

Conditionals, Literal Content, and ‘DeRose’s Thesis’: A Reply to Barnett

KEITH DEROSE

Yale University

keith.derose@yale.edu

Against Barnett (????), I argue that the theory I advance in DeRose (2010) is best construed as one on which “‘were’ed-up’ future-directed conditionals like ‘If the house were not to be painted, it would soon look quite shabby’ are, in ways important to how they function in deliberation, different in literal content from their ‘straightforward’ counterparts like ‘If the house is not painted, it will soon look quite shabby’. I also defend my way of classifying future-directed conditionals against an attack by Barnett by defending a standard (among philosophers) approach to the basic structure of some conditionals. Finally, I counter Barnett’s charge that by using the concept of deliberation in my account of the meaning of future-directed conditionals, I put an implausibly demanding constraint on what it takes to understand such sentences.

1. Importantly different in content!

One of the main issues I address in ‘The Conditionals of Deliberation’ (DeRose 2010, henceforth ‘CoD’) is the relation between what I call ‘straightforward’ future-directed conditionals and their “‘were’ed-up’ counterparts. David Barnett, quite fittingly, then, begins his discussion of my paper with three scenarios featuring such future-directed conditionals (‘FDCs’). His scenarios all involve an agent deciding whether to blow on the candles on a birthday cake in front of her. The first two scenarios are the same, except that the protagonist deliberates on the basis of the straightforward ‘If I blow on the candles, I will look like a fool’ in the first one, but on the basis of the ‘were’ed-up ‘If I were to blow on the candles, I would look like a fool’ in the second. These are examples in which these two conditionals at least seem to amount to much the same thing; Barnett writes: ‘Scenarios 1 and 2 are hard to tell apart, for they differ only with respect to **Will** and **Would**, which are themselves hard to tell apart’ (p. ???). But these conditionals turn out to be hard to tell apart in meaning only in certain situations; Scenario 3 is supposed to be one of the different situations that show that these conditionals can amount to very different things, and where their difference in meaning is important to how they get used in deliberation. Now, I think Barnett’s third case fairly clearly fails to show what he wants it to show,¹ but we should not be

¹ Barnett’s key verdicts about his Scenario 3 are that the character would (and presumably should) have a low degree of confidence in **Would**, a high degree of confidence in **Will**; and that she should be guided in her practical

distracted by that. Straightforward conditionals can come apart from their ‘were’ed-up cousins in the ways Barnett is seeking to show (as we will see a bit later). Indeed, examples from CoD show this, and a better constructed example would have worked for Barnett’s conditionals.

So, at least where reflection involves a better example than his Scenario 3, I was right on board with the moral Barnett drew from his opening examples:

So, while **Will** and **Would** are initially hard to tell apart, on reflection there is an important difference between them, one that has significant implications for decision theory. (p. ???)

I thought that this was setting up our common ground, before getting to what might separate us. Imagine my shock, then, when what immediately followed was ‘Keith DeRose (2010) disagrees,’ and my dismay as I read on and discovered that this was no easy-to-factor-out glitch, but that much of what followed was an attempt to show that I was wrong to disagree with something that I in fact not only agreed with, but actually took to be a pretty nice statement of a key view of mine!

Barnett writes, ‘On DeRose’s thesis, **Will** is the very same idea as **Would**’ (p. ???). I believe, then, that I don’t accept ‘DeRose’s thesis’!² My uncertainty here stems from a lack of

deliberation by her low confidence in **Would**, rather than her high confidence in **Will**, and she should therefore blow on the candles. My reactions, and that of some very smart folks to whom I put Barnett’s case, is that the example is underscribed in a way that makes it unclear whether the agent should blow on the candles and very unclear how confident she should be in either **Would** or **Will**, and that neither of them is at all clearly deserving of more confidence than the other.

² ‘DeRose’s thesis’ is first discussed in the last four paragraphs of Barnett’s introductory material (before his section 1), where it seems to refer to the claim (a) that straightforward FDCs and their ‘were’ed-up counterparts express the same ideas. I’m inclined to reject that thesis. Then in the second paragraph of his section 1, Barnett specifies ‘DeRose’s thesis’ to be a combination of two claims: (b) that straightforward FDCs and their ‘were’ed-up counterparts have the same meaning, except for two differences that I specify, and (c) that both types of FDCs should be classified as indicative conditionals. Barnett rejects all of (a)-(c). I do not reject either (b) or (c). [I have no strong allegiance to the letter of (b): I would not be surprised or particularly dismayed if there turned out to be other differences between the two types of FDCs. My main focus is on the second difference I posit, and my main allegiance here is to the idea that this difference in meaning can well account for the differences in behavior of the two types of conditionals in the kinds of uses that I am looking at. The first difference I posit is (so far as I can see) a fairly unrelated difference that explains some quite different phenomena that I might well have missed had I not considered these other phenomena before. I would be far from shocked if there turned out to be other such differences that have escaped my notice.] And since Barnett may intend (b) and (c) to together constitute his official construal of ‘DeRose’s thesis’, I may not reject ‘DeRose’s thesis’, officially so-designated. But what Barnett does with (b) is argue that it leads to (a), and then much of his paper is spent directly attacking (a). (This also applies to his introductory material, where we have not yet encountered (b), and Barnett is already claiming that ‘There is much, however, that DeRose’s thesis cannot explain’ (p. ???), and aiming this at (a), rather than (b).) In the last few pages of his paper (section 3), Barnett gives ‘reasons to reject DeRose’s thesis’ that target (c) and that target (b)

confidence in my grasp of what Barnett means by ‘idea’, but given how I’m inclined to understand him, my view is that ‘were’ed-up conditionals express quite *different* ideas from their straightforward counterparts; that this difference can be *very* important to how these conditionals are useful in deliberation (in such a way that the ‘were’ed-up ones are ‘especially well-suited’ (CoD, p. 1) for use in deliberation); but that in many other situations, the differences don’t matter, so the conditionals can then *appear* to be similar or the same in meaning. I sought a theory that could explain all that.

Here’s how Barnett explains his usage:

Philosophers sometimes use ‘literal content’, ‘proposition literally expressed’, or ‘sense’ to designate the idea literally expressed by an utterance. In this paper, I deliberately avoid the term ‘proposition’ to avoid the presumption that what is expressed by a sentence has any of the features typically associated with propositions so-called, for instance, objective truth-conditions. Instead, I use the term ‘idea’. (p. ???)

Not certain just how much Barnett’s ‘ideas’ are supposed to differ from propositions, I have to be somewhat tentative in my denial that ‘were’ed-up and straightforward FDCs express the same ‘ideas’. But I think a reasonably good case can be made, anyway.

In a *somewhat* similar vein, in CoD I remained neutral on the propositionality of indicative conditionals. This largely because I think there are good reasons for doubting that indicative conditionals express propositions—though there are also ways in which they behave (and least relative to certain assumptions) as if they do express propositions (as we will see shortly).

Why does Barnett think I hold that, or am committed to it that, a ‘were’ed-up FDC expresses the same idea as does its straightforward counterpart? He attempts to derive this from my proposal that ‘were’ed-up FDCs be understood as having the same meaning as their straightforward counterparts, except for two differences. He then argues that neither difference I

directly. So he does spend much of his paper targeting (a), and for the purposes of evaluating his arguments, many of which only target (b) indirectly, it’s handy to take (a) to be the important part of ‘DeRose’s thesis’.

Barnett’s response (pp. ??-??) to the claim that I don’t accept (a) is then a bit puzzling: He writes the question is ‘ultimately moot’ because at the end of his paper there is a section that targets (c). (Though he doesn’t mention it at pp. 6-7, the last section also contains objections that are aimed directly at (b). We will briefly consider one of these at the end of this paper.) But since so much of his paper is aimed at (a), one would think it would be quite a disappointment to learn that nobody, at least so far as I know, actually holds (a).

posit *entails* a difference in literal content or in what idea gets expressed by the sentence, from which he concludes that, on my view, these two types of sentence have the same literal content and express the same ideas (pp. ??-??).

Based on my grasp of what Barnett means by ‘literal content’ and ‘idea’, I am strongly inclined to agree with him that the first difference in meaning I suggest—that “‘were’-ing up a conditional can serve the function of calling attention to the possibility that its antecedent is (or will be) false’ (CoD, p. 37)—would not make for a difference in literal content or in what idea gets expressed. But things are very different with the second, ‘more serious’ difference I posit, which is a somewhat involved difference in warranted assertability conditions between conditionals of the two types. I explain this difference at pp. 37-8 of CoD, but this *fairly* compact explanation makes use of terms (‘deliberationally useless’, ‘depending on backtracking grounds’) that have been explained elsewhere in CoD. Fortunately, we don’t have to spell this all out now, because the exact nature of this difference doesn’t seem to matter for the points that need to be made here.

Now, Barnett is certainly right that there can be a difference in warranted assertability conditions between two sentences without there being a difference in literal content. Where we’re dealing with sentences that do express propositions, this can happen when the sentences, despite a difference in warranted assertability conditions, have the same truth-conditions and express the same propositions. But we shouldn’t miss the fact that on other occasions, of course, a difference in warranted assertability conditions *is* due to a difference in literal content. If one of the sentences has a warranted assertability condition that the other lacks, that warranted assertability condition *can* also be a truth condition that the one sentence has but the other lacks. (It’s somewhat as if I had said that there has been ‘no change over the past year in what pets I have, except for the addition of one cat and one dog’, and Barnett took that to commit me to there having been no change in what collies I have, on the grounds that I said that a new cat and a new dog were the only changes, and having neither a new cat nor a new dog *entails* a change in collies. To which the response is: Um, still, that dog *could* be a collie, you know.)

But how can we discern those differences in warranted assertability conditions that are *also* differences in literal content from those that are not, when we are dealing with sentences that we are not presuming to even have truth conditions or to express propositions? (I assume that we do not want to use ‘literal content’ or ‘idea’ here in such a way that all such sentences that lack truth conditions and/or don’t express propositions count as having the same literal content and as expressing the same idea.)

I would suggest we seek tests in the vicinity of the ways that our conditionals behave (at least relative to some assumptions) as if they expressed propositions. As I wrote in CoD, ‘I do believe that even those who, like me, deny the propositionality of indicative conditionals have to make good sense of our pervasive talk about our asserting, accepting, believing, knowing, etc., these things’ (p. 39). So, if one can believe that, say,

(O) If Pete plays, he will win

without believing that

(Ow) If Pete were to play, he would win,

we could take that to indicate that (O) and (Ow) differ in their literal content.³ Or maybe the possibility of a difference in *rational* belief, or in *knowledge* would be better tests. We needn’t choose among the possible tests here: I think that straightforward FDCs and their ‘were’ed-up counterparts pass all these tests (as well as several other related tests) for having different literal content—and several of the cases I use in CoD show this. Take, for example, my first version of the Sly Pete story,⁴ at CoD p. 21 (see CoD for more details about the story, but many readers will be able to gather the essentials of the case from what I’m about to write). In this example, Sigmund has signalled to Pete which card Gus is holding, and has received Pete’s return signal confirming that Pete got the message and knows that Gus is holding card 83. Sigmund does not know which card Pete is holding, so he doesn’t know who is holding the higher card, but he does

³ If the sentence ‘If Pete plays, he will win’ did not express a proposition, then how could Sigmund believe that if Pete plays, he will win? My *suspicion*:

(BO) Sigmund believes that if Pete plays, he will win

attributes to Sigmund a *conditional belief*, involving *two* propositions—that Sigmund believes the proposition that Pete will win, where this is conditional on the proposition that Pete will play, and where Sigmund’s holding such a conditional belief is not a matter of there being any new, conditional proposition corresponding to (O) that Sigmund believes, but is irreducibly an attitude toward a pair of propositions. But this is to dip a toe into the very matters I am remaining neutral on in CoD. To wade in *just a bit* further, *perhaps*

(BOw) Sigmund believes that if Pete were to play, he would win

attributes to Sigmund a *different kind* of conditional belief (we might call it a ‘w-conditional’ belief) aimed at the *same* two propositions—saying that Sigmund believes that Pete will win, where this is *w*-conditional on the proposition that Pete will play, and where Sigmund’s holding such a *w*-conditional belief is not a matter of there being any new, conditional proposition corresponding to (Ow) that Sigmund believes, but is irreducibly an attitude toward a pair of propositions. On this possibility, the difference in meaning between (O) and (Ow) affects which variety of conditional belief is invoked by the verb ‘believes’ in attributions like the above.

⁴ The original Sly Pete story is of course due to Allan Gibbard (1981, pp. 226-9, 231-4). In CoD, I tell a few different versions of this story, with various modifications to the original made to suit it for the purposes for which I’m using it.

know that Pete knows what both cards are. And Sigmund knows that Pete is not stupid enough ‘play’ if he is holding the lower card; but Sigmund also knows that Pete will ‘play’ if he has the higher card, in which case he will win. Here, it seems clear that Sigmund believes that, and reasonably believes that, and knows that (O) if Pete plays, he will win. But it seems (though this is somewhat less clear) that Sigmund does *not* believe that, does *not* reasonably believe that, and does *not* know that (Ow) if Pete were to play, he would win. Indeed (to hint at another possible test), since the card that Sigmund knows Gus is holding is quite high (83, where there are 101 cards numbered 0-100), Sigmund should think it quite unlikely that Pete would win if he were to play, while finding it quite likely indeed that Pete will win if he plays.⁵

Since I’m not presuming that the FDCs here express propositions or have truth conditions, I can’t claim that straightforward and ‘were’ed-up FDCs express different propositions, or that they have different truth conditions. That they not only have different warranted assertability conditions, but also pass the above tests for having different literal contents is about the most one could hope for in the way of reasons for saying that they differ in literal content and in the ideas they express. And since I do accept those differences between them, I am inclined to deny that straightforward FDCs have the same literal content as, and that they express the same ideas as, their ‘were’ed-up counterparts.

That is by far the most important point I wish to make in response to Barnett’s discussion; I am very anxious to not be so thoroughly misunderstood in a matter so fundamental to my position on the meaning of these two types of conditionals. So, though there is much else in Barnett’s paper I would dispute, I thought it was worth allotting much of the space in this fairly brief reply to addressing this central matter of DeRose interpretation. Also, I hesitate a bit to address anything else, for fear of distracting attention from that most important point. But I should at least write something fairly quickly about the other major dispute between Barnett and myself, even though I won’t be in any position to settle this matter here, as I will do in the section 2, and then in section 3 I will fail to resist saying something quick in response to one very interesting criticism Barnett raises at the very end of his paper.

⁵ See CoD, pp. 32-3 for some related discussion, where I present Gibbard’s similar rulings concerning (O) and (Ow).

2. An underlying difference in approach

Barnett also argues against my key claim that both straightforward and ‘were’ed-up FDCs should be classified as indicative conditionals. This dispute cannot be settled quickly, for it crucially involves a fundamental underlying disagreement about the structure of conditionals, but I can *fairly* quickly indicate what that underlying disagreement is. Consider the (famous, due to E.W. Adams) subjunctive conditional:

(B) If Oswald hadn’t shot Kennedy, someone else would have.

If, say, a possible worlds analysis of such a claim asks us to consider the nearest possible worlds in which the antecedent of (B) is true, to discern if the consequent is also true in those worlds, philosophers know to ponder the nearest worlds in which

(~K) Oswald didn’t shoot Kennedy

is true, checking to see if

(L) Someone else did [shoot Kennedy]

is also true in those worlds. We take (~K) to be the antecedent of (B), and (L) to be (B)’s consequent. This is no innovation by me; it is the standard approach among philosophers. As I understand it, this approach is based on the idea that the right way to understand (B) is to view it as built from a ‘sentence frame’ like

If it had been the case that ... , it would have been the case that ... ,

where the blanks are filled in with (~K) and (L), respectively, yielding

(Br) If it had been the case that Oswald didn’t shoot Kennedy, it would have been the case that someone else did [shoot Kennedy]

as the correct ‘regimentation’ of (B). On this standard approach, the difference between (B) and the (also famous, due again to Adams) indicative conditional,

(A) If Oswald didn’t shoot Kennedy, someone else did,

is not that these two conditionals forge the very same kind of connection between two somewhat different antecedent/consequent pairs, but that they connect in somewhat different ways the same

antecedent/consequent pair. For (A)'s regimentation results from plugging that same antecedent/consequent pair—(\sim K) and (L)—into a somewhat different conditional sentence frame. On this approach, though the differences in words used between (A) and (B) occur in the embedded clauses, the function of these differences (if we say that (B) is a 'souped-up' version of (A), the function of the 'souping up') is not to change what things are being conditionally connected, as things may initially appear, but to change just what type of operator is being applied to the same pair of propositions.⁶

Looking at a *somewhat* parallel approach to different kinds of sentences *may* help here. Consider these two possibility statements:

(Pi) It is possible that the Cubs won today

(Ps) It is possible that the Cubs should have won today

I take (Pi) to express the epistemic possibility of

(Won) The Cubs won today

(Ps) on the surface looks like it expresses the same thing (that 'it is possible') about a different proposition, one expressed by

(SHW) The Cubs should have won today

And indeed, I think (Ps) *can* express the epistemic possibility of (SHW). (To employ terminology I've seen in a few places, when this happens, (SHW) is said to have or to keep its 'stand-alone meaning' when it occurs in (Ps). For a situation in which (Ps) has such meaning, imagine a discussion or an argument over whether the evaluative claim (SHW) is true. At some point, somebody rather sceptically concludes that at least they don't know that (SHW) is false, and expresses something in the vicinity of this conclusion of ignorance by means of (Ps).) But I think (Ps) can also express, and more often and more standardly does express, a different kind of possibility—some more 'metaphysical' possibility—of the proposition expressed by (Won), the same thing that is said to be possible by (Pi). When we compare (Ps), where it bears this second, more common reading, with (Pi), though the differences between the words in the possibility sentences occur after, and not in, the phrase 'It is possible that', the function of these differences

⁶ I am following Lewis 1973, pp. 2-3 and Edgington 1995, pp. 237-8 in my procedure here, and I am following Edgington in my terminology—in my use of 'sentence frame' and 'regimentation'. See also, for example, Bennett 2003, pp. 5-6, where Bennett calls this approach, with which he agrees, the 'common opinion'.

(if we speak of (Ps) as a being ‘souped-up’ version of (Pi), the function of the souping up) is not to change what is being said to be possible, but to change the kind of possibility that is being asserted of the same proposition, (Won). I can understand someone rejecting this approach, but, right or wrong (and *I think it’s right*), the approach certainly seems perfectly intelligible.

Barnett rejects that (\sim K) is the antecedent of (B), and with it, apparently, the standard approach to conditionals explained two paragraphs above, and this disagreement over approach plays a pivotal role in his argument against my case for my way of classifying conditionals. This is especially so in his claim that I misapply (Barnett’s rendition of) a test I offer for classifying a conditional as indicative (at pp. 11-2), but it also seems to me to underlie several important differences between us. But Barnett’s reasons for rejecting the standard approach, given in footnote (n. ?? at p. ???), seem to me meagre. He writes:

But if ‘antecedent’ means ‘clause embedded in the if-clause of a conditional sentence’, then this assumption [that (\sim K) is the antecedent of (B)] is patently false (which is not to say that it is not standard among philosophers). But what else can ‘antecedent’ mean?

Barnett then goes on to explore another possible meaning for ‘antecedent’ that I agree is unpromising, before closing the case on my assumption with a negative verdict. Note that we can give Barnett the term ‘antecedent’ (and ‘consequent’),⁷ and express the approach described above with different terms—perhaps the likes of ‘deep antecedent’? I could then phrase my classifying tests in those new terms. The question isn’t whether ‘antecedent’ and ‘consequent’ can be used in the way Barnett wants to, but whether they (or perhaps new alternatives to them) cannot be correctly used in the way those terms are employed in the approach two paragraphs above, a story given in the note (n. 18, CoD p. 19) where I explain my assumption about what the antecedents and consequents of the likes of (B) are, which is the note that Barnett discusses in rejecting that assumption. The meanings of these terms are given in the story: The ‘antecedent’ of a conditional is the clause that fills in the ‘if’-gap of the sentence frame from which that conditional is derived in order to yield the regimentation of that conditional. (Employing some standard sloppiness, we can slide conveniently between using ‘antecedent’ to refer that clause and to refer to the proposition that clause expresses.) One can of course be unhappy with that story in various ways, but I don’t see that Barnett here gives us any

⁷ I’m not sure exactly how Barnett would use these terms—for instance, how he would apply them to Lewis’s example, ‘No Hitler, no A-bomb’ (1973, p. 4). Would Barnett claim that the antecedent here is just ‘Hitler’, or in this case would he opt for something like ‘Hitler had not existed’?

independent reason to think that such an approach is wrong beyond his professed bafflement (which itself is baffling to me) at what ‘antecedent’ could possibly mean on this scheme.

Of course, none of this shows that the standard (at least among philosophers) approach to conditionals that I assume in CoD is right. Here I just wanted to make the point that Barnett has given no serious reason to think it is wrong, and to identify for readers a fundamental difference in approach to conditionals that underlies some points of disagreement between Barnett and myself. I am more than happy to have my arguments in CoD to rest on *this* bit of philosophical orthodoxy about conditionals.

3. Understanding FDCs without having certain concepts?

Finally, one of the Additional ‘Reasons to Reject DeRose’s Thesis’ in Barnett’s third and final section is quite interesting and, I think, worth commenting on. Barnett writes that a reason to reject my account is that

by adding the concept of deliberation to the meaning of [‘were’ed-up FDCs], the thesis places an implausible constraint on what is required to understand the sentences. (p. ???)

It is certainly true that the notion of deliberation plays a key role in my explanation of the meaning of ‘were’ed-up FDCs. But that’s not the half of it! The notion figures in my explanation of a key feature of the meaning of *straightforward* FDCs, too. (It figures in my account of the related but importantly different warranted assertability conditions for the two types of FDCs.) And if *that* worries Barnett, he might be truly scandalized by this: In my account of the meaning straightforward FDCs, I make use of a highly technical notion of a ‘deliberating agent’ with respect to a conditional.

To get a feeling for just how technical this notion is, but also for why it might be needed, consider this second (in this paper) version of the Sly Pete story. Project yourself into the role of Sigmund. You are helping Pete to cheat. You’ve seen Gus’s card, which is quite high: 83. But you haven’t signalled what Gus’s card is to Pete (as Sigmund did in the first version of the story that we considered in section 1). Rather, here in Story 2, you wrote a note saying that Gus’s card is 83, and slipped the note to Pete. Pete has not yet read the note, so he does not yet know who has the higher card, but it is certain that he will read the note before he has to decide whether to

‘play’, at which point he will know what both his own card is and what Gus’s card is. And, as in Story 1, it’s certain that Pete will not be stupid enough to play if his own card is lower than Gus’s, and it’s also certain that he will play if his own card is higher than Gus’s, in which case Pete will win. Now I take it, and trust you will also intuit, that based on these certainties, you are in a position to assure your boss, whom we assume is not and will not be in contact with Pete, that

(O) If Pete plays, he will win.

That’s old news. But note this interesting contrast. I take it, and trust you will also intuit, that if you get a chance to whisper something to him, you are *not* in a position to assert to the deliberating Pete this second-person version of (O):

(O-2) If you play, you will win

In case you think the explanation for this difference is that Pete is the agent referred to in the above straightforward FDCs while the boss is not, try adjusting the story a bit. Make Story 3 like Story 2, except that, as you realize, the boss is considering intervening in Pete’s game to help him to win by phoning him (Gus doesn’t object to Pete’s taking phone calls during the game!) with advice on whether or not to ‘play’. Now I take it that, and trust you will agree that, though the boss is still not referred to in (O), here in Story 3 you cannot assert (O) to the boss. (After all, you think it’s quite likely that Pete holds the lower card.)

Why not? Why can you assert (O) to the boss in Story 2, but you cannot assert (O-2) to Pete in Story 2, and you cannot assert (O) to the boss in the Story 3? If you think you can give a good explanation for this while using only terms that are tame enough that we can feel confident that all ordinary competent users of straightforward FDCs already have concepts for them, then all I can say to you is ‘Good luck!’ Yet it seems that competent users of these conditionals will know that they are in a position to assert (O) to the boss in circumstances like Sigmund’s situation in Story 2, but that they are not in a position to assert (O) to Pete in that situation, nor to assert (O) to the boss in circumstances like Sigmund’s situation in Story 3. Though it’s hard to say exactly what’s going on here in the vicinity of the paradox of analysis, it does seem that ordinary competent users of these sentences are *somehow* guided by *some* notions that they in some sense don’t have concepts for, at least to the extent that these are notions they don’t have any explicit awareness of, and yet can figure in good explanations for why the speakers will use the terms when they do, and won’t use them when they won’t.

At any rate (however the we solve this paradox of analysis), a key part of my explanation is that, in my technical sense, the boss is not a ‘deliberating agent’ with respect to (O) in Story 2, while Pete is a deliberating agent with respect to (O) (and (O-2)) in Story 2, and the boss is a deliberating agent with respect to (O) in Story 3. To get a feel for how technical this notion of a ‘deliberating agent’ with respect to a conditional is, note that we can imagine the boss to be deliberating over something, like which derivative bets to place on the results of Pete’s game, even in Story 2, and may even be using (O) in those deliberations. The sense in which he is *not* a deliberating agent with respect to (O) in Story 2 is my technical sense that he is not there involved in deciding whether to (try to) make the antecedent of (O) true in order to promote or resist the consequent being made true. (Recall that in Story 2 the boss is not in contact with Pete and will not intervene in Pete’s game.) Not so tame, I know. And that of course is only part of my explanation, and there are more technicalities involved in the other parts. See CoD for the account.

Here I wish to acknowledge that my account uses terms that ordinary competent users of the relevant conditionals have no explicit concepts for. But I also wish to express my conviction that that’s how it goes generally when giving interesting analyses of the meanings of philosophically interesting bits of ordinary language. I think it’s unreasonable to expect otherwise. What is reasonable to hope for is that, if ordinary speakers had the account explained to them and they came to understand it, it would seem to them a plausible account of why they use the term in question as they do. I would hope to be able to pass that reasonable test.

References

- Adams, E. W. 1975: *The Logic of Conditionals*. Dordrecht: D. Reidel.
- Barnett, David ????: ‘Future Conditionals and DeRose’s Thesis’. *Mind*, ???, pp. ???-???
- Bennett, Jonathan 2003: *A Philosophical Guide to Conditionals*. Oxford: Oxford University Press.
- DeRose, Keith 2010: ‘The Conditionals of Deliberation’. *Mind*, 119, pp. 1–42.
- Edgington, Dorothy 1995: ‘On Conditionals’. *Mind*, 104, pp. 235–329.
- Gibbard, Allan 1981: ‘Two recent theories of Conditionals’. In Harper, Stalnaker, and Pearce 1981, pp. 211–47.

Harper, W. L., R. Stalnaker, and G. Pearce (eds) 1981: *Ifs: Conditionals, Belief, Decision, Chance, and Time*. Dordrecht: D. Reidel.

Lewis, David 1973: *Counterfactuals*. Oxford: Basil Blackwell.