Conditional Acceptance

Ofra Magidor, University of Oxford

Forthcoming in Oxford Studies for the Philosophy of Language

Introduction

Providing an adequate semantics for indicative conditionals is a notoriously difficult task. In this paper I present a new challenge to the task of providing semantics for indicative conditionals, and the related epistemic question of when one ought to accept a conditional statement.

The challenge takes the following form: In **§1**, I present a case involving an utterance of a certain indicative conditional *C*. In **§2**, I argue that at least each of three prominent theories of conditionals predict that the in this case, you should assign a credence of at least 0.9 to *C*. In **§3**, I argue that this prediction is wrong: given the case, it is entirely permissible for you to assign a credence which is lower than 0.9 to *C*. In **§4**, I discuss what conclusions we can draw from the argument both to the semantics of conditionals, but also to epistemology more generally. The appendix contains the more technical details of the probabilistic analysis of the case.

§1 The case

You are visiting an exotic island, and you are certain of (that is, you have credence 1 in) the following set-up:

1. <u>The Island</u>: the island contains exactly two kinds of inhabitants: Randomers and Reliabilists. You do not have any way to distinguish between the two kinds of inhabitants other than by relying on the features explained below.

2. <u>Randomers:</u> Randomers are odd people: they have their own beliefs but they keep them entirely to themselves. When they speak, they pick some subject matter *P* and flip a fair coin in order to decide whether to utter *P* or to utter ∼*P*. Which sentence they utter is neither a reflection of their credences nor of the facts in the world.[1]

[1] Some might find such cases of insincere utterances inappropriate for semantic theorising. In that case, here is a variant which will do just as well for my purposes: Randomers only utter sentences they are highly confident in, but they form their beliefs entirely randomly. (On this variant, you should either maintain that the beliefs of Randomers are not probabilistically coherent, or alternatively, that the entire probabilistic state of a Randomer changes frequently by mechanisms other than conditionalization.)

3. <u>Reliabilists</u>*:* Reliabilists, on the other hand, utter (declarative) sentences only when they are highly confident in them. Moreover, while Reliabilists are not omniscient and sometimes do utter falsehoods, they are highly rational and extremely well-informed. You take them to be experts, and when you are certain someone is a Reliabilist, you attempt to defer to their beliefs, including partial beliefs.[2]

For a start, let us set up the case so that, subject to the usual caveats, your credence in *P* conditional on *X* being a Reliabilist who utters *P* is precisely 0.9 (call this 'The Simple Case').[3] How might we defend such conditional priors? For a start note that on the standard Bayesian picture your conditional prior must be some precise number or other: there seems to be no reason we cannot set this number to 0.9 in particular. Moreover, such conditional priors might be motivated by the following picture: suppose you think that Reliabilists only utters *P* if their credence in *P* is at least 0.9 (and have no additional information about their credences). One way of trying to defer to their expertise is to adopt the (somewhat cautious) strategy of always adopting a conditional credence of 0.9 in such cases.

Some, however, seem to find this constraint on your conditional credences too strict. I will thus also present another version of the case (call it 'The Generalized Case'), where your credence in *P* conditional on X being a Reliabilists who uttered *P* are *at least* 0.9.[4] (As we shall see, the shift from the Simple to the Generalized case makes no difference to my argument in §2, and while it does somewhat complicate the argument in §3, the argument ultimately goes through equally well.)

There is a final complication here: the utterance which will be at the centre of my case is an indicative conditional. This fact makes no difference to the set-up given any theory of conditionals which maintains that they express propositions. But one of the prominent theories I will discuss below (The Suppositional Theory) maintains that conditionals do not express propositions. Moreover, while proponents of the theory often avail themselves of talk of 'the probability of a conditional', there are good reasons to think that such "probabilities" cannot be fully represented

---

[2] To make the case a bit more realistic, we can suppose that there are some subject-matters on which you don't take Reliabilists to be experts. Thus for example, you might not defer to Reliabilists on the question of whether you have a tooth-ache. But I will assume that the conditional which is at the centre of this case is one of those many subject-matters on which you do defer to Reliabilists.

[3] The caveats are that since $Pr(A|B) = \frac{\Pr(A \wedge B)}{\Pr(B)}$, the conditional probability cannot be equal to 0.9 if $Pr(A)=1$, or if $Pr(A \wedge B) = 0$ (and consequently if either $Pr(A) = 0$ or $Pr(B) = 0$).

[4] This time, we'll only need the caveat that your prior credence in '*P* and *X is a Reliabilist who uttered P*' is not null.

using the standard Bayesian picture as above.[5] Indeed, on one way of interpreting the Suppositional Theory, conditionals do not strictly speaking have probabilities at all. A conditional statement 'If A then B' can be more or less "accepted", and it accepted to precisely the degree of one's conditional credence in B given A. Relatedly, on this interpretation a conditional is assertible in so far as the speaker's conditional credence in B given A is high, and the role of the assertion is to get the hearer to also adopt a similarly high conditional credence.

While this picture does not allow us to (strictly speaking) accept my assumption that conditional on X being a Reliabilist who utters a conditional you should have a credence of (exactly or at least) 0.9 in the conditional, it can allow for a very close analogue of this assumption: suppose you are certain that Reliabilists only assert conditionals if they have a high conditional credence in B given A. Thus when you are certain that X is a Reliabilist who uttered a conditional 'If A then B', you should (subject to the usual caveats as above), also adopt a high conditional credence in B given A. Relatedly, if you revise your current credences by conditionalizing on the claim that X is a Reliabilist who uttered *P*, the resulting credence function would be one on which your degree of acceptance of the conditional is 0.9 (The Simple case), or at least 0.9 (The Generalized case).

4. <u>The distribution of islanders:</u> The intuitions in the case become most vivid, I think, if you suppose that the vast majority of the islanders are Randomers, and a small minority are Reliabilists. As it turns out, however, for the purposes of the Simple Case, *any* non-trivial distribution is sufficient for my argument, and for the purposes of the Generalised Case, a distribution of for example 50% Reliabilist/50% Randomers will do. (See §3 and the Appendix for discussion).

5. <u>The utterance:</u>

   Walking around the island you meet a stranger S. The stranger utters the following conditional:

   (*C*) If I am a Reliabilist then I have high-blood pressure.

This completes the set-up of the case. The question is now whether, upon hearing S's utterance, you ought to adopt a high credence in *C*. In what follows, I argue that each of three highly prominent

[5] For example, on the Suppositional Theory, the "probability" of a conditional, conditional on the negation of its antecedent is undefined, even if the latter receives positive probability.

theories on the semantics of conditionals predict that you should adopt a credence of at least 0.9 in *C*, but that this prediction is incorrect.[6]

§1. Prediction of the prominent theories

There are three general positions concerning the semantics of indicative conditionals. The first, maintains that the conditional is *truth-functional* (the truth-value of a conditional depends only on the truth-values of the antecedent and consequent) and it is easy to show that the only candidate truth-function here is that of the material conditional.[7] According to this view, 'If A then B' is true just in case '¬A∨B' is true.[8]

The second maintains that the conditional is not truth-functional, but does have truth conditions. Probably the most influential theory of this form, at least among philosophers, is that of Robert Stalnaker.[9] On Stalnaker's view, 'If A then B' is true if and only if B is true in the selected A-world. Generally, the relevant A-world is selected to be the "closest" one to the actual world.[10] However, a crucial constraint on the ordering-function in the case of indicative conditionals is that if there are any A-worlds in the context-set (roughly: the set of worlds that are considered open possibilities in the conversation), then the selected *A-world* must be in context-set. Thus if *A* is not ruled out in the conversation, 'If A then B' is true just in case B is true in the closest A-world in the context set.

The third position maintains that indicative conditionals do not even have truth-conditions (or at least, they have at most partial truth-conditions). On this view, although conditionals might have systematic semantic values, these do not take the form of standard propositions. The most prominent variant of this view is the Suppositional Theory.[11] According to the Suppositional Theory, the central principle

[6] Some readers might have the suspicion that this case is related to the semantic paradoxes, perhaps in particular Curry's paradox. I do not think this objection is on the right track. In so far as this is merely raised as suspicion rather than a concrete argument, it is hard to diffuse - but let me make a few comments that might help appease these worries. First, I see no reason at all to suppose the conditional in question is paradoxical. Second, there needn't be anything in the content of *C* that (either directly or indirectly) refers to *C* or its truth-value. (It might be helpful to realise that the labels 'Reliabilist' is merely a helpful notation – the group need not be defined by its reliability properties – it might refer to members of a particular gender or religion. Furthermore, the case only depends on the fact that you have credence one in the claim that Relibilists have the reliability properties described above, nothing in the case requires them to *actually* have them). Finally, there is no reason to suppose that *C* is in some way ungrounded. Of course, the question of what exactly grounds *C* depends on which semantics for conditionals is adopted, but the same kind of facts that make standard indicative conditionals such as 'If John is British, he has high blood pressure' true or acceptable, should serve equally for the case of *C*.

[7] See Edgington (1995), p. 242 for a defence of this claim.

[8] Notable defenders of the truth-functional view include Grice (1989) and Jackson (1987). See Edington (2009) for a summary of the main arguments in favour of truth-functionality.

[9] See Stalnaker (1968) and Stalnaker (1975).

[10] I write 'closest' in quotes because the ordering function might be highly context sensitive, and in particular depend on the conditional in question.

[11] For defences of this theory see Adams (1975), Edgington (1995), and Bennett (2003).

governing indicative conditionals is that the credence one should assign to the conditional 'If A then B' should be the credence one assigns to the consequent B conditional on the antecedent A (this principle is often summarised as 'the probability of the conditional is the conditional probability' and is also referred to as 'Stalnaker's Thesis' or 'Adams's Thesis').[12] Adam's Thesis is connected to the view that conditionals are non-propositional in two ways: first, because the principle places a direct constraint on one's credence in the conditional, rather than offering truth-conditions coupled with the general principle that one's credence in a proposition is the credence that it is true. This means that the Suppositional Theory is at least compatible with a non-propositional approach. Second, and more importantly, a range of triviality results show that (subject to some plausible background assumptions), one cannot satisfy Adam's Thesis while also maintaining that conditional express propositions.[13]

What do each of the three views predict about your credence in *C* following the utterance in the case?[14]

Start with the material conditional view. According to this view, *C* is true just in case either S is not a Reliabalists or S has high blood pressure. Let 'R' denote the proposition that S is a Reliabilist, and let '*P*' denote your credence function *after* updating on S's utterance (i.e. your posterior credence function). For a start, note that by the Law of Total Probability $P(C) = P(C|R) \cdot P(R) + P(C|{\sim}R) \cdot P({\sim}R)$. Since $P(R) + P({\sim}R) = 1$, if both $P(C|R) \geq 0.9$ and $P(C|{\sim}R) \geq 0.9$, it follows that $P(C) \geq 0.9$. And this is indeed the case on the material conditional view: conditional on S being a non-Reliabilist, the conditional is trivially true and should receive credence 1. Conditional on S being a Reliabilist, since S uttered *C*, you should (in light of the set-up) give *C* a credence of at least 0.9. So your overall credence in the conditional should be at least 0.9.[15]

Next, consider the Suppositional Theory (I leave the discussion of Stalnaker's view to the end, because it's the trickiest case). According to the Suppositional Theory, the credence (or whatever stands proxy

[12] Though as I noted in the introduction, on some interpretations of the view, we shouldn't think of agents as strictly speaking having credences in conditionals. If you prefer this interpretation, replace 'degree of acceptance' with 'credence' in the conditional in the discussions of the view.
[13] See Lewis (1976) for the original triviality results, and Bennett (2003), §5 for a helpful summary of various extensions.
[14] I will discuss the Simple and the Generalised case in tandem here, as the difference between them does not matter for the purposes of the current discussion.
[15] It is worth noting that in addition to the truth-conditions, Jackson and Grice each offer some pragmatic constraints on utterances of a conditional. One might suggest that we should take these into account when considering your credence in *C*, at least for the case where you suppose S is a Reliabilist. But in the case of Jackson, the constraint is simply that the speaker's credence obey Adam's thesis, which would still give us a credence of 0.9 (see the discussion of the Suppositional Theory below). And in the case of Grice, one can easily cancel the relevant conversational implicatures (e.g. by adding some special pragmatic explanation for why S uttered *C* rather than one of its disjuncts. One simple explanation is that you simply asked S specifically about *C*).

for credence) that you should assign to *C* is the conditional credence you have in the consequent ('S has high blood pressure' – abbreviate this with 'HBP') conditional on the antecedent (namely, R). Let P_R be the credence function you obtain from P by conditionalizing on *R.* On the assumption that S is a Reliabilist, you ought, given the set-up, to have a credence (or credence-proxy) of at least 0.9 in *C,* or in other words, you should have a conditional credence of at least 0.9 in the consequent conditional on the antecedent. This entails that $P_R\ (HBP\ |R) \geq 0.9$. On the other hand, by the definition of P_R , $P_R(HBP|R) = P_R(HBP) = P(HBP|R)$. Thus, $P(HBP|R) \geq 0.9$ as required.

Finally, consider Stalnaker's view. According to the view, the conditional is true just in case in the closest context-set world where S is a Reliabilist (assuming there is one), S has high-blood pressure. First, note that there certainly are context-set-worlds where S is a Reliabilist. (You have no reason to rule out that S is a Reliabilists prior to hearing S's utterance, and you also have no reason to rule it out after the utterance: after all you have no reason to assume S's utterance is false in this case, and even if you did, that would not give you grounds for ruling out because Reliabilists sometimes utter falsehoods.) The second thing to note is that following S's utterance, it is part of the common ground in the conversation that S uttered *C.*[16] It follows that all worlds in the context-set are ones where S utters *C*.

Now first let us consider the set of context-set worlds in which S is a Reliabilist (that is, let us assess *C* conditional on S being a Reliabilist). As we've established, each of these is a world where S utters C. Given the set up, you have 0.9 credence that a world where S is a Reliabilist and S utters *C* is a world where *C* is true. Moreover, on Stalnaker's semantics, any world where S is a Reliabilist is a world where *C* is true just in case the consequent of *C* is true. Thus the worlds in which S is a Reliabalist are distributed so that 0.9 of the region they occupy is a HBP region.

Next, consider the probability of *C* conditional on S being a Randomer. Here things become somewhat trickier. As Lewis pointed out, on Stalnaker's semantics the probability of a conditional is not always equal to the conditional probability, precisely because the probability of the conditional over worlds where the antecedent is false might differ from its distribution over worlds where it is true.[17] Here is an example which illustrates how the two could come apart: you are certain that a person in the other room was holding a cup and dropped it (but you cannot see the cup). Now suppose you think there are two possibilities: either the cup is made of glass or it is made of pewter. Conditional on its being made of glass, you have a credence of 0.95 that it shattered. But suppose you also know that if the cup is made of pewter it must have a particular shape (say an oval shape), and moreover you think

[16] Note that Stalnaker is explicit about endorsing the claim that one (typically) adds to the context-set the fact that the particular utterance has been made (see e.g. Stalnaker (1999), p. 86).
[17] See Lewis's discussion of 'imaging' in Lewis (1976), pp. 308-12.

that glass cups which have an oval shape are much less fragile than usual (they have only 0.5 chance of shattering). Now consider the conditional: 'If the cup is made of glass it shattered'. On the Suppositional Theory, your credence in this conditional ought to be 0.95. But Stalnaker's view predicts otherwise: assuming that the closest world to each pewter-world is one where the cup's shape is held fixed, then conditional on the cup being made of pewter, the conditional 'If glass then shattered' will receive a credence of 0.5, and the overall probability of the conditional will be lower than the conditional probability (namely 0.95).

Return to our case. What I would like to argue is that, as long as we add an unproblematic background assumption, this case is importantly different then the pewter/glass case just described. All we need to add is that any feature which is held fixed in closest worlds is one that is probabilistically independent for you of the question of whether S has high blood-pressure. For example, you might think that for each Randomer world *w*, the closest Reliabilist world is one where S's gender is the same as it is in *w*. But that would make no difference if you thought the probability of S having high blood-pressure conditional on S being a Reliabilist is identical to the probability of S having high blood pressure conditional on S being a Reliabilist with a particular gender. Assuming such independence, then, for each world in the context-set where S is a Randomer, your probability that the closest-context-Reliabilist-world is a high-blood pressure world should just be 0.9. With the independence assumption in place, this case is thus analogous to a much simpler version of the cup-case: suppose that you think that glass cups are 0.95 likely to shatter when dropped, and otherwise think that the likelihood of shattering is only dependant on the material the cup is made of. In that case, your probability for the conditional 'If the cup is made of glass it shattered' would be 0.95, and the probability of the conditional would be identical to the conditional probability. While Stalnaker's view does not vindicate Adams's thesis across the board, it does validate it in a restricted range of cases, and with the proper set up it will validate it in our particular case.

A final worry might arise: there is one feature that is certainly *not* probabilistically independent for you of whether S has high-blood pressure: S's blood pressure levels. Suppose we require that the closest world to each world *w* matches *w* in S's blood-pressure levels? (After all, wouldn't a world which matches *w* in this way count as closer to *w* than a world which doesn't match it?...) If we had such a matching constraint then then your probabilities on S's blood-pressure should affect your credence that the closest *w*-world is a high-blood pressure world. Suppose for example that, on the supposition that S is a Randomer, you are 50-50 on whether S has high-blood pressure. Given the matching constraint, it follows that your credence that the in the closest-Reliabalist-world S has high blood-pressure should be 0.5, and our overall credence in C might be significantly lower than 0.9.

I have two responses to this worry. First response: I think it is highly implausible that we ever interpret 'closeness' in a way which requires a matching constraint on the truth of the consequent. Consider another variant of the cup-example: you have a cup which is in fact made of pewter. You drop it and (you are certain that) it didn't shatter. You also know that cups made of glass which are dropped in a similar fashion are .95 likely to shatter. Now it's not obvious how we should interpret closeness in these cases, but I think at lead a good indication is to look at the corresponding counterfactual judgements (we certainly shouldn't simply assess closeness *tout court* given how context sensitive the notion is, at since it's far less controversial that the semantics of counterfactuals involves something like Stalnaker's closest-world analysis this seems like a good place to look…). So consider the counterfactual 'If the cup had been made of glass it would have remained intact'. If closeness worked so as to make worlds that matched the actual world in facts about the consequent (i.e. shattering facts), then we should assign this counterfactual a credence of 1, but we certainly do not assess the counterfactual in this way. Similarly, suppose I toss a fair coin and it in fact lands on heads. Now consider the counterfactual 'If the coin had been 0.9 biased towards tails, it would have landed heads'. If closeness favoured matching the actual worlds in facts about the consequent, we assign this counterfactual a credence of 1, but clearly we do no such thing.[18] It seems, then, that matching facts about the consequent of a conditional simply cannot be requirement on how closeness is interpreted.

Second response to the worry: if despite my remarks above one still insists that closeness can involve matching facts about blood-pressure levels, then we can simply add to the set-up of the case the claim that you are certain that all Randomers have high-blood pressure (though you do not think there is any interesting connection between being a Reliabilist and blood-pressure levels). This will not make a difference to the rest of my argument[19], but now if we expect the closest worlds to match up in blood-pressure facts, this would only increase the probability that one should assign to *C*. (In this case, Stalnaker's theory would predict that conditional on S being a Randomer, your credence in *C* is 1 which is over 0.9, and thus the overall probability of *C* would be at least 0.9).

I conclude, then, that all three views predict that you should assign a credence of at least 0.9 to C.[20]

---

[18] There is of course a tricky question of which credence we assign to 'If the coin had been 0.9 biased towards tails it would have landed *tails*', but whatever you think of *that* question, it's clear that we do not assign a very high credence to the counter-factual with 'lands heads' in its consequent.

[19] It makes no difference at all to the predictions of the material conditional and Suppositional Theory, and will make no difference to my argument that these predictions are wrong in the next section.

[20] What about other views concerning the semantics of conditionals? Some notable alternatives that take conditionals have truth-conditions are ones that take 'If A then B' to be true just in case the corresponding material conditional is true in all worlds representing some contextually determined epistemic state (See Kratzer (1986) and Rothschild (2011) for views that take this form). These views are phrased in terms that are too general to predict what they would say about this specific case, but these theories should, I think, have the resources to avoid

§3. Why the prediction is wrong

In this section I argue that given the case, it is entirely permissible for you to have a credence of less than 0.9 in *C.* In §3.1 and §3.2 I argue for this conclusion assuming The Simple Case, and in §3.3 I explain how to extend my argument to The Generalized Case.

A preliminary about the dialectical force of my argument is in order: my argument in this section does not rely on any particular theory for the semantics of conditionals (after all, what are the correct semantics is precisely what is at stake here). Nor do I maintain that the considerations I rely on are fully consistent with the theories I reject (after all, the conclusion I reach is inconsistent with the prediction of these theories, we should expect the assumptions in the argument to be inconsistent as well…). Rather, I wish to tease out some pre-theoretic intuitions we have about conditionals which show that the prediction of the three theories is incorrect. I will do so, by arguing in turn for two claims:

<u>Claim One</u>: Prior to S's utterance, it is permissible for you to have a credence which is lower than 0.9 in *C*.

<u>Claim Two:</u> If your credence prior to the utterance is lower than 0.9, then your credence after the utterance should also lower than 0.9.

Together, these claims entail that it is permissible for your posterior credence to be lower than 0.9.

§3.1 Defending Claim One

Claim One is, I take it, highly intuitive. Think of the sorts of situations where one would give a high credence to *C*: one such situation is where you think that all Reliabilists have high-blood pressure. Another, perhaps, is where you think that merely a large-majority of Reliabilists have high blood-pressure. Other cases involve having high credence in S having some particular feature that is positively correlated, amongst Reliabilists, with high blood pressure (for example, you might be certain that all *tall* Reliabilists have high blood pressure, and be highly confident that S is tall). Perhaps you also give *C* high credence in the case where you merely have a very high credence that the conjunction

---

the problematic predictions here (cf. my discussion in §4). Another kind of view which might be classified as non-propositional is a trivalent account of the conditional (see Belnap (1970) and Rothschild (*forthcoming*)), – again, without further details it is hard to determine what such accounts will predict about this case, but in so far as these views wish to respect Adams's thesis (as Rothschild (*forthcoming*) proposes) they will deliver the same prediction as the Suppositional Theory.

of the antecedent and consequent are true (even if you don't think there is any interesting connection between the two conjuncts).

The crucial point, however, is that one can easily set up the case so that none of the above apply. Suppose that prior to the utterance you do not have a particularly high credence in either the antecedent or the consequent. Furthermore, you don't think Reliabilists are particularly prone to high blood pressure and there is no other feature (at least no other feature you have high credence in S possessing) which you think of as positively increasing the likelihood of S's having high blood pressure (if they are Reliabilists). Perhaps, in the absence of any such factors you would have credence of 0.5 in *C,* or perhaps you should adopt a much lower credence.[21] But either way it would be very odd to *require* that you must have a prior credence of over 0.9 in *C*.[22]

§3.2 Defending Claim Two

Here is why we should accept Claim Two: assume your prior credence in *C* is lower than 0.9. After hearing S's utterance, there are two cases to consider: first, suppose that S is a Randomer. On that hypothesis, the fact that S uttered *C* gives you no information concerning *C*, so you should stick to your prior credence in *C* (which, by assumption, is lower than 0.9). Second, suppose that S is a Reliabilist. On that hypothesis, upon hearing the utterance, your credence in *C* should be 0.9.[23] But since you are less than certain that S is a Reliabilist, your overall posterior probability should be less than 0.9.

Let me put this line of reasoning in slightly more precise probabilistic terms: let 'R' represent the proposition that S is a Reliabalist, and 'UC' the proposition that S uttered *C*. Let '*Pr*' denote your *prior* credence function (i.e. your credence function prior to S's utterance).[24] If you update your beliefs using standard conditionalization, your credence in *C* after the utterance should be $Pr(C|UC)$. But the Law of Total Probability, $Pr(C|UC) = Pr(C|UC \wedge R) \cdot Pr(R|UC) + \Pr(C|UC \wedge {\sim}R) \cdot \Pr({\sim}R|UC)$. But first, due to the set-up, $Pr(C|UC \wedge R) = 0.9$. Second, I claim that $Pr(C|{\sim}R) = Pr(C)$. Why is that? First, recall that we are discussing your *prior* credence function, thus S's uttering *C* should not make any difference here. One might try to argue that as *R* appears in the antecedent of *C*, learning *R* (and equally, learning ${\sim}R$) should affect the probability of *C*. But that is not very plausible: barring very special circumstances, the credence you should assign to a conditional should not depend on the

[21] Those who accept conditional excluded middle might be more tempted towards the view that the credence should be 0.5.
[22] It is worth noting that each of the theories discussed in §2 can also accept Claim One (though in the case of the material conditional view, one would need to assume that your prior in S being a Randomer isn't extremely high).
[23] Recall that for now I am focusing on the Simple Case.
[24] Note that this in contrast with the credence functions discussed in §2, which represented your *posterior* credences.

credence you assign to its antecedent.[25] For example, consider the conditional 'If John presses the button, the bomb will explode'. The credence you should assign to this conditional typically depends on factors such as whether you think the button is wired to the bomb, the bomb is live, and so forth. But how likely you think it is that John will decide to press the button should *not* impact your credence in the conditional (indeed, one of the chief criticisms of the material conditional view is that it blatantly violates this independence). Third, I maintain that $Pr(C|{\sim}R \wedge UC) = Pr(C|{\sim}R)$. The reason is that, supposing S is a Randomer, the fact that they uttered *C* should have no effect on your credence in C (after all, S's utterance of *C* is akin to reporting the result of a coin flip). Thus $Pr(C|{\sim}R \wedge UC) = Pr(C|{\sim}R) = Pr(C)$, which by assumption is less than 0.9. Thus as long as $Pr(R|UC) < 1$, we get that $Pr(C|UC) < 0.9$ (See the Appendix for the calculations of the precise probabilities *C* receives in the Generalized case).[26]

Having given my argument for Claim Two, it is instructive to see where the three theories discussed in the previous section get things wrong. Both the Material Conditional view and Stalnaker's theory get things wrong by assuming that *C gets* a high probability not only conditional on $UC \wedge R$ but also conditional on $UC \wedge {\sim}R$. The Material Conditional view gets this result simply because it assumes that, on your prior credence function, *C* is *not* independent of its antecedent: $Pr(C|{\sim}R) = 1$ and hence that $Pr(C|UC \wedge {\sim}R) = 1$. Stalnaker's view gets things wrong because, given the relevant background assumptions, it entails that $Pr(C|UC \wedge {\sim}R) = Pr(C|UC \wedge R)$.[27] Finally, The Suppositional Theory gets this wrong by assume that $Pr(C|UC \wedge {\sim}R)$ is entirely irrelevant to the posterior probability of the conditional (the view takes this probability to be undefined as it involves assessing a conditional whose antecedent is certain to be false), and takes the probability of the conditional to be directly equal to $Pr(C|R)$.[28]

[25] Cf. Rothschild (2011) and the discussion in §4 above. (As I show there, your credence function *after* conditionalizing on UC constitutes one such special circumstance.)

[26] Certain variants of epistemic externalism might maintain that you shouldn't update your credences by conditionalization in this case: if S is a Reliabilists, you should have a high credence in their utterance (whether or not you are certain that S is a Reliabilist). But even if we accept such views, one can simply stipulate that in the case (unbeknownst to you) S is a Randomer, and thus you ought not have a high credence in *C*.

[27] Here is another way to highlight why views such as Stalnaker's, which predict you should give *C* a high posterior because you give *C* a high posterior conditional on S being a Reliabilist are highly problematic. Consider a slightly modified version of the case, where instead of Randomers the island contains 'Yay-Sayers' – which, whenever asked about a claim *p*, they assent to it. Now suppose you are planning to ask S about *p*. In this case, $Pr(C|UC \wedge {\sim}R) = Pr(C|{\sim}R)$, because the claim that R is a not a Reliabilist entails that S is a Yay-Sayer and thus will utter *C*. But then, if your prior in S being a non-Relibilists is sufficiently high (e.g. because you think the island contains a lot more Yay-Sayers then Reliabilists), then your *prior in C* (i.e. your credence in C even before it was uttered), would be extremely high, which seems highly implausible.

[28] The fact that, on the Suppositional Theory, conditionals are undefined when conditionalizing on the negation of the antecedent follows from the more general feature that, according to the theory, conditionals are undefined when the antecedent receives probability zero. Here is the typical defence one gets from proponents of the theory for this general feature (from Bennett (2003), p.56): "conditionals are devices for intellectually managing states

§3.3 The Generalized Case

So far, I have discussed The Simpler Case. As we have seen, the posterior probability of C is the weighted average of $Pr(C|UC \wedge R)$ and $Pr(C|UC \wedge {\sim}R)$. In the Simpler Case, we know that the former is exactly 0.9 and the latter is lower than 0.9, and thus neither the precise values of these two nor the precise way they are weighted mattered to the defence of Claim Two. The same is not true for The Generalized Case: since in the Generalized Case $Pr(C|UC \wedge R)$ might be higher than 0.9, the weighted average might end up being higher than 0.9. However, even in the Generalized case, the overall posterior probability of the conditional will be lower than 0.9 provide that either your prior in *C* was sufficiently low (and thus $Pr(C|UC \wedge {\sim}R)$ is equally low) or your prior in *S* being a Randomer was sufficiently high (and thus $\Pr({\sim}R|UC)$ is sufficiently high).

In the Appendix, I give a full probabilistic analysis of the case. One can substitute various values, but here are a couple of sample values that are of interest. (See the Appendix for a full explanation of these results.) Let us suppose that $Pr(C|UC \wedge R) = 1$ (this is the "worst case scenario" for the argument in the Generalized Case). If your prior in the conditional is 0.5 (or lower) then any distribution of the islanders in which more than 20% of them are Randomers would suffice for the posterior probability of *C* to be lower than 0.9. Alternatively, if your prior distribution has 95% of the islanders be Randomers, then as long as your prior in the conditional is lower than 0.890626 the posterior would still be lower than 0.9. Since on the one hand, I can set the case to involve whatever (non-trivial) distribution if islanders we wish, and on the other hand, the same considerations in §3.1

of partial information, and for preparing for the advent of beliefs that one does not currently have. For an A that you regard as utterly ruled out, so that for you *P(A)*= 0, you have no disciplined way of making such preparations, no way of conducting the Ramsey test; you cannot say what the upshot is of adding to your belief system something you actually regard as having *no* chance of being true".
The problem with this defence is that it is irrelevant to cases (such as the one I am considering in the paper) where you are not genuinely giving the antecedent null credence, but rather you are assessing the conditional under the *supposition* that the antecedent is false. In that case you do not (unconditionally) regard the antecedent as having no chance of being true, and there is no reason why you cannot use your (unconditional) epistemic state to assess the plausibility of the conditional. Perhaps a helpful way to see the point is by analogy to the case of epistemic modals. For example, if I am genuinely certain that not-A, then 'might A' would trivially get credence zero, but the same would not be obviously true if I am *merely supposing* not-A. (Consider a case where I know that the murderer is either Jack or Jill but I don't know which. The following seems to have an acceptable reading: "Suppose Jack is not the murderer but Jill framed him by leaving lots of misleading evidence that he is. So each of Jack and Jill *might* be the murderer"...).Relatedly, if I am genuinely certain that Jack is not the murderer, then plausibly we cannot define my conditional probability of Jack going to jail conditional on the claim that he might be the murderer. But if I have a non-zero credence that Jack is the murderer, that arguably does not bar me from having some conditional probability on Jack going to jail, conditional on the supposition that he is not the murderer but might be the murderer...).

suggest that our prior in C should not be higher than 0.5 (and certainly not as high 0.890626!), the argument in the Generalized Case goes through as well.[29]

§4 Lessons and conclusions

§4.1 The semantics of conditionals

The main lesson of this paper concerns the semantics of indicative conditionals: each of the three prominent theories has been shown to generate the wrong prediction in this case, and any adequate theory needs to provide the correct prediction. Moreover, my discussion provides a novel kind of counter-example to Adams's thesis.[30] This is important, first – because Adams's thesis (or at least a restricted form of Adams's Thesis) is a consequence of a range of theories. But second, rejecting Adams's thesis doesn't only provide a general objection to views such as the Suppositional Theory. Rather it undermines the central motivation for adopting such non-propositional views (the idea was that in light of the triviality results, one can only maintain Adams's Thesis buy rejecting the claim that conditionals have truth-conditions. But if Adams's Thesis should anyhow be rejected, this motivation for going non-propositional no longer holds any sway).

This point also connects to a more specific issue concerning the semantics of conditionals. Rothschild recently argued that Lewis's triviality proof crucially relies on the following assumption: an indicative conditional is probabilistically independent of its antecedent, i.e. $Pr(If A\ then\ B|A) = Pr(If\ A\ then\ B)$.[31] Moreover, Rothschild maintains that while the assumption that the conditional is probabilistically independent of its antecedent is in most cases correct, it is not always correct: there are some counterexamples to this assumption and these counterexamples make for cases where Adam's Thesis fails.

Rothschild's discussion is very interesting, but I think the specific counterexample to independence he proposes is not entirely convincing (indeed, he himself expresses some doubts about it[32]). Here is his example: suppose you know the following: a certain type of car either has a defect in the airbag system, in which case whenever the car crashes the airbag fails to inflate; or it lacks the defect, in which case whenever the car crashes the airbag inflates. You are confronted with a car of the relevant

[29] Moreover, consider the conditional *C**: 'If S is a Reliabilist, then S does not have high blood pressure'. If one accepts Conditional Excluded Middle, then at least one of *C* and *C** should have a probability less than or equal to 0.5, and one can easily run the argument with S uttering *C** instead of *C*.
[30] See the discussion in n. 31 below on how my case differs from those of Kaufmann (2004) and Rothschild (2011).
[31] That is, Lewis's proof relies on Adams's thesis, but Rothschild argues that specific instance of the thesis that is used in Lewis's proof, amounts to accepting the weaker independence claim (see Rothschild (2011)).
[32] See his f.n. 31.

type, but you do not know whether the car has defect or not, or whether or not the car will crash. It seems that your credence in the conditional 'If this car crashes, the airbag will not inflate' is entirely independent of your credence you give to the antecedent (whether or not you assume the car in fact crashes, your credence in the conditional will simply be your credence in the car has the defect). But now suppose that you also know that cars which have the defect are also somewhat more likely to crash. Rothschild maintains that in that case the probability of the conditional is no longer independent of its antecedent. His reasoning is roughly the following: the probability you assign to the conditional is simply the probability you assign to the car having a defect. But if we conditionalize on the car crashing, then your credence on the car having the defect (and hence of the conditional) increases. The problem with this reasoning, however, is that it is not obvious that in the revised case (where you know cars with the defect are more likely to crash) the probability that you assign to the conditional (before conditionalizing on the antecedent) is still equal to the probability that you assign to the car having the defect: it is hard to know exactly how to assess such conditionals (and, as Rothschild suggests, they may well have different readings in different contexts), but it's not implausible that (at least one reading) your probability in *C* already takes into account the thought that situations in which the car crashes are situations where it is more likely to have the defect. Putting things otherwise: we know that $\Pr(C) = Pr(C|crashing) \cdot Pr(crashing) + Pr(C|not\ crashing) \cdot Pr(not\ crashing)$, so the crucial question is whether $Pr(C|not\ crashing) < Pr(C|crashing)$, and intuitions might go either way in this case.

Interestingly, though, the case I propose in this paper provides a cleaner counterexample to independence: since your posterior probability in *C* is (assuming a set-up as n §3.1) lower than 0.9, but your probability in *C* conditional on its antecedent is at least 0.9, independence fails. And this example seems to be clearer, because here we do have a specific argument for why the probability of *C* conditional on the negation of its antecedent should be low.[33]

[33] Rothschild's case also seems to be an instance of the challenges to Adams's thesis which are raised in Kaufmann (2004). But the case presented in this paper seems crucially different from Kaufmann's cases. On the Kaufmann style cases, the probability of the conditional is basically calculated as the conditional probability, *except* that one effectively assumes that some variable of the case is held fixed to the way it actually is (and thus that the probability of that variable is calculated via actual probabilities, rather than probabilities conditional on the antecedent). To take Rothschild's example above, the relevant variable is the car having or lacking a defect. Thus assume for example that you have a credence of 0.5 in car having the defect unconditionally, of 0.9 on its having the defect conditional on its crashing (and you are certain that when the car crashes, the airbag inflates if and only if it lacks the defect) Your conditional probability of the antecedent on the consequent is calculated as: $Pr(not\ inflate|carashes \wedge defect) * Pr(defect|crash)$+ $\Pr(not\ inflate|carashes \wedge no\ defect) * Pr(no\ defect|crash)$ = $1 * 0.9 + 0 * 0.1 = 0.9$. But if we're holding fixed the question of whether or not the car has a defect, then the probability of the conditional would be calculated as $\Pr(not\ inflate|carashes \wedge defect) * \Pr(defect)$+$\Pr(not\ inflate|carashes \wedge no\ defect) * \Pr(no\ defect)$ = 1*0.5 + 0*0.5 = 0.5.

§4.2 Conditional acceptance in philosophical arguments

I want to conclude with a somewhat more tentative note on some implications that this case might have to issues that go beyond the semantics of conditionals. Conditionals play a crucial role in philosophical arguments: often, philosophers find it difficult to defend outright philosophical views, and are much more confident about conditional claims of the form 'If this assumption is correct, then these consequences follow'. In particular, philosophers might be interested in conditionals where the antecedent recommends some logical theory, method of reasoning, or procedure for justification, and the consequent draws some conclusion from adopting this theory or method.

But while many such conditionals are surely acceptable, one should note that these types of conditionals are precisely at risk of having a similar structure to the one presented in this paper: if the reasoning which leads us from the antecedent to the consequent relies on employing the very method recommended by the antecedent, then our acceptance of the conditional will not be independent of our acceptance of the antecedent - and adopting such conditionals would not be neutral between all parties to the debate. (These will precisely be cases where one's credence in the conditional, conditional on the negation of the antecedent might be low, and those parties who have a very low prior in the antecedent, might thus have a low overall credence in the conditional as a whole.)[34]

Thus while the example presented in this paper might appear initially seem somewhat recherché, it has profound implications both to the semantics of conditionals as well as to philosophical arguments more generally.[35]

The crucial point is that the failure of Adams's thesis that is presented in this paper cannot be subsumed under the same analysis. The only relevant variable that one might try to "hold fixed" is the question of whether or not S has high blood-pressure. But, as with the discussion of Stalnaker's view in §2, we simply assume that your credence in Randomers having high blood-pressure is extremely high (if you want, 1), in which case holding fixed S's blood pressure will not help explain the particular failure of Adams's thesis argued for in §3.

[34] One particular case of this sort involves a variant of an argument proposed by Dummett against strict finitism (see Dummett (1975), Magidor (2007), and Magidor (2012)). Strict Finitism is an extreme form of constructivism in philosophy of mathematics, which takes as its key notion construability *in practice*. One upshot of the view is that proof-skeletons that would be too long to fill out in practice (e.g. ones that would take 2^{100} steps to fill out) do not count as legitimate proofs by the finitist's lights. Call a proof skeleton 'short enough' (SE) is it short enough to be filled-out in practice. I cannot go into the details here, but it turns out that one can provide, for each n, an allegedly fintistically acceptable argument for the conditional claim that 'If SE(2^n) then SE(2^{n+1})', and that moreover, accepting these conditionals leads to contradiction, even by the finitist's own lights. The interesting fact, though, is the form that each of these arguments take: it start with the supposition that SE(2^n), then uses 2^n steps, and reaches the conclusion that SE(2^n). The idea is that since each argument uses 2^n steps, then *under the supposition* of the antecedent (namely that 2^n is short enough), the argument should be finitstically acceptable and thus the finitist should accept the consequent. My contention, however, is that the finitist should not accept the conditional on this basis, for precisely the same reason that you are not required to accept *C* in the current case: since for some values of n they don't in fact accept that SE(2^n) they shouldn't accept the relevant conditional.

[35] Thanks to audiences in Carnegie Mellon, Krakow, London, NYU, Rutgers, and York as well as to Cian Dorr, John Hawthorne, Nicholas Jones, Sarah Moss, Daniel Rothschild, James Studd, and Timothy Williamson for helpful discussion of this paper. Thanks also to the Leverhulme Trust for financial support.

Appendix: A probabilistic analysis of the case

In this Appendix show how to calculate the precise probability of *C* in this case. Let *'Pr'* represent your prior probability function. As above 'R' represents the proposition that S is a Relibilist and 'UC' the proposition that S uttered *C*. Let *r* be what you assume is the percentage of Reliabilists in the island (i.e. $r = Pr(R)$), *c* your prior credence in *C* (i.e. $c = Pr(C)$). Finally, let k be your credence in *C* assuming that it was uttered by a Reliabilist, i.e. $k = Pr(C|UC \wedge R)$. In the Simpler Case we know that $k = 0.9$, but in the Generalized case we merely know that $k \geq 0.9$.[36]

As discussed in §3 above, your posterior credence in C should be $Pr(C|UC) = \frac{Pr(C \wedge UC)}{Pr(UC)}$. For simplicity of the discussion, assume that you are certain in advance that S will utter either *C* or ∼C.[37]

Now let us partition the space where UC is true into four regions: $UC \wedge C \wedge R$; $UC \wedge C \wedge {\sim}R$; $UC \wedge {\sim}C \wedge R$; $UC \wedge {\sim}C \wedge {\sim}R$. Note that $Pr(C \wedge UC)$ is the sum of probabilities of the first two regions, and that $Pr(UC)$ is the sum of probabilities for all four regions. Now let's calculate the probabilities of each of these regions in turn.

<u>Step 1: calculating $\boldsymbol{UC \wedge C \wedge R}$</u>

We start by noting that $Pr(UC \wedge C \wedge R) = Pr(UC \wedge C|R) \cdot Pr(R)$

Now, let Pr_R be the probability function we get from Pr after conditionalizing on R.

So $Pr(UC \wedge C|R)$ = $Pr_R(UC \wedge C)$= $Pr_R(C|UC) \cdot Pr_R(UC)$

But we know from the setup of the case that $Pr_R(C|UC) = k$, so $Pr_R(UC \wedge C) = k \cdot Pr_R(UC)$.

We also know that $Pr_R(C|{\sim}UC) = 1 - k$ (because if S does not utter C, S utters ∼C and then ∼C has a probability of k, see n. 34).

Finally, we know that $Pr_R(C \wedge {\sim}UC) = Pr_R(C) - Pr_R(C \wedge UC)$.

[36] I will not assume that your credence in a proposition conditional on it being uttered by a Reliabilist is always uniform, but to simply the calculations I will assume that the probability of ¬*C* conditional on a Reliabilist uttering ¬*C* is also *k*.

[37] This assumption makes no difference at all to the calculations. Let '*DC*' be the proposition that S discussed *C* (i.e. uttered either *C* or ∼C). Then since UC entails DC, $Pr(UC) = Pr(UC|DC) \cdot \mathrm{Pr}(DC)$ and $Pr(UC \wedge C) = Pr(UC \wedge C|DC) \cdot \mathrm{Pr}(DC)$, so $Pr(C|UC) = \frac{Pr(UC \wedge C|DC) \cdot \mathrm{Pr}(DC)}{Pr(UC|DC) \cdot \mathrm{Pr}(DC)} = \frac{Pr(UC \wedge C|DC)}{Pr(UC|DC)}$, so we can assume to have conditionalized on *DC*.

So we have:

$(1-k) = \ Pr_R(C|{\sim}UC) = \frac{Pr_R(C\wedge{\sim}UC)}{Pr_R({\sim}UC)} = \frac{Pr_R(C)-Pr_R(C\wedge UC)}{1-Pr_R(UC)} = \frac{Pr_R(C)-k\cdot Pr_R(UC)}{1-Pr_R(UC)}$

We thus get that $Pr_R(UC) = (Pr_R(C) + k - 1))/2k - 1$.

But, as argued in §3.2, on your prior credence function, in the conditional is independent of antecedent, so $Pr_R(C) = Pr(C|R) = Pr(C) = c$.

Thus $Pr_R(UC) = (c + k - 1))/(2k - 1)$

And to conclude we have:

$Pr(UC \wedge C \wedge R) = Pr(UC \wedge C|R) \cdot P\,r(R) = k \cdot Pr_R(UC) \cdot r = \ k \cdot (c + k - 1))/(2k - 1) \cdot r$

Step 2: calculating $\boldsymbol{UC} \wedge \boldsymbol{C} \wedge {\sim}\boldsymbol{R}$

We start by noting that $Pr(C \wedge UC \wedge \neg R) = Pr(UC|C \wedge {\sim}R) \cdot Pr(C \wedge {\sim}R)$.

Now, on the supposition that S is not a Reliabilist, S is a Randomer and thus flips a fair coin in order to decide whether to utter C or ~C. So $Pr(UC|C \wedge {\sim}R) = 0.5$. Moreover, since (as argued above), on the prior credence function *C* and *R* are independent, $Pr(C \wedge {\sim}R) = Pr(C) \cdot P\,r({\sim}R)$

Finally, we know that $Pr({\sim}R) = \ 1 - r$

So we conclude that:

$Pr(C \wedge UC \wedge {\sim}R) = \ 0.5 \cdot c \cdot (1 - r)$.

Step 3: calculating $\boldsymbol{UC} \wedge \neg\boldsymbol{C} \wedge \boldsymbol{R}$

$\mathrm{Pr}(UC \wedge {\sim}C \wedge R) = Pr(UC \wedge {\sim}C|R) \cdot Pr({\sim}R)$

But using our calculations in Step 1, we have $Pr(UC \wedge {\sim}C|R) = Pr_R(UC \wedge {\sim}C) = \ Pr_R({\sim}C|UC) \cdot Pr_R(UC) = (1 - k) \cdot (c + k - 1))/(2k - 1)$.

So we conclude that

$\mathrm{Pr}(UC \wedge {\sim}C \wedge R) = (1 - k) \cdot (c + k - 1))/(2k - 1) \cdot (1 - r)$.

Step 4: calculating $\boldsymbol{UC} \wedge {\sim}\boldsymbol{C} \wedge {\sim}\boldsymbol{R}$

Using similar considerations to those in Step 2 we get:

$$Pr(UC \wedge {\sim}C \wedge {\sim}R) = Pr(UC|{\sim}C \wedge {\sim}R) \cdot P\,r({\sim}C \wedge {\sim}R) = Pr(UC|{\sim}C \wedge {\sim}R) \cdot P\,r({\sim}C) \cdot Pr({\sim}R)$$
$$= 0.5 \cdot (1-c) \cdot (1-r)$$

<u>Step 5: calculating $\mathbf{Pr}(\boldsymbol{C}|\boldsymbol{UC})$</u>

Putting everything together we get:

$$\Pr(C|UC) = \frac{k \cdot (c+k-1))/(2k-1) \cdot r + 0.5 \cdot c \cdot (1-r)}{k \cdot (c+k-1))/(2k-1) \cdot r + 0.5 \cdot c \cdot (1-r) + (1-k) \cdot (c+k-1))/(2k-1) \cdot (1-r) + 0.5 \cdot (1-c) \cdot (1-r)}$$

This ratio is somewhat complicated, but can be simplified once particular values are substituted for *c*, *k*, and *r*. The ratio simplifies particularly nicely if we assume that *k*=1 (as in the most difficult version of the Generlized Case). In that case we get:

$$\Pr(C|UC) = \frac{c \cdot r + 0.5 \cdot c \cdot (1-r)}{c \cdot r + 0.5 \cdot c \cdot (1-r) + 0.5 \cdot (1-c) \cdot (1-r)} = \frac{c \cdot r + c}{2c \cdot r + 1 - r}$$

It is now easy to see why (as noted at the end of §3.3), if $c = 0.5$ then $Pr(C|UC) < 0.9$ just in case $r < 0.8$, and that if $r = 0.05$ then $Pr(C|UC) < 0.9$ just in case $c < 0.890625$.

References

Adams, E. (1975): *The Logic of Conditionals: an application of probability to deductive logic*, Dordrecht.

Belnap, N. (1970): 'Conditional assertion and restricted quantification', *Nous* 4: 1-12.

Bennett, J. (2003): *A Philosophical Guide to Conditionals*, Oxford University Press.

Dummett, M. (1975): 'Wang's Paradox', *Synthese* 30: 301-23.

Grice, P. (1989): 'Logic and conversation', in his *Studies in the Way of Words*, Harvard University Press.

Edgington, D. (1995): 'On conditionals', *Mind* 104: 239-273.

Edgington, D. (2009): 'Conditionals'¸ in Edward Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Spring 2009 edition.

Jackson, S. F. (1987): *Conditionals*, Basil Blackwell.

Kaufmann, S. (2004): 'Conditioning against the grain: abduction and indicative conditionals', *Journal of Philosophical Logic* 33: 583-606.

Kratzer, A. (1986): 'Conditionals', *Chicago Linguistics Society* 22: 1-15.

Lewis, D. (1976): 'Probabilities of conditional and conditional probabilities', *Philosophical Review* 8: 297-315.

Magidor, O. (2007): 'Strict finitism refuted?', *Proceedings of the Aristotelean Society* 107: 403-411.

Magidor, O. (2012): 'Strict finitism and the happy sorites', *Journal of Philosophical Logic* 41: 471-491.

Putnam, H. (1981): *Reasons, Truth, and History*, Cambridge University Press.

Rothschild, D. (2011): 'Do indicative conditionals express propositions?', *Nous* 47: 49-68.

Rothschild, D. (*forthcoming*): 'Capturing the connection between conditionals and conditional probability in trivalent semantics', *forthcoming* in *Journal of Applied non-classical logics*.

Stalnaker, R. (1968): 'A theory of conditionals', in Nicholas Rescher (ed.), *Studies in Logical Theory*, Blackwell: 98-112.

Stalnaker, R. (1975): 'Indicative conditionals', *Philosophia* 5: 269-286.

Stalnaker, R. (1999): *Context and Content*, Oxford University Press.