is not trouble-free. There are concerns about the general legitimacy of mixed semantics, for example. One could also exploit the cases where The Equation is not obeyed to make trouble. On this proposal, your degree of belief that a fair coin will come up heads if flipped should be 0, since the conditional

(34) The coin will come up heads if it is flipped.

says, roughly, that you ought to have a high confidence in heads given that the coin is flipped. Since you know the coin is fair, you should think that content highly improbable, a result that is, at least prima facie, in tension with the thought that the right answer to the question "What is the probability that the coin will come up heads if flipped?" is "50%".[12] But whether or not such concerns are surmountable, the proposal at least serves its illustrative purpose. Other ways of executing the same basic strategy may do better.

[11] It is noteworthy that this strategy, if successful as a way out of Lewis' triviality result, also provides a way out of Edgington's (1995) argument for Non-Propositionality. Her argument depends on the following inference: because being certain of $A \vee C$ but not of A is sufficient for being certain that $\neg A \rightarrow C$, it must not be possible for $\neg A \rightarrow C$ to be false when $\neg A \supset C$ is true. On the mixed semantics, however, this inference is invalid. The fact that being certain of $A \vee C$ but not A is sufficient for certainty that $\neg A \rightarrow C$ only tells us that every combination of norms and evidence that warrants certainty in $A \vee C$ but not A also mandates a high $p(C|A)$. And this is true so long as we assume that all eligible norm-systems respect the probability axioms, which we have done just in virtue of using probability functions to represent norm-systems.

[12] Thanks to Daniel Nolan for making this point.

5.3 More Pragmatics

Those who prefer to avoid these semantic shenanigans might yet be able to get by with a pragmatic story, if they can just tell one that has more teeth than the factual-basis story we criticized earlier. After all, there is a distinct feeling that (1) contains an antecedent whose relevance to the consequent is far from clear, while the relevance in cases like (2) is patent. Something like this observation seems to drive Gibbard towards the factual-basis story, and telling a pragmatic story that gets the right cuts may just be a matter of developing the point more carefully. To that end, let us consider how clarity of relevance might suffer in the course of applying the Ramsey Test.

According to the Ramsey Test, we evaluate a conditional by supposing the antecedent, adjusting our epistemic household accordingly, and checking to see whether the consequent is believed after the adjustment is made. It would be surprising if this procedure always yielded a determinate result. Can it really be that it is always clear what the appropriate adjustments are and what fate they hold for the consequent? So long as we see the Ramsey Test in terms of probabilistic conditioning, we get the illusion that there will always be a determinate answer. If our epistemic household is represented by a probability function, and the process of adjustment is conditionalization, then the posterior probability of the consequent is always determinate. But this thoroughgoing determinacy is very likely an artifact of Bayesian idealization. Real agents will often have to work with step-by-step inferential reasoning, and for them the Ramsey Test will more closely resemble evaluating something in premise-conclusion format.

So let's take this as an alternate model of the Ramsey Test: we take the conditional, break it up into premise-conclusion format, and see whether the resulting argument is any good.[13] Sometimes it will be obvious whether the conclusion

[13] We don't intend for the notion of 'good argument' to bear much weight. As Daniel Nolan has pointed out to us, there may be all sorts of reasons to infer a conclusion from a conditional premise and background knowledge such that the premise does little work; maybe you infer the conclusion because you already knew it before hearing the conditional premise. Notice that this wouldn't help a great deal with explaining why (1) is bad. Even if you already know

could plausibly be inferred from the premise, but other times it won't be. If the resulting argument is badly enthymematic, for example, the verdict may be indeterminate. There will be various ways of 'connecting the dots' between the premise and the conclusion, with varying degrees of effectiveness and plausibility, none of which is obviously what the speaker has in mind. In such cases, evaluating the conditional becomes a messy process, with interpretability taking a substantial hit. Perhaps this is what happens in cases like (1), but not (2).

The challenge for this story is to get the cuts in the right places; not just between (1) and (2), but also between (1) and (8). Let's look at (1) and (2) first. If we break up the main conditional into an argument structure, we have:

(P1) If Strawson was there, Kripke was there.

(C1) Anscombe was there.

as compared with:

(P2) If John works hard, his company will buy him a car.

(C2) They care about their employees.

One wouldn't ordinarily infer (C1) from (P1) as they presumably would (C2) from (P2). However, if one knew that Anscombe liked going to conferences when she knew a Strawson-Kripke conditional to be true, (C1) becomes a much less objectionable inference. Correlatively, (1) sounds sensible with this background in place; in fact, it sounds true! So far so good.

The story is less obviously effective when comparing (1) and (8). Take the biconditional:

(35) If Jason is coming iff Sarah is coming, then Barry is coming.

The following seems like a funny inference to draw:

---

that Anscombe is coming, you won't thereby think the conditional is licensed. So the premise has to do some work in justifying the conclusion. This isn't wholly satisfactory but we think an intuitive grasp of 'good argument' is sufficient in these cases to at least point to further research.

(P3) Jason is coming iff Sarah is coming.

(C3) Barry is coming.

So why is the *if and only if* case interpretable but not the *if/then* case? Here we might productively combine the pragmatic story with the parsing story told in 5.1 to explain the difference: unlike their conditional cousins, the *if and only if* cases are sufficiently easy to parse that one can supply a story pragmatically to connect the dots. But parsing causes enough trouble in the *if/then* cases that, in the presence of the additional pragmatic burden of connecting the dots, interpretability fails altogether.

A worry: if (1) is uninterpretable in part because of indeterminacy in the Ramsey Test, why does such indeterminacy only show up when we apply the Ramsey Test to nested conditionals? Why doesn't the same indeterminacy infect simple conditionals? Here we can only offer a very tentative conjecture: that supposing simple antecedents is a common and familiar epistemic procedure, while supposing conditional states of affairs is highly irregular, and thus uncertain. Most of our knowledge and evidence comes to us in the form of simple facts — Kripke was at the conference, John's company bought him a car, the cat is on the mat, etc. Cases where we acquire knowledge of a conditional directly, as opposed to inferring it from other facts, are comparatively rare. So it is to be expected that adjusting our epistemic household to accommodate a conditional is an unfamiliar task, and we are often much less clear on how to go about it.[14]

[14]It is telling, in this connection, the rarity and controversiality of cases where probabilities must be updated directly on some conditional probability, as opposed to updating on some evidential proposition (whether certain, as in standard conditionalization, or uncertain, as in Jeffrey conditionalization). For example, van Fraassen (1981) argues that conditionalization does not apply in his Judy Benjamin case, and so we need some rule for updating on conditional probabilities directly. But Grove and Halpern (1997) object that the Judy Benjamin case has a straightforward treatment in terms of standard conditionalization, one that avoids the complications that come with trying to update on conditional probabilities directly in the Judy Benjamin case. Bradley (2005) offers other cases to motivate updating directly on conditional probabilities, and he suggests Adams Conditioning as a rule for handling them. Whatever position we take on such cases, the point is that they are hard to come by, and that the correct way to handle them is not obvious.

6. Conclusion

These considerations are obviously programmatic rather than substantial. However, we think that the embedding of conditionals is a very difficult topic and the data is extremely interesting and subtle, even if ultimately it does not provide a vindication of expressivism about conditionals. Our aim has been to separate the linguistic facts about embedding from the Lewis triviality proof (and its cousins). This should in no way decrease the interesting and challenging questions posed by sentences such as (1).

References

Bennett, Jonathan. 2003. *A Philosophical Guide to Conditionals*. Oxford University Press.

Edgington, Dorothy. 1995. "On Conditionals." *Mind* 104:235–329.

Edgington, Dorothy. 2000. "General Conditional Statements: A Reply to Kölbel." *Mind* 109:109–16.

Gibbard, Allan. 1981. Two Recent Theories of Conditionals. In *Ifs*, ed. William L. Harper, Robert Stalnaker and Glenn Pearce. Reidel.

Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Harvard University Press.

Hájek, Alan and Ned Hall. 1994. The Hypothesis of the Conditional Construal of Conditional Probabilities. In *Probability and Conditionals*. Cambridge University Press.

Lewis, David. 1976. "Probabilities of Conditionals and Conditional Probabilities." *The Philosophical Review* LXXXV(3):297–315.

Nolan, Daniel. 2003. "Defending a Possible-Worlds Account of Indicative Conditionals." *Philosophical Studies* 116(3):215–69.