

The Conditionals of Deliberation

KEITH DEROSE

Yale University
keith.derose@yale.edu

Practical deliberation often involves conditional judgements about what will (likely) happen if certain alternatives are pursued. It is widely assumed that the conditionals useful in deliberation are counterfactual or subjunctive conditionals. Against this, I argue that the conditionals of deliberation are indicatives. Key to the argument is an account of the relation between 'straightforward' future-directed conditionals like 'If the house is not painted, it will soon look quite shabby' and 'were'ed-up' FDCs like 'If the house were not to be painted, it would soon look quite shabby': an account on which both of these types of FDCs are grouped with the indicatives for semantic treatment and on which, while conditionals of both types are properly used in means/ends deliberations, those of the 'were'ed-up variety are especially well suited for that purpose.

1. My thesis: the conditionals of deliberation are indicatives

Everyone knows that the conditionals of deliberation are counterfactuals, right? Here, for example, is a very typical statement, by Allan Gibbard and William L. Harper, as they set up a paper, before we are supposed to have reached any controversial territory:

We begin with a rough theory of rational decision-making. In the first place, rational decision-making involves conditional propositions: when a person weighs a major decision, it is rational for him to ask, for each act he considers, what would happen if he performed that act. It is rational, then, for him to consider propositions of the form 'If I were to do *a*, then *c* would happen'. Such a proposition we shall call a *counterfactual*. (Gibbard and Harper 1978, p. 153)

That is from more than 30 years ago, but the widespread assumption it expresses remains in force today: the conditionals of deliberation are counterfactuals.¹

¹ Given only the passage I quote, Gibbard and Harper can be read as simply stipulating that they will use 'counterfactual' to label *whatever* class of conditionals those of the form 'If I were to *a*, then *c* would happen' happen to fall in. In that case, since I agree conditionals of that form are conditionals of deliberation, I would agree that 'counterfactuals' (so understood) can be conditionals of deliberation. But it is clear that in calling these conditionals 'counterfactuals', Gibbard and Harper mean to be grouping them in with other conditionals that go by that name — like the paradigmatic counterfactual, (B), that we are about to encounter. I will

Going against all that, my subversive thesis is that the conditionals of deliberation are on the other side of the great divide between the types of conditionals: they are indicative conditionals.

By ‘conditionals of deliberation’, I mean conditionals that play the role in deliberation that Gibbard and Harper described above. We will look at some examples of conditionals playing that role in the following section.

But, first, the other matter that must be explained to understand my thesis is what is meant by ‘indicative conditionals’. For our purposes (as for many philosophical purposes), the best way to use the classificatory scheme of ‘indicative’ conditionals, on the one hand, versus ‘subjunctive’ or ‘counterfactual’ conditionals, on the other, is by reference to suitable paradigm examples of each. For that purpose, our paradigms will be E. W. Adams’s famous pair:

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

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

We can all sense the marked difference in meaning between these two. (Which is shown by the fact that those who think that Oswald was acting alone, with no back-up, will typically accept (A) but reject (B).) (A) will serve as our paradigm ‘indicative’ conditional; (B) as our paradigm ‘subjunctive.’ To call a conditional ‘indicative’ is to say that its meaning is such that it should be grouped with (A) for semantic treatment. To call a conditional ‘subjunctive’ or ‘counterfactual’ is to say its meaning is such as to be grouped with (B). Though (A) and (B) are our official paradigms, other conditionals are clearly similar enough to these paradigms in the relevant respects to be *clearly* subjunctive or indicative conditionals in our sense; I will say that such a conditional is a ‘paradigmatic’ subjunctive or indicative conditional, as the case may be.

One might naturally think that the way to test my thesis would be to look at the meaning of indicative conditionals, and then determine whether conditionals with such a meaning are or could be useful in the requisite way in deliberation. However, as wise students of indicative conditionals will tell you, we are in no position to follow that course, for the matter of the meaning of indicative conditionals is about as controversial as philosophical issues get.²

argue that they are wrong to think that it is conditionals semantically like (B) that are the conditionals of deliberation.

² This will be discussed in section 13.

So we will instead follow the more promising strategy of looking at particular examples. Some conditionals are quite clearly useful in deliberation in the way in question, and we do not need a settled view on their semantics to discern that. And while all theories as to the truth conditions, or lack thereof, of indicative conditionals are extremely controversial, paradigmatic indicative conditionals display certain types of ‘behaviour’ closely connected with their meaning (and which behaviour forms crucial data which theories about their meaning try to explain), and by seeing that conditionals useful in deliberation display the same characteristic ‘behaviour’, we can build a good case that they should be grouped with the indicatives for semantic treatment.

2. Examples of conditionals in deliberation, divine and human

As promised, we now look at some examples of conditionals at work in contexts of deliberation. We start with an example from philosophy of religion — in fact, from one of the hottest topics in current analytic philosophy of religion, the issue of whether God possesses ‘middle knowledge’.

So: suppose you are God, and you want to create a primo world. After creating lots of great stuff, you discern in your infinite wisdom that what your world really needs, to top it all off, is a free creature performing just one good action with libertarian freedom. (Of course, it is more realistic to suppose that if one libertarian-free action is good, then a really primo world would have many such actions, but we will keep the issue simple by supposing that you desire only one such good action.) So you decide to create such a creature — let us call her Eve — and put her in one of the exactly right possible situations (we will call this situation S_1) where she is free to choose between two courses of action, one good and one bad, and where her performing the good action is exactly what is needed for the perfect completion of your world. If she performs the bad action, however, that ruins everything: you could have had a better world by not having any free creatures at all than you get with Eve messing up. Since Eve must exercise libertarian freedom for you to get the primo world you want, you cannot cause her to do the right action, nor can you set off a series of causes that will causally determine her to do the right thing. What’s a god to do? Is there any way of getting the primo world

you desire without taking any real risk that everything will get messed up?

Perhaps you can use your knowledge, your Divine omniscience, to avoid any real risk. It is fairly widely agreed that you cannot here utilize any *simple foreknowledge* you might have of such non-conditional propositions as *Eve will sin (in S₁)*. For suppose that Eve in fact will sin in S₁, and you foreknow this. The supposition that you use this foreknowledge to avoid the trouble your world is headed for is itself headed for trouble. For if you then decide to put Eve in some other situation, say S₂, where she may fare better, or to put some other, more favourably disposed, possible free creature into S₁, or if you decide to skip the whole free creature idea and make do with a pretty good, though not primo, completely deterministic world, then, though it looks as if you have thereby avoided the trouble, it also looks like you did not after all know that Eve would sin, since it turns out not to be true that Eve sins, and you cannot have known what was not true.

So, as it seems to most who study the issue, the knowledge that God at least arguably might have that could be used to avoid any real dice-throwing risks in an indeterministic world is not simple foreknowledge of such propositions as *Eve will sin in S₁*, but is rather ‘middle knowledge’ of certain *conditionals*. But which conditionals? Ignoring here how the participants in the middle knowledge debate would themselves respond,³ here is the natural answer to

³ Strangely, the ur-examples put forward as potential objects of God’s middle knowledge are Alvin Plantinga’s

If Curley had been offered \$20,000, he would have accepted the bribe. (Plantinga 1974, p. 174) and Robert Adams’s

If President Kennedy had not been shot, he would have bombed North Viet Nam. (Adams 1977, p. 109)

These seem strange examples because middle knowledge is clearly assumed also to be knowledge that would help God to exercise a kind of providential control over an indeterministic world which we will here just cryptically call ‘Molinistic control’, yet these conditionals are past-directed sentences that would be in place only after it was too late to control the relevant events: they appear to be useless examples of ‘Monday morning quarterbacking’, as E. W. Adams aptly pointed out about similar past-directed conditionals, back in the days when quarterbacks at least sometimes called their own plays on Sundays (Adams 1975, p. 133). Presumably, those involved in the middle knowledge debate thought that these ur-examples were somehow past tense versions of FDCs (perhaps the likes of our (C) and (D), but more likely conditionals like (Cw), all of which we will encounter a bit later in this article) such that the past-directed conditionals are true (and known to God) after the time the antecedent would have occurred if and only if the corresponding FDCs are true (and known to God) before the time of the antecedent. But since it is highly uncertain that any such relation really

this question: what would really help you in your Divine predicament is knowledge of something like

(C) If I put Eve into situation S_1 , she will sin

or, less personally

(D) If Eve is put into situation S_1 , she will sin

Suppose you know those conditionals. Then you will know not to put Eve into S_1 , and the supposition that you so use this ‘middle knowledge’ to avoid trouble does not itself lead to the trouble that we hit when we assumed you used simple foreknowledge to the same effect. For if there is some other situation S_2 that is such that you foresee that Eve will do the right thing if she is put into S_2 , and you therefore put her into S_2 rather than S_1 (or if you create some other possible free creature, or none at all), and you thereby avoid trouble, we can still consistently claim that you knew (C) and (D). If, on the other hand, what you know (with your usual divine certainty) is the happier complement⁴ of (C),

(Cc) If I put Eve into S_1 , she will not sin

holds, the question of whether God could have middle knowledge splits into two separate questions that those involved in the debates have wrongly thought to amount to the same thing:

- (1) Are the past directed counterfactuals that are put forward as the potential objects of middle knowledge true and known to God?
- (2) Are the conditionals the knowledge of which would enable God to exercise ‘Molinistic control’ over an indeterminist world true and known to God?

The distinction between these two questions is particularly noticeable to me, since I give them different answers: 1: Yes, those past-directed counterfactuals are often true and known to God — though this gets a bit complicated, due to the fact that the meaning of these counterfactuals is context-dependent and their meaning in context is not always such that either a ‘counter-factual of freedom’ or its complement is always true. Still, it is often the case that one of these counterfactuals of freedom expresses a proposition that is true, and God can and does know that that proposition is true. 2: But no, the conditionals the knowledge of which would enable God to exercise ‘Molinistic control’ (FDCs of the types we will be discussing in this article) are not true and could not be known by God. Veterans of (both sides of) the Middle Knowledge Wars inform me that, given those answers, I am best classified as being against middle knowledge. That seems right to me, since question 2, which I answer negatively, seems the one most important to the middle knowledge debate.

⁴ I will call (A) and (Ac), pairs of conditionals sharing the same antecedent but having opposite consequents — conditionals of the forms $A \rightarrow C$ and $A \rightarrow \neg C$ — ‘complements’ of one another. In so doing, I am not assuming that they are contradictories of one another — that exactly one of them must be true. Nor am I even assuming that they are inconsistent — that at most one of them can be true. (Arguably, in some cases of indicative conditionals, both members of a pair of ‘complements’ can be true.)

and (D)'s complement, then you know that you are free and clear to put Eve into S_1 , without worrying that she will mess everything up. Of course, in situations where the agent acts with libertarian freedom, it is very controversial whether even omniscient God can have knowledge of the relevant conditionals—indeed, that just is the hot controversy over whether God has 'middle knowledge'.⁵ But at least these seem to be the conditionals that it *would* be helpful for you, as God, to know.

Like God, we lesser agents also use conditionals in deciding which courses of action to pursue. Indeed, it often is useful for us to know conditionals about what people will freely do if we do something:

(E) If I offer Eve £2,000 for her car, she will accept

Of course, not being God, Divinely Certain knowledge is certainly not in the cards for us (whether or not God might have it). Maybe even knowledge simpliciter of matters such as this is unattainable. Still, such a conditional seems like the kind of thing it would be helpful to know, and, failing that, to have beliefs about that are very likely to be correct. And such beliefs, whether or not they amount to knowledge, seem to actually guide our action: the reason (or at least part of the reason) why I might offer Eve £2,000 for her car may be that I believe that, or believe it is probable that, if I offer her that amount, she will accept. And, of course, beliefs about what others will freely do in various situations form only one kind—and perhaps a particularly problematic kind—of the conditional beliefs that so guide our action. In the relevant situations, it is helpful to know, or have probable beliefs, about the following:

(F) If I try to drive to work without first filling up the tank, I will run out of petrol

(G) If I start walking to my meeting only 5 minutes before it starts, I will be late

(H) If the house is not painted, it will soon look quite shabby

These all seem to be conditionals that would be useful in deliberation: to the extent that I have reason to believe one of them, then in so far as I desire its consequent to be true, I have reason to make (or to try to make, in cases where the truth of the antecedent is not completely up to me) its antecedent true. And to the extent I believe

⁵ Or at least, this is *half* (but the more important half) of the 'middle knowledge' controversy: See note 3.

one of these conditionals and want its consequent not to be true, I have reason to try to avoid the truth of its antecedent. Hence, all these conditionals seem to be conditionals of deliberation, playing the role in deliberation that is described in the Gibbard and Harper quotation at the very opening of this article.

What is more, all of (C)–(H) at least appear to be indicative conditionals, as I am informed by many who think they can tell what camp a conditional falls in just by quickly looking at its quasi-grammatical features. More importantly, there are strong reasons for thinking this appearance is correct, and that the meaning of (C)–(H) is such that they should be grouped with (A), as we will see below in sections 5 and 6. So it looks like we will soon have good reason to think that at least some conditionals of deliberation are indicatives.

3. Straightforward versus ‘were’ed-up future-directed conditionals

But wait! Though all of (C)–(H) are quite naturally used in deliberation in the way described by Gibbard and Harper, they are not of the form Gibbard and Harper specified as the conditionals that play that role. Recall that we were told that the conditionals rational for us to consider in deliberation are those of the form, ‘If I were to do *a*, then *c* would happen.’ Thus, Gibbard and Harper (as well as many others) would probably counsel God to consider this ‘were’ed-up version of (C), rather than (C) itself, in deliberation:

(Cw) If I were to put Eve into situation S_1 , she would sin

And the ‘conditionals of deliberation’ we humans should consider in deliberation would be identified not as the likes of (E)–(H) themselves, but rather their ‘were’ed-up counterparts:

(Ew) If I were to offer Eve £2,000 for her car, she would accept

(Fw) If I were to try to drive to work without first filling up the tank, I would run out of petrol

(Gw) If I were to start walking to my meeting only 5 minutes before it starts, I would be late

(Hw) If the house were not to be painted, it would soon look quite shabby

These ‘were’ed-up conditionals, the darlings of various decision theorists, are also evidently conditionals of deliberation: they can

play the relevant role in deliberation. And what with all the ‘were’s and ‘would’s inserted in them here and there, they can seem somehow more subjunctive than what we will call their ‘straightforward’ cousins, the likes of (C)–(H), and are likely to be classified as counterfactuals by those who think they can tell easily just by looking.

But, in part precisely because they are future directed, neither of these types of future-directed conditionals (henceforth FDCs) — neither the ‘straightforward’ nor the “‘were’ed-up’ ones — are very similar to either of the paradigms of our two ‘camps’ of conditionals. The paradigmatic indicatives and subjunctives are past directed. Since none of our FDCs is paradigmatically in either camp, it will take some investigation to decide how to group them. And, despite the names given to the two camps — ‘indicative’ versus ‘subjunctive’, as the second camp is often titled — the question of how to classify our various FDCs is not ultimately a question about the ‘moods’ of the verbs they contain, but about whether their meanings are such that they should be grouped with the paradigmatic indicatives (like (A)) or with the paradigmatic subjunctives (like (B)) — or perhaps whether they belong in neither of these two camps.

4. A preliminary look at the relation between straightforward and ‘were’ed-up FDCs

It is of course possible that the two types of FDCs we are dealing with belong in different camps. Indeed, as I have noted, some who think they understand what ‘moods’ of verbs amount to, and think that these moods are a good indicator of the semantic ‘camp’ individual conditionals occupy, might quickly classify our straightforward conditionals as indicatives and the ‘were’ed-up conditionals as subjunctives. Since conditionals of both types are ‘conditionals of deliberation’ as we are using that phrase, this would mean that conditionals of deliberation can be either indicatives or subjunctives.

However, the relation between one of these straightforward FDCs and the analogous ‘were’ed-up FDC at least often does not *seem* to be much like the relation between (A) and (B). The difference between (A) and (B), as we noted, is quite sharp. In contrast, when we compare, for instance, these two conditionals that we have already considered:

- (C) If I put Eve into situation S_1 , she will sin

(Cw) If I were to put Eve into situation S₁, she would sin

there seems to be nothing like the sharp contrast we sense between (A) and (B). In fact, in at least many contexts, including many where the speaker is deliberating about whether to put Eve into situation S₁, these two, so far from being sharply different, can at least seem to be something like equivalent. It can seem decidedly odd to conjoin an assertion of either with the assertion of the other's complement:

(C+Cwc) If I put Eve into situation S₁, she will sin; but if I were to put her into situation S₁, she would not sin

and

(Cw+Cc) If I were to put Eve into situation S₁, she would sin; but if I put her into situation S₁, she won't sin

both sound extremely awkward. Indeed, they produce something of the feeling of a contradiction. And it is even odd to combine an assertion of either of these conditionals with a question about the acceptability of the other. It is hard to make sense of either of the following, except perhaps to understand the speaker as, in the second half of each, thinking twice about and throwing into question what she has just asserted in the first half:

(C+Cw?) If I put Eve into situation S₁, she will sin. But would she sin if I were to put her into situation S₁?

(Cw+C?) If I were to put Eve into situation S₁, she would sin. But will she sin if I put her into situation S₁?

When one considers some straightforward FDC and the corresponding 'were'ed-up FDC, the preliminary hypothesis that can spring to mind as to their relation is that they mean the same thing, but for the fact that the 'were'ed-up version also somehow signals that its antecedent is improbable, where the type of signalling in question is such that the conditional is not rendered false or wrong if the antecedent is not actually improbable. This suggests itself because many of the situations where one would use the 'were'ed-up FDC are ones where one believes the antecedent is improbable and is looking to emphasize that fact. (Of course, we often use straightforward FDCs even when their antecedents are improbable if we do not want to be calling special attention to that fact.) But when one more carefully considers the range of situations in which one might opt for the

‘were’ed-up version, I think one will be led to postulate a somewhat more general hypothesis, and say instead that the function of ‘were’-ing an FDC up is to call attention to the possibility that the antecedent is (or will be) false, where one reason one might have for calling attention to the possibility that the antecedent is (or will be) false is that it is quite likely that it is (or will be) false.

But so far this is all just a preliminary look. As we will see, the preliminary hypothesis mentioned above does not capture all the differences between straightforward and ‘were’ed-up FDCs. I have noted here that straightforward FDCs often seem to be equivalent or at least very close in meaning to their ‘were’ed-up counterparts.⁶ However, as we will see (and as others have already shown), there are other situations in which a straightforward FDC and its ‘were’ed-up counterpart come apart enough that the former seems correct while the other seems wrong.⁷ Thus, despite how similar the two types of FDCs often seem to be, and despite the fact that, as I will argue, both types of conditionals should be grouped with the indicatives, the differences in meaning between a straightforward FDC and its ‘were’ed-up analogue, can importantly go beyond what is allowed for in the above paragraph.

What I do want to take from our preliminary look is that there is some close relation between straightforward FDCs and their ‘were’ed-up counterparts. This close relation at least often seems to be quite different from the sharp contrast between the likes of (A) and (B). A main desideratum of an account of the relation between these two types of FDCs is that it somehow make sense of this sense of a very close relation between them—that it makes sense of the fact that ‘were’ed-up FDCs often seem so close in meaning to their straightforward counterparts, while also making sense of where and how the two differ, in the situations where they do come apart. In section 12 of this article, I will propose a hypothesis as to the relation between these two types of FDCs that does just that.

⁶ William Lycan has claimed that there do not seem to be ‘Adams pairs’ of FDCs; see Lycan 2001, pp. 162–6, for discussion. The examples I mention in the next note and discuss later in the article seem to be counter-examples to Lycan’s claim, though Lycan indicates awareness that there might be such counter-examples in cases involving backtracking at Lycan 2001, p. 163, n. 18. Lycan’s claim is testimony to how close the straightforward and ‘were’ed-up FDCs can appear to be when one is considering examples that occur in contexts of deliberation and/or do not involve backtracking (to anticipate terms that will be explained below).

⁷ Gibbard (1981) presents an example that we will discuss (in a somewhat modified form) at the start of section 11 in which a straightforward FDC (in our modified example, this will be (O)) seems correct while its ‘were’ed-up counterpart (Ow) seems wrong. Dorothy Edgington (1995) presents a good example that I will present below in note 30 and discuss in note 32.

5. Straightforward FDCs are indicatives: assertability conditions

Our first item of business, though, is to argue that straightforward FDCs are indicatives — that semantically, they belong with paradigmatic indicatives like (A).⁸ This will be done in this and in the following section.

⁸ Here I find myself sharply at odds with another rebel about conditionals, Vic Dudman, who contends that straightforward FDCs belong with (B). Dudman argues:

The amalgam of 'indicative' doctrine thus far is untenable. For

(DA) Someone shot Kennedy. Therefore if Oswald did not shoot Kennedy someone else did

encodes a valid argument, while

(DB) Someone will shoot Kennedy. Therefore if Oswald does not shoot Kennedy someone else will

does not. (Dudman 1994, p. 21; changing Dudman's '(A)' and '(B)' to '(DA)' and '(DB)' to avoid confusion with my own (A) and (B))

But, at least intuitively, (DA) and (DB) strike me the same way with regards to the issue of validity: both strike me as *intuitively* valid, though this is far from an overpowering intuition in either case. There are concerns that hinder me from following intuition in both cases (including one Dudman raises that we are about to look at), but the concerns are roughly equal in the two cases. Dudman continues (immediately after what is quoted above):

Persuaded that some unidentified illwisher, acting alone, is poised to shoot Kennedy, and with excellent prospects of success, X may well assent that someone will shoot Kennedy, but she will dissent from the future conditional inferred in (DB) if she thinks the lone malefactor could be Oswald, unless she already knows of some other plot to shoot Kennedy. (DA), on the other hand, encodes a straight out deduction. (pp. 21–2)

But what Dudman says here against the validity of (DB) (all but the last sentence of the above quotation) can be mimicked with as much success against the validity of (DA), when we imagine a speaker in the relevantly analogous evidential situation (her very strong, but only, reason for thinking that someone shot Kennedy is her knowledge of a plot by someone or other that should already be completed, and the fact that she has not heard one way or the other that the deed has been done is not evidence that it has not been), or so it seems to me:

Persuaded that some unidentified illwisher, acting alone, was poised to shoot Kennedy, and with excellent enough prospects of having succeeded by now, X may well assent that someone shot Kennedy, but she will dissent from the past conditional inferred in (DA) if she thinks the lone malefactor could be Oswald, unless she knows of some other plot to have shot Kennedy.

And Dudman's bald assertion of (DA)'s validity in the last sentence of the passage quoted above adds nothing to the case, since the same assertion, with (at least to my thinking) about the same degree of plausibility (not all that much, in either case, given the consideration Dudman raises against (DA) and the parallel consideration that can be raised against (DB)), could just about as well be baldly declared concerning (DB). The considerations for and against the validity of (DA) and (DB) Dudman puts in play here run so parallel to one another that his argument not only fails to convince me, but also it actually serves to reinforce my confidence in my opposite verdict that the conditionals that are the conclusions of (DA) and (DB) should be grouped together. For related critical discussion of Dudman, see Cross 2002.

Paradigmatic indicatives like (A) have certain remarkable assertability conditions. A powerful reason for thinking straightforward FDCs are indicatives is that they share these assertability conditions with the likes of (A). The problem is that there are different formulations, all roughly in the same ballpark, about what these assertability conditions for indicatives are. This will complicate our discussion a bit, but I hope that whatever formulation of the assertability conditions of the likes of (A) one prefers, one will be able to see that straightforward FDCs have the same assertability conditions as do the paradigmatic indicatives.

Frank Jackson formulates an account as follows:

The assertibility of an indicative conditional is the conditional probability of its consequent given its antecedent. (Jackson 1987, p. 11)⁹

In quickly citing supporting evidence for this account, a page later Jackson writes:

Or take a conditional with 0.5 assertibility, say, 'If I toss this fair coin, it will land heads'; the probability of the coin landing heads given it is tossed is 0.5 also. (Jackson 1987, p. 12)

Jackson does not argue that the conditional in the above quotation is assertable to degree 0.5. That is just an observation that Jackson makes. I find this extremely puzzling. To the extent that I can just intuit the degree to which the conditional is assertable, I would give it a value *much* lower than 0.5. (Forced to assign a number, I would go for something like 0.06.) After all, it is a fair coin. So I have *no idea* which of its two sides it will land on if I toss it. I would have to say that I am in no position to assert either that it will land heads if I toss it, or that it will land tails if I toss it. And it does not seem a close call: neither conditional seems close to being half-way assertable. It is tempting to say that I am in no position *at all* to assert either conditional, which might tempt one to give them both a flat 0 on the

⁹ Note that Jackson uses 'assertibility' — a slight modification on the usual spelling of 'assertability' — as a semi-technical term to denote those aspects of assertability that have to do with whether an assertion is 'justified or warranted' — in the epistemological sense, not in a purely pragmatic one' (Jackson 1987, p. 8). Thus, an assertion is assertible in this semi-technical sense where one is in a good enough epistemic position with respect to what one is asserting to be able to assert it. Thus, if I am in a library where I am not supposed to talk at all, 'I am in a library' is not assertable for me (because I am not allowed to assert anything by the library's rules), but it is assertible (supposing I am in a good enough epistemic position with respect to my location to assert that). While I do not employ Jackson's device to mark it, my own use of the normally spelled 'assertable' and 'assertability' is like Jackson's use of 'assertible'/'assertibility'.

assertability scale. But then, I suppose that when I compare Jackson's conditional with 'If I toss this fair die, it will land 6', the latter seems even less assertable, suggesting the former should not just be given a 0. Still, 0.5 seems way too high. Indeed, I suspect the only way someone would reach the conclusion that Jackson's conditional has an assertability of 0.5 is if one were already assuming that its assertability was equal to the relevant conditional probability, which we know to be 0.5. (But in that case of course one should not be seeking to support Jackson's hypothesis by citing that assertability value as an independent observation that matches the theory's prediction.)

So I do not have much sympathy for Jackson's hypothesis. Still, if one is inclined to think that the assertability of paradigmatic indicatives like (A) are equal to the conditional probability of their consequents on their antecedents, then, hopefully, one will also think that the assertability of a straightforward FDC is equal to the conditional probability of its consequent, given its antecedent. And, indeed, Jackson himself thinks so: the example he uses in the above quotation is a straightforward FDC, which he takes to be in the indicative camp, and which he does explicitly say fits his hypothesis.

David Lewis has a closely related, but different and superior, account. Lewis claims that the assertability of an indicative conditional 'goes ... by the conditional subjective probability of the consequent, given the antecedent' (Lewis 1976, p. 297). Note that this motto could be adopted by Jackson as well; on both theories, the assertability of an indicative conditional 'goes by' the relevant conditional probability. But Lewis posits a different, more plausible, connection. He does not claim that the degree to which the conditional is assertable is equal to the conditional probability of its consequent on its antecedent. Rather, according to Lewis, an indicative conditional is assertable if the conditional probability of its consequent on its antecedent is very high — sufficiently close to 1.¹⁰ Presumably, 0.5 is not sufficiently close to 1.

¹⁰ Lewis writes:

The truthful speaker wants not to assert falsehoods, wherefore he is willing to assert only what he takes to be very probably true. He deems it permissible to assert that A only if $P(A)$ is sufficiently close to 1, where P is the probability function that represents his system of belief at the time. Assertability goes by subjective probability.

At least, it does in most cases. But Ernest Adams has pointed out an apparent exception. In the case of ordinary indicative conditionals, it seems that assertability goes instead by the conditional subjective probability of the consequent, given the antecedent. (Lewis 1976, p. 297)

The best way to interpret Lewis here is as holding the thesis I ascribe to him. He ends the first paragraph of the above quotation by writing that in most cases, 'Assertability goes by

Lewis's hypothesis is quite plausible, and works for most examples. Note that his hypothesis seems to apply plausibly to our paradigm indicative, (A), but not at all plausibly to our paradigm subjunctive, (B). I trust that those who accept Lewis's account, and hold that, say, (A), is assertable when the conditional probability of *Someone else shot Kennedy* given *Oswald did not shoot Kennedy* is sufficiently close to 1, will also find that straightforward FDCs are assertable when the conditional probability of their consequents on their antecedents is sufficiently close to 1.

But while Lewis's hypothesis seems close to right, and gets most cases right, I think it gets some cases wrong.¹¹ My own favoured account of the assertability conditions of indicative conditionals is a version of the Ramsey test (Ramsey 1931). I will start with a standard description of the Ramsey test, by William Lycan, who is using '>' as his sign for a conditional connective:

To evaluate $A > C$, add A hypothetically to your current belief set, make such revisions in your new total belief set as would be rationally required to preserve coherence while retaining A, and see whether C would be a member of the revised set. If it would, the conditional may be asserted; if not, not. (Lycan 2001, p. 48)

While this seems on the right basic track as an account of the assertability conditions of indicative conditionals, it seems too permissive. Much depends here on what it is for a proposition to be in one's 'belief set'. But on my understanding of that, there seem to be many simple, non-conditional propositions that are in my belief set that I am in no position to assert, because, though I do believe them, I am just not well-enough positioned with respect to them to be in a

subjective probability', where this summarizes the observation that propositions are assertable when their probability is sufficiently close to 1. Thus, when in the second paragraph he writes that in cases of indicative conditionals, assertability 'goes by' conditional probability, it seems natural to give a similar 'sufficiently close to 1' reading of 'goes by'.

Both Jackson and Lewis take themselves to be following E. W. Adams in their hypotheses about the assertability conditions of indicative conditionals. But if I am right that their accounts are significantly different from one another's, then, unless Adams gives two very different accounts, they are not both following Adams exactly.

¹¹ I am thinking primarily of lottery cases here. Note that lottery cases also provide plausible counter-examples to the simple high-probability account of non-conditional assertions: no matter how many tickets there are in a standard lottery situation, and so no matter how close to 1 the probability of 'I lost' is (suppose the drawing has already taken place, but the speaker has not heard the results of the drawing), 'I lost' still seems unassertable. Similarly, no matter how close to 1 is the conditional probability of 'If the drawing has been held, I lost', the speaker seems unable to properly assert that conditional.

position to flat-out assert them. If that is right, then this standard Ramsey test account of the assertability conditions of indicative conditionals seems too weak. Suppose that the result of adding A hypothetically to my belief set would result in C becoming part of my revised belief set alright, but only as one of the members of that set that I would not be in a position to assert. Then it seems that I am not in a position to assert the conditional $A \rightarrow C$.

If we knew what are the conditions of the assertability of regular, non-conditional propositions, that would guide us in working out a Ramsey test account of the assertability of indicative conditionals. The account Lycan articulates above—which can be called a ‘conditional belief set account’—seems plausible if, but only if, this ‘simple belief set’ account holds of regular, non-conditional assertions: you are in a position to assert that P if and only if P is in your belief set. If instead, like Lewis, one accepts a probability account of simple assertion, on which one is positioned to assert that P if and only if P’s probability for you is sufficiently close to 1, then you will want to apply a Ramsey test, not by asking whether C would become part of one’s belief set when one adds A to that set, but whether C’s probability would then become sufficiently close to 1. Against both of those accounts of simple assertability, I am a committed advocate of the knowledge account of assertion, on which one is positioned to assert what one knows.¹² This suggests a version of the Ramsey test on which we ask whether adding A as a certainty to one’s belief set would put one in a position to *know* that C.

But for our current purposes (and for many other purposes as well), we can bypass all this uncertainty about general assertability by simply accepting a ‘conditional assertability’ account of indicative conditionals, on which one is positioned to assert $A \rightarrow C$ if and only if adding A as a certainty to one’s belief set would put one in a position to assert that C. If it would, then $A \rightarrow C$ is assertable for one; if not, not. We then leave open the further matter of what it takes generally to be in a position to assert some non-conditional proposition. That is the version of the Ramsey test account that I here advocate. It seems to correctly articulate the assertability conditions of paradigmatic indicative conditionals, like (A), but not of subjunctives, like (B). And it seems just as plausible when applied to straightforward FDCs as it is when applied to paradigmatic (past-directed) indicatives—except for a strange but important sort of counter-example we will encounter

¹² For discussion and defence, see DeRose 2002, especially Sects 2.1 and 2.2, pp. 179–83.

and discuss in section 10. Ignoring those counter-examples for now, this impressive match in the assertability conditions of straightforward FDCs and paradigmatic indicative conditionals gives us good reason to think straightforward FDCs are indicatives.

6. Straightforward FDCs are indicatives: the paradox of indicative conditionals

‘Indicative’ conditionals like (A) display another remarkable property: they are subject to what Frank Jackson has dubbed the ‘Paradox of Indicative Conditionals’.¹³ While it is widely recognized that indicatives like (A) have this property, I am not aware of anyone using the presence of this property a classifying device, but it seems a good device, and a nice complement to the test that we used in the previous section.¹⁴ There we used the conditions under which the sentences in question seem assertable. Another genus of semantic classificatory guides is what inferences involving a sentence are—or at least seem to be—valid. Our new test is of this second variety.

Before Jackson gave it its name, the Paradox of Indicative Conditionals was nicely set up by Robert Stalnaker (1975), using as his example the paradigmatically indicative conditional,

($\neg I \rightarrow J$) If the Butler didn’t do it, the gardener did

The Paradox consists in two apparent facts about ($\neg I \rightarrow J$); it is a *remarkable* paradox in that these apparent facts are quite simple, and the intuitions that they are indeed facts are each intuitively quite powerful, yet the intuitions cannot both be correct. First, ($\neg I \rightarrow J$) seems to be entailed by the disjunction

($I \vee J$) Either the butler did it, or the gardener did it

If someone were to reason

($I \vee J \therefore \neg I \rightarrow J$) Either the butler did it or the gardener did it;
therefore, if the Butler didn’t do it, the gardener did

¹³ See Jackson 1987, pp. 4–8. I am at least unaware of anyone using this terminology before Jackson.

¹⁴ I have some qualms about this test, which I will explain in note 15, below. My problems involve indicatives that do not clearly pass this test for being an indicative. But while a conditional’s not clearly passing this test is not a secure sign that it is not an indicative, a conditional’s passing this test still seems to be strong grounds for thinking it is an indicative, and that is what I am relying on here.

they would certainly seem to be performing a perfectly valid inference.¹⁵ However, the strong intuition that $(I \vee J \therefore \neg I \rightarrow J)$ is valid clashes with a second strong intuition, namely, that $(\neg I \rightarrow J)$ is *not* entailed by the opposite of its antecedent

(I) The butler did it

The reasoning

(I $\therefore \neg I \rightarrow J$) The butler did it; therefore, if the Butler didn't do it, the gardener did

so far from being valid, appears to be just crazy. (Only a philosopher, dazed by over-exposure to \supset s, would actually reason in that way.) But at least one of these strong intuitions — that $(I \vee J \therefore \neg I \rightarrow J)$ is valid or that $(I \therefore \neg I \rightarrow J)$ is invalid — must be wrong. Given that (I) entails $(I \vee J)$, and given the transitivity of entailment, it just cannot be that $(\neg I \rightarrow J)$ is entailed by the 'weaker' $(I \vee J)$ but fails to be entailed by the 'stronger' (I).

This suggests a test: if a conditional, $\neg A \rightarrow C$, has the remarkable property of being subject to the 'Paradox of Indicative Conditionals' — that is, if it seems to be entailed by $A \vee C$ but also seems not to be entailed by A — then it should be classified with the indicatives.¹⁶ Note that not only are we using highly suspect intuitions in applying this test, and note also that we are not in any way relying on our intuitions being correct. Indeed, whenever a conditional does elicit the two intuitions that by the current test indicate it should be classified with the indicatives, we know that at least one of those

¹⁵ For some reason, this clear appearance of validity seems (at least to me, at least often) to largely vanish when the first disjunct of the premiss contains a negative and the antecedent of the conclusion is positive. In contrast with $(I \vee J \therefore \neg I \rightarrow J)$, which clearly seems valid,

$(\neg I \vee J \therefore I \rightarrow J)$ Either the butler did not do it or the gardener did it; therefore, if the Butler did it, the gardener did it

does not produce a clear intuitive appearance of validity, at least to me. And this does not seem to be due to a feature of this particular example (like that we tend to assume that only one of these people could have 'done it'). Often, at least, $(\neg A \vee C \therefore A \rightarrow C)$, with the negative in the premiss seems (at least to me) to fail to produce the clear appearance of validity that $(A \vee C \therefore \neg A \rightarrow C)$, with the negative in the conclusion, produces. Perhaps the way to apply this test to indicatives with positive antecedents, $A \rightarrow C$, is to test whether $\neg A \rightarrow C$ is subject to the Paradox of Indicative Conditionals, and if it is, take that as a sign that both $\neg A \rightarrow C$ and $A \rightarrow C$ are indicatives.

¹⁶ See note 15, above, for remarks about testing a conditional with a positive antecedent.

intuitions must be wrong.¹⁷ We are using how inferences involving conditionals intuitively strike us as a classifying device, even where we know that at least some of the intuitions are misleading.

Applying this test to the ur-examples of the types of conditionals, we find that the test works here. For the ‘indicative’ (A) is subject to the Paradox, while the ‘subjunctive’ (B) is not. (A) does indeed seem to be entailed by

(K \vee L) Either Oswald shot Kennedy, or someone else did

but not by the ‘stronger’

(K) Oswald shot Kennedy

That is, the reasoning

(K \vee L \therefore A) Either Oswald shot Kennedy, or someone else did; therefore, if Oswald didn’t shoot Kennedy, someone else did

while not exciting, certainly gives a very strong appearance of being valid. But

(K \therefore A) Oswald shot Kennedy; therefore, if Oswald didn’t shoot Kennedy, someone else did

intuitively seems about as crazy as does (I \therefore \neg I \rightarrow J).

On the other hand, as we would expect, the ‘subjunctive’ conditional (B) is not subject to the paradox, for, though (B) gives the intuitive appearance of not being entailed by (K), and thus produces one of the two intuitions needed to pass our test, (B) fails to produce the second intuition needed to pass our test, because it does not seem to be entailed by (K \vee L); the inference

(K \vee L \therefore B) Either Oswald shot Kennedy, or someone else did; therefore, if Oswald hadn’t shot Kennedy, someone else would have.

¹⁷ Well, at least one must be wrong if validity is understood in the usual way — as the impossibility of the premiss being true while the conclusion is false. Those who do not think indicative conditionals have truth conditions will often propose other relations between premisses and conclusions to stand in for validity, as understood above, and some such relations may be such that they really do hold for A or C \therefore \neg A \rightarrow C, but not for A \therefore \neg A \rightarrow C. On this way of proceeding, one has to move one’s focus from validity/invalidity to some reasonable facsimile thereof, but, in return, one hopes to secure a reasonable division of the sheep from the goats among inferences.

in contrast to $(K \vee L \therefore A)$, does not give the strong appearance of being valid.¹⁸

When we apply this test to straightforward FDCs, we find that these are subject to the Paradox.

(H) If the house is not painted, it will soon look quite shabby

is subject to the paradox, for

$(M \vee N \therefore H)$ Either the house will be painted, or it will soon look quite shabby; therefore, if the house is not painted, it will soon look quite shabby

does seem unexciting but valid,¹⁹ while

$(M \therefore H)$ The house will be painted; therefore, if the house is not painted, it will soon look quite shabby

¹⁸ In applying this test, I make what I take to be the standard (at least among philosophers) and reasonable assumption that the antecedent of

(B) If Oswald hadn't shot Kennedy, someone else would have is

$(\neg K)$ Oswald didn't shoot Kennedy and that the consequent of (B) is

(L) Someone else did [shoot Kennedy]

(Thus, for instance, when we are told to consider the nearest possible world(s) in which the antecedent of (B) is true, we all know that what we are to ponder is the closest possible world(s) in which $(\neg K)$ is true.) I assume, that is, that the right way to get at the structure of (B) is to view it as built from the sentence frame

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

where the blanks in that frame are filled in with $(\neg K)$ and (L), respectively, to yield

(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). (I follow Lewis 1973, pp. 2–3 and Edgington 1995, pp. 237–8 in my procedure here, and I follow Edgington in my terminology — in my use of 'sentence frame' and 'regimentation'.)

Though I make the above assumption, I do not think my current argumentative purposes here depend on it. I know of no plausible understanding of what (B)'s antecedent and consequent are on which (B) passes our test. (And if (B) really has no antecedent and no consequent, then it also fails our test.) Taking $(\neg K)$ and (L) to be its antecedent and consequent is to give (B) its best plausible chance of passing, and it fails on its best plausible chance. (Similar points may apply to the application of our previous first test to (B).)

¹⁹ In connection with notes 14 and 15, above, note that

Either the house won't be painted, or it will look much better; therefore, if the house is painted, it will look much better

does not seem to produce a clear intuitive appearance of validity.

intuitively seems invalid—to about the extent that $(I \therefore \neg I \rightarrow J)$ and $(K \therefore A)$ seem invalid.

On the basis of our two tests, we have good grounds for thinking that straightforward FDCs are indicative conditionals. They should be classified as indicatives because they have the assertability conditions of indicative conditionals, as we saw in the previous section, and because they are subject to the Paradox of Indicative Conditionals, as we have just seen in this section.

That straightforward FDCs are indicatives in turn gives us reason to think that at least some conditionals of deliberation are indicatives, since, as we have already observed, it is clear, at least on the face of it, that straightforward FDCs can play the role in deliberation that would make them, in our terminology, ‘conditionals of deliberation’.

‘Were’ed-up FDCs also seem to be conditionals of deliberation. We will soon turn to the issue of how they should be classified.

7. Sly Pete and the problem of bad advice

But first we will look at a couple of reasons some might cite for thinking that straightforward FDCs should not be used in the way in question in deliberation: first, that doing so can issue in bad advice for deliberators; second, that it can lead to conflicting advice. Answering these objections will shore up our case that straightforward FDCs are conditionals of deliberation, and will also help to set up our consideration of how to understand and classify ‘were’ed-up FDCs.

Both problems are nicely illustrated by Allan Gibbard’s tale of the gambler, Sly Pete, which we will modify for our current purposes.²⁰

Sly Pete is playing a new card game called *Risk It!* against Gullible Gus. Largely because your henchmen have been hovering about the game and helping him to cheat, the unscrupulous Pete has already won £1,000 from Gus as they move into the final round of the game. The final round of this game is quite simple. A special deck of 101 cards, numbered 0–100, is brought out, shuffled, and one card is dealt to each of the two players. After each player gets a chance to view his own card, but not his opponent’s, the player who is leading going into

²⁰ Gibbard 1981, pp. 226–9, 231–4. In Gibbard’s story, Pete is playing Poker. Some readers, however, do not know much about Poker, and rather than explaining that game, I am using a simpler, made-up game, where the relevant rules are easier to explain. Also, in Gibbard’s story, your two henchmen each hand you a note, and you are unable to tell which note came from which person. I have changed that to accommodate the different philosophical lessons I am looking to draw from the story.

the final round — in this case, Pete — gets to decide whether he wants to ‘play’ or ‘quit’. If he decides to ‘quit’, then he simply keeps the money he has won before this final round — in this case, £1,000. If he instead decides to ‘play’, then his winnings are either doubled or cut to nothing depending on which player holds the higher card: both players show their card, and if the leader’s (Pete’s) is the higher card, the leader’s winnings are doubled — in this case, to £2,000. But if the leader decides to play, and his card is the lower one, he walks away with nothing.

In our first version of the story, your henchman Sigmund (the signaller) has seen what card Gus is holding, has signalled to Pete that Gus’s card is 83, and has received Pete’s return sign confirming that Pete got the message, and knows that Gus is holding 83. Sigmund does not know what card Pete is holding, and so does not know which player holds the higher card, but because he knows that Pete knows what both cards are, and because he is certain that Pete is not stupid enough to ‘play’ if his card is the lower one, it is clear that Sigmund knows that, and is in a position to report to you that:

(O) If Pete plays, he will win

Such information is helpful to you, because, we may suppose, you are making derivative bets on the results of Pete’s game.

But though *Sigmund* seems to know that, and seems in a position to report to you that, Pete will win if he plays, Pete cannot use this conditional that Sigmund knows in *Pete’s* deliberation about whether or not to play. If Pete overhears Sigmund reporting to you that ‘If Pete plays, he will win’, it would be disastrous for Pete to reason as follows: ‘Well, Sigmund seems to know that I’ll win if I play; so, I should play.’ And if Pete knows what Sigmund’s grounds are for his claim, Pete will know not to reason in that disastrous way, if he is a competent consumer of indicative conditionals. This is a case where using a straightforward FDC as a conditional of deliberation leads to trouble: where the conditional would constitute bad advice if used in deliberation.

There are other cases where it seems that indicatives constitute bad advice if used as conditionals of deliberation. I will not try to specify this range of cases exactly, but many of the cases are the types of situations which motivate what is known these days as ‘causal decision theory’. So, for instance, if Sophie is deciding between going to seminary or joining the army, and knows that (even after she has heard about the connection between her career choice and the likelihood of

her having the condition) her choosing to go to seminary would be very strong evidence that she has a certain genetic condition that, if she has it, will almost certainly also result in her dying before the age of 40 years, she has strong grounds to accept that, very probably

(P) If I go to seminary, I will die before the age of 40

Yet, as most can sense, this, plus her desire not to die young, provides her with no good reason to choose against the seminary, for she already either has the genetic condition in question or she does not, and her choice of career paths will not affect whether she has the condition.

It is worth mentioning one other example where using indicatives in deliberation seems to indicate a course of action that I at least accept as irrational—though I suppose that judgement is controversial: it can seem that letting indicatives be your guide would lead one to be a one-boxer in Newcomb's problem. As David Lewis writes:

Some think that in (a suitable version of) Newcomb's problem, it is rational to take only one box. These one-boxers think of the situation as a choice between a million and a thousand. They are convinced by indicative conditionals: if I take one box, I will be a millionaire, but if I take both boxes, I will not ...

Others, and I for one, think it rational to take both boxes. We two-boxers think that whether the million already awaits us or not, we have no choice between taking it and leaving it. We are convinced by counterfactual conditionals: If I took only one box, I would be poorer by a thousand than I will be after taking both. (We distinguish normal from back-tracking counterfactuals, perhaps as in [Lewis 1979], and are persuaded only by the former.) (Lewis 1981, p. 377)²¹

I am a committed two-boxer, like Lewis. So if, as Lewis seems to suppose, letting indicatives be our guide led to choosing one box in Newcomb's problem, I would take that as a serious objection to letting indicatives be our guide in deliberation.

8. Sly Pete and the problem of conflicting advice

Consider a second version of the Sly Pete story. Here, it is your henchman Snoopy (the snooper), rather than Sigmund, who is on the scene. Snoopy does not know the signals, so, though he was able to see Gus's card—which again is 83—he was not able to report that to Pete. But

²¹ For an explanation of Newcomb's problem (due to William Newcomb), see the essay by which the problem was introduced to the philosophical world, Nozick 1969.

Snoopy is able to help *you*, for he moves around so that he sees Pete's card as well as Gus's. Because Snoopy knows that Pete is holding the lower card — 75, let us say — he knows that, and is able to report to you that:

(Oc) If Pete plays, he will not win

Now, consider a third version of the story that combines the first two versions. Pete is indeed holding the lower card, as was specified in version 2, and as was left open in version 1. Sigmund does his signalling and reporting of (O), as in version 1, and leaves the scene, and then Snoopy does his snooping and reporting of (Oc), as in version 2, but each is unaware of what the other has done. As in version 1, Sigmund does know that Pete knows what Gus's card is, and so, since he also knows that Pete will not be stupid enough to play if his is the lower card, Sigmund seems to be speaking appropriately and truthfully when he reports to you that 'If Pete plays, he will win' (O). And as in version 2, Snoopy knows that Pete holds the lower card, and so seems to be speaking appropriately and truthfully when he reports to you that 'If Pete plays, he will not win' (Oc). Are we to suppose that both of these reports *are* true, and that you know both that Pete will win if he plays and also that Pete will not win if he plays? This would appear to be a violation of the 'Law' of Conditional Non-Contradiction — the Law that $A \rightarrow C$ and $A \rightarrow \neg C$ cannot both be true.^{22,23}

There are excellent reasons, roughly of the type that Gibbard gives,²⁴ for thinking that both reports are true — or at least that neither is

²² In our terminology of note 4, this is a 'Law' to the effect that a conditional and its complement cannot both be true. I use scare-quotes because it is very controversial whether this 'Law' is actually true of indicative conditionals. Those who hold that indicative conditionals are equivalent to material conditionals, for instance, will deny this 'Law.'

²³ Why not just say that this would *be* a violation of the 'Law'? Some would try to preserve the Law, while retaining the truth of both reports, by appealing to context-sensitivity: *If Pete plays he will win* is (somehow) Sigmund-true; *If Pete plays he will not win* is (somehow) Snoopy-true.

²⁴ Gibbard 1981, bottom paragraph of p. 231. Gibbard is arguing for the non-falsehood of slightly different, past-directed indicative conditionals. He relies on the point that neither henchman is making any relevant mistake, but does not go on to *explicitly* argue, as he could have, and as I will, that the relevant facts of which they are ignorant are incapable of rendering their statement false. However, I take Gibbard to be at least hinting at this additional argumentative manoeuvre when he makes sure to point out that not only is it the case that neither henchman has any false beliefs about relevant matters of fact, but also that 'both may well suspect the whole relevant truth'.

false. Because they are competent speakers using the relevant assertions in an appropriate manner, we should not charge either Sigmund's or Snoopy's claim with falsehood unless there is some relevant fact which they are getting wrong, or are at least ignorant of, and their mistake about or ignorance of this relevant fact explains why they are making a false assertion. But, as Gibbard points out, neither henchman is making any mistake about any underlying matter. To be sure, each is *ignorant* about an important fact: Snoopy does not realize that Pete knows what Gus's card is, and Sigmund does not know that Pete is holding the lower card. But in neither case does this ignorance on the speaker's part make it plausible to suppose he is making a false claim.

Since for most who hear the story, it is Sigmund's report of (O) that seems the more likely candidate for being false (though perhaps reasonable), let us work this out in his case. Pete in fact holds the lower card, and Sigmund is indeed unaware of that fact. And it seems a very relevant fact: anyone (including Sigmund) who comes to know this fact will thereby become very reluctant to say what Sigmund says—that Pete will win if he plays. However, while Sigmund does not know that Pete holds the lower card, he does recognize the substantial possibility that that is the case. In fact, from Sigmund's point of view, the probability that Pete's card is lower than Gus's is quite high (0.83). (Recall that Sigmund knows that the card Gus holds is 83, but does not know which of the remaining 100 cards Pete holds.) So, *if* this fact—that Pete holds the lower card—were enough to make Sigmund's claim false, then from Sigmund's own point of view, his claim had a very high probability of being false. But a speaker cannot appropriately make a claim that from his own point of view is very probably false. But Sigmund *does* appropriately assert that Pete will win if he plays. So the fact that Pete holds the lower card must not render Sigmund's claim false. But, then, what underlying fact about which Sigmund is ignorant or mistaken does render his claim false? None—there are no good candidates.²⁵ Likewise for Snoopy and his ignorance of the fact

²⁵ There are of course other facts, very closely related to the fact we are focusing on (that Pete holds the lower card) and to each other, that, like the fact we are focusing on, are not known by Sigmund but are extremely relevant here: (1) Pete holds the losing hand; (2) By the rules of the game, given what cards each player is holding, Pete's winnings should be cut to nothing if he plays, etc. But these are also facts that, like the fact that Pete holds the lower card, Sigmund thinks are quite likely, and so the above argument should apply with equal force to these facts as well.

that Pete knows what Gus's card is, for we may suppose that Snoopy thinks it quite likely (though not certain) that Pete knows what Gus's card is, and it will remain the case that Snoopy is in a position to assert (Oc). It is controversial whether indicative conditionals are truth-evaluable. But if your henchmen's conditional reports to you are the sort of things that can be true or false, we must conclude that they are both true.²⁶

And if indicative conditionals are *not* the sort of things that can be true or false, then we must conclude that both of your henchmen's reports have whatever nice property can be assigned to them in lieu of truth—assertable, as opposed to unassertable; assertable and not based on an underlying factual error, as opposed to unassertable or based on underlying error; probable, as opposed to improbable; acceptable, as opposed to unacceptable; or what not.

Thus, indicatives seem not only to give what seems to be bad advice in some cases of deliberation, but can also give conflicting advice: Sigmund's conditional would lead Pete to play; Snoopy's conditional would counsel him not to; both claims are true (or at least 'true-like'), and certainly neither is false.

Which should Pete heed?

9. The solution to the problems: deliberationally useless conditionals

Actually, that is not at all a hard question to answer: of course, Pete should respect Snoopy's (Oc) and not play. (Oc), not (O), is, we will say, *deliberationally useful*—it is the one the agent involved should make use of in deliberating over whether to (try to) make the antecedent true as a way of promoting (or resisting) the consequent being made true. (O), in contrast, is *deliberationally useless*. In the following section, we will discuss the conditions under which a straightforward FDC is deliberationally useless.

But for now, what is vital for us to observe is that, as normal, competent speakers, we demonstrate an awareness of the fact that

²⁶ Note that those who hold that indicative conditionals are equivalent to material conditionals will be quite happy with this story, as they reject the Law of Conditional Non-Contradiction, anyway. In fact, the reasoning you will perform, if you are clever enough, upon receiving both henchmen's reports, is precisely what a material conditional reading of indicative conditionals would indicate: $A \rightarrow C$; $A \rightarrow \neg C$; therefore, $\neg A$ —Pete will not play!

some conditionals, while perhaps useful for other purposes, are not deliberationally useful, for we will not inform a deliberating agent of such conditionals, even though we will so inform others. Note this crucial difference between Sigmund and Snoopy. Based on his knowledge of what both players' cards are, Snoopy is not only in a position to knowingly inform *you* of (Oc), but could also assert (Oc) to the deliberating Pete: if Snoopy had a chance to quickly whisper a conditional to Pete as Pete deliberated over whether to play, he could change (Oc) to the second person and tell Pete, 'If you play, you won't win.' Sigmund, on the other hand, while he knows that (O), and is in a position to inform *you* of (O), cannot inform the deliberating Pete of (O). If Sigmund is a competent speaker, he knows *not* to tell the deliberating Pete that (O), for he knows that (O) is not deliberationally useful.

In saying that (O) is not 'deliberationally useful', I do not mean to be saying that it is useless for *all* deliberations. In our story, (O) may be very useful to *you* as you decide—deliberate about—which derivative bets to place on Pete's game, and, in keeping with that, Sigmund feels very well positioned to inform you of (O). In saying that (O) is not deliberationally useful, I mean more narrowly that it is not useful for someone involved in deciding whether to (try to) make the antecedent true in order to promote or resist the consequent being made true. (And when I write of a 'deliberating agent', or of someone considering a conditional 'in the context of deliberation', I will be using those phrases narrowly, to designate agents in contexts where they are deliberating about whether to make the antecedent of the relevant conditional true as a way of promoting or resisting the conditional's consequent being made true. In this narrow, technical usage, Pete is a deliberating agent with respect to the conditionals (O)/(Oc) in our story, while you are not.) Because he can tell that (O) is not in our narrow understanding 'deliberationally useful', Sigmund cannot competently inform Pete, who is a deliberating agent with respect to (O), of (O), but can report (O) to you, who are not a deliberating agent with respect to that conditional.

Some might suspect that the crucial difference between you and Pete in this story that explains why (O) can be told to you, but not to Pete, is that Pete is the agent mentioned in the antecedent, while you are not. But that the real key difference is that Pete is considering whether to (try to) make the antecedent true in order to promote its consequent being made true can be shown by these new variants of

our cases. Suppose first that, as you tell Sigmund, you are considering calling Pete on his cell phone, to tell Pete whether or not to play, and it is in that connection that you are wondering whether Pete will win if he plays. (Gullible Gus, true to his name, does not object at all to Pete taking phone calls during the game.) Once you have thereby made it clear that you are in our narrow sense a deliberating agent with respect to (O), then, even though you are still not the agent mentioned in (O), Sigmund can no longer inform you of (O), as we can all sense. On the other hand, suppose that in a second new variant of our story, Sigmund does not give Pete any signals, but rather hands Pete a note that says what Gus's card is. Pete has not read the note and so does not yet know what Gus's card is, but he does know that he *will* know Gus's card, as well as his own, when he has to decide whether to play. Now, Pete, like Sigmund, has grounds sufficient for a (first-person version of) (O) that will enable him to use (O) for certain purposes: he knows he will not play unless he has the higher card, so he knows that he will win if he plays. For instance, suppose Pete's wife, who has heard that Pete has won £1,000, but is worrying that he might lose that money in the final round, calls him. Pete can now assure her, by telling her he is *certain* that he will win if he plays. (After all, Pete is certain that he will know what Gus's card before he decides whether to play.) So, here is a case where Pete, though he is the agent involved in (O), can assert the deliberationally useless (O). But of course, he cannot use it in a context of deliberation, as we are narrowly using that phrase: in deciding whether to play, he cannot reason to himself: 'Well, as I already know full well, and as I just told my wife, I will win if I play. So, I should play. I don't even have to read the note!'

The observation that it is a component of linguistic competence not to use a deliberationally useless conditional in the context of deliberation is vital because it provides the solution to our problems of bad advice and of conflicting advice. Yes, some straightforward FDCs, like Sigmund's (O) in the story of Sly Pete, are deliberationally useless. How then can straightforward FDCs function as conditionals of deliberation without causing all kinds of trouble? Because competent speakers/users of these conditionals know not to, and will not, assert/use them in contexts of deliberation where they are deliberationally useless. We do not give or use the bad advice in the cases where an FDC constitutes bad advice. And in cases like our third version of the Sly Pete story, where straightforward FDCs would

otherwise give conflicting advice, it is only the good advice of the deliberationally useful conditional that we give or take.²⁷

When we consider whether Sigmund can assert (O) to the deliberating Pete we have a case in which the assertability of a straightforward FDC diverges from what is predicted by the supposition that such conditionals have the distinctive assertability conditions of indicative conditionals: even where those assertability conditions are met, a straightforward FDC is unassertable in a context of deliberation where it is deliberationally useless. But it seems clear that this divergence in assertability conditions should not make us take back our judgement that straightforward FDCs are indicatives, for, even supposing that straightforward FDCs belong with the indicatives, it is the reverse of surprising that it would be wrong to assert them to deliberating agents in contexts of deliberation when they are deliberationally useless.

10. When FDCs are deliberationally useless: dependence on backtracking grounds

When are FDCs deliberationally useless? A consideration of a lot of cases, most of which I will not present in this article, suggests this answer: FDCs are deliberationally useless when they are based on *backtracking grounds*. We will make do with just three examples,

²⁷ A referee asks whether this aversion of ours to deliberationally useless conditionals might just be a matter of Gricean pragmatics and an instance of our general attempt to say only what is helpful to say. At this point, my goal is to counter the objection that we would get into bad trouble if we let these conditionals be our guides, and I am doing so by observing that we in fact have the ability to recognize when conditionals are deliberationally useless and the tendency not to misuse them in the situations in which doing so leads to trouble (while we do use them in situations in which such use does not lead to trouble). For this purpose, it is fine if this turns out to be susceptible to what some might term a broadly Gricean pragmatic explanation — so long as this is consistent with its accounting for not only our refusal to *assert* the conditionals in the relevant situations, but also our refusal to silently rely on them in our own practical deliberations when they are deliberationally useless. Ultimately, I am sceptical about whether this can, in any helpful sense, be termed a matter of Gricean pragmatics. For, as we will see, we will ultimately want to account for why in certain situations (where the conditionals in question are deliberationally useless, but one is not in a ‘context of deliberation’), a straightforward FDC is appropriately asserted while its ‘were’ed-up analogue is not. It seems unpromising to explain this as a difference in the conversational implicatures generated by the assertions of the conditionals, for conversational implicatures are typically generated by general rules of proper assertion together with the content of the assertions in question. Thus, if the two conditionals generate different conversational implicatures, this will typically be grounded in a difference in their content or meaning. So we will have to somehow get at this difference in content or meaning, and cannot make do with a *merely* pragmatic (in at least some good senses of ‘pragmatic’) explanation.

revisiting cases that we have already discussed in section 8, starting with the Sly Pete story.

Compare the kind of grounds Snoopy has for the deliberationally useful (Oc) with Sigmund's grounds for the useless (O). The reasoning that supports (Oc) for Snoopy involves his beliefs about how things are at and before the time at which the event reported in the antecedent of (Oc) would occur. He then adds to these beliefs the supposition of the antecedent of the conditional—he supposes that Pete will play—and then reasons forward in time and in the causal order, asking what will happen if the antecedent of the conditional is made true, given how he thinks the world is and will be at times prior to and at the time of the antecedent. (Since he knows Pete holds the lower card, adding the supposition that Pete plays leads to the conclusion that Pete loses.) In contrast, Sigmund's knowledge of (O) is based on backtracking grounds. His reasoning, as it would be naturally expounded, involves something like the 'that would be because' locution, which, along with 'that will mean that', are good signs that backtracking reasoning is going on. His reasoning is something like this: 'If Pete plays, that will be because he has the higher card; and then, of course, he will win.' Note that Sigmund does not actually believe that Pete has the higher card. In fact, from Sigmund's point of view, the probability that Pete has the higher card is quite low—0.17. But after he provisionally supposes the antecedent of (O), he reasons backward in the temporal and causal order of things, and conditionally revises his view of what is happening before the time of the antecedent, and then reasons forward in time, using his conditionally revised view of the relevant state of affairs before the time of the antecedent, together with the supposition of the antecedent, to arrive at a conditional view of what will (probably) happen after the time of the antecedent.

Sophie's grounds for (P)—which we can sense is deliberationally useless to her—likewise involve this backtracking pattern of reasoning. After provisionally making the supposition that she goes to seminary, she then reaches backward in the causal order to conditionally alter her view of what her genetic condition is (from agnostic to supposing that she (probably) has the lethal condition), to explain how that antecedent (likely) would become true, and she then conditionally reasons forward to her untimely death.

And the one-boxer's reasoning for

(Q) If I take only one box, I will be a millionaire

at least seems to display that same backtracking pattern — though, again, much is controversial here. At least as it seems to me, his reasoning would have to go something like this: having provisionally supposed he will take just one box, the one-boxer then reasons backward in the temporal and causal order to conditionally determine how much money was (probably) put in the box given that supposition, and then forward to his winning a fabulous fortune.

Our three cases, Sigmund's (O), Sophie's (P), and the one-boxer's (Q) are based on backtracking reasoning, and in each case the conditional is deliberately useless. (I take this to be pretty clear in all three cases, and uncontroversial in the first two.) What is more, I believe a look at many more cases shows that it is conditionals based on backtracking grounds that are blocked from being used in deliberation (in our narrow sense), though they can be asserted in non-deliberating contexts and can be used for various other purposes. A look at various examples will reveal that the mere presence of backtracking grounds does not render a conditional deliberately useless, so long as those backtracking grounds are not needed for the agent to know the conditional. It is *dependence upon*, and not the mere *presence of*, backtracking grounds that render a conditional deliberately useless.²⁸ To the extent that your knowledge of a conditional depends on backtracking grounds, that conditional is deliberately useless to you. If your knowledge comes from someone informing you of the truth of a conditional, then you have sufficient non-backtracking grounds for the conditional only if your informant does. To the extent that you do not know the nature of your informant's grounds, you do not know whether the conditional can be properly used in deliberation.

²⁸ To consider one example illustrating this, make these two modifications to our Sly Pete story. First, imagine that Pete now holds the *higher* card, and that Snoopy sees this to be so. Second, suppose that Snoopy knows that Sigmund has informed Pete what Gus's card is. Now Snoopy has the same backtracking grounds for (O) that Sigmund has: like Sigmund, he is able to reason that since Pete knows what Gus's card is, Pete will play only if he knows that he holds the higher card, and in that case, he will win. But because in this variation Snoopy also has non-backtracking grounds for (O) (he also knows that Pete holds the higher card), he, unlike Sigmund, is in a position to assert that (O) even in contexts of deliberation, as we can intuit. Snoopy's mere possession of backtracking grounds for (O) does not block the assertability of (O) for him in contexts of deliberation, since he also has sufficient non-backtracking grounds for (O), and so is not depending on his backtracking grounds.

There is *much* that needs to be—but, alas, will not be—worked out here. For instance: is it just any old dependence on backtracking grounds that renders a conditional deliberationally useless, or are the cases in which a conditional is deliberationally useless some proper subset of the cases where the speaker's knowledge of it is dependent on backtracking grounds? And, of course, more needs to be said about just what constitutes 'backtracking grounds' in the first place. Yes, it is all a bit messy right now. But it is worth noting that it does not appear that you can avoid this messiness (or at least messiness very much like that which we face here) by supposing instead that counterfactual conditionals should be our guides. Here, it is worth revisiting the quotation from Lewis that we looked at towards the end of section 7. Recall, in particular, this part:

We two-boxers ... are convinced by counterfactual conditionals: If I took only one box, I would be poorer by a thousand than I will be after taking both. (We distinguish normal from back-tracking counterfactuals, perhaps as in [Lewis 1979], and are persuaded only by the former.)

Here, while expressing his conviction that counterfactuals should be our guide, Lewis is quick to point out that it is only 'normal', as opposed to 'back-tracking', counterfactuals that should so guide us. It looks like, however messy the process might be, backtracking must somehow be excluded, whether you think it is counterfactuals or indicatives that are the conditionals of deliberation. Now, there are differences between Lewis's and my proposed exclusion of backtracking in deliberation that are apparent already at this early stage of working out a detailed account. Lewis writes of 'back-tracking counterfactuals', and seems to treat the difference between them and 'normal' counterfactuals as a difference in the meaning or content of the conditionals. (To take or intend a counterfactual as being of the 'back-tracking' variety seems for Lewis to be a matter of taking or intending it to be governed by a particular non-standard type of ordering in the closeness of possible worlds.) In contrast, following what seems to me a more promising path, I focus on the grounds for, rather than the meaning of, the conditionals in question, identifying the problematic (for the purposes of deliberation) conditionals to be those the speaker's or considerer's possession of which is dependent on a certain type of grounds. But in both cases, some exclusion of backtracking needs to be worked out. I hope the beginnings of an account that I have given in this section will be helpful to others who might seek to work this out better.

11. Towards an account of the relation between straightforward and 'were'ed-up FDCs: the standard position

Let us return to our friend Sigmund from the Sly Pete story. Sigmund has signalled to Pete what Gus's card is, and having received the confirmation signal from Pete, Sigmund knows that Pete knows what card Gus holds. Knowing that Pete will play only if Pete holds the higher card, in a non-deliberating context, Sigmund informed you, as we recall, that:

(O) If Pete plays, he will win

Now Sigmund has left the scene, and is thinking things over to himself, or perhaps discussing matters with a friend. (To be able to keep our conditionals future-directed, we will suppose that we have not yet reached the time at which Pete is to announce whether he will play: suppose that, to build suspense, there is mandatory half-hour waiting period between when the cards are distributed in the final round of the game, and when the leader announces whether he will play, and that Sigmund's thoughts and conversations that we are now considering take place during the waiting period.) So long as he is not contemplating intervening in Pete's decision, Sigmund can think and assert (O), and use it for deriving various conclusions, like, for instance, that Pete will leave the game with at least £1,000.

But, in this non-deliberational context, does Sigmund similarly accept, and can he assert, the following 'were'ed-up version of (O):

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

Making only minor, non-substantial changes (explained in the attached note), here is Gibbard's own treatment of (O) and (Ow) in Sigmund's situation:

My informal polls on whether Sigmund accepts (O) have been inconclusive, but most people I have asked think he does. Thus (O) seems to be read as an epistemic conditional, and thus semantically like the future of an indicative rather than a subjunctive conditional. (Ow) is generally treated as a nearness conditional: it is regarded as unlikely, given the information available to Sigmund, and as true if and only if Pete has a winning hand. (Gibbard 1981, pp. 228–9)²⁹

²⁹ At p. 228, Gibbard presents this numbered proposition:

(29) If Pete calls, he'll win

which is in philosophically important respects like our (O). In the key passage, Gibbard discusses his (29), and also two other sentences, the first of which is just like our (Ow),

My own sense, and my own informal polling, match what Gibbard reports here: (Ow) seems wrong for Sigmund, while, as I have already mentioned, (O) seems fine (recalling that it is part of the situation as we are understanding it that Sigmund is not contemplating intervening further in Pete's situation).³⁰

Gibbard's response is to place (O) and (Ow) on opposite sides of the great semantic divide among conditionals. This I take it would be the standard way to treat the relation between such conditionals: whatever exactly one says about the meanings of the conditionals in each of the two camps, and whatever label one gives to the two camps, (O) and (Ow) are on opposite sides of the big division.

Going back to our observations in section 4, above, there are a couple of prima facie problems with such a treatment. First, though I agree with Gibbard that (O) seems right and (Ow) wrong for Sigmund, the difference between the two conditionals seems slight and subtle. They seem to mean at least approximately the same thing, which, together with the sense that one seems right and the other wrong here produces a bit of a sense of puzzlement about the situation. (Perhaps this puzzlement is reflected in the non-conclusive nature of Gibbard's polling results.) At any rate, the relation between (O) and (Ow) seems very different from the very sharp contrast we

and the second of which I have omitted. Here is how Gibbard puts it (except that 'Sigmund' has been substituted in place of Gibbard's name for the relevant character):

My informal polls on whether [Sigmund] accepts (29) have been inconclusive, but most people I have asked think he does. Thus (29) seems to be read as an epistemic conditional, and thus semantically like the future of an indicative rather than a subjunctive conditional.

If Pete were to call, he would win

If Pete called, he would win

are generally treated as nearness conditionals: they are regarded as unlikely, given the information available to [Sigmund] at t_0 , and as true if and only if Pete has a winning hand.

³⁰ Edgington presents another case in which a straightforward FDC seems correct while its 'were'ed-up counterpart seems wrong:

[T]here are two prisoners, Smith and Jones. We have powerful evidence that one of them will try to escape tonight. Smith is a docile, unadventurous chap, Jones is just the opposite, and very persistent. We are inclined to think that it is Jones who will try to escape. We have no reason to accept:

(3c) If Jones were not to try to escape tonight, Smith would

However, we could be wrong in thinking that it is Jones who will escape:

(3b) If Jones doesn't try to escape tonight, Smith will. (Edgington 1995, p. 239)

sense between the likes of (A) and (B). This makes it *prima facie* implausible to place (O) and (Ow) on opposite sides of the great semantic divide among conditionals, the way (A) and (B) are.

Second, on such a view, which are the conditionals of deliberation? The standard answer to this seems to be that only the subjunctives, like (Ow), are conditionals of deliberation. And it would seem a bit surprising that conditionals from both of these very different camps would play this same role in deliberation. But, as we have observed, straightforward FDCs like (O) often seem quite clearly to play the role in deliberation that would make them conditionals of deliberation.

I present these only as *prima facie* problems. There are possible ways around them. One can hold that the semantics for one or the other, or both, of indicative and subjunctive conditionals are in some way or ways quite flexible, and that when they are flexed in certain ways, the meanings of conditionals from opposite camps can approach each other quite closely, while their meanings can also be flexed in other ways to produce sharp contrasts. Perhaps there is something about conditionals being future-directed that interacts with the machinery of the semantics of the conditionals involved to pretty consistently produce pairs of conditionals for the opposing camps that are nevertheless quite close to one another in meaning in the case of FDCs, which could explain both the lack of a sense of sharp contrast in pairs of FDCs and also the appearance (or to some extent, the reality) of straightforward FDCs' usefulness in deliberation: since in future-directed cases indicative conditionals are somehow close in meaning to the corresponding subjunctives, which latter are conditionals of deliberation, one might not go far wrong in using the indicatives in deliberation. (Such a treatment might appeal to the occasional problems one can encounter in using indicatives in deliberation that we have looked at here for support: because future-directed indicatives come close in meaning to the real conditionals of deliberation, you will often do alright following the indicatives, but the problems show that they are only proxies, and you are better off sticking with the real thing.)

I do not find such an attempt to finesse these problems with the standard approach promising, but my purpose is not to conclusively close down such possibilities, but to present an alternative approach which I think will in the end provide a better understanding of the relation between straightforward and 'were'ed-up FDCs, and which, among other things, makes sense of the role both types can play in deliberation.

12. A new account of the relation between the two types of FDCs: ‘were’ed-up FDCs as souped-up indicatives

Recall our discussion in section 5 of the assertability conditions of paradigmatic indicative conditionals like (A) and of straightforward FDCs. As you will recall, I opted for the ‘conditional assertability’ account of the assertability conditions for paradigmatic indicative conditionals, on which one is positioned to assert $A \rightarrow C$ if and only if adding A , as a certainty, to one’s belief set would put one in a position to assert that C — though if you prefer one of the cousins of this account, you can adjust what I am about to say accordingly. You will recall that I then observed that straightforward FDCs have those same assertability conditions — except for a type of counter-example discussed in section 9. But what about ‘were’ed-up FDCs? Checked against *most* cases, they too seem to have those tell-tale assertability conditions of indicative conditionals — which should strongly tempt us to conclude that they too should be semantically classified with the indicatives.

But the generalization that ‘were’ed-up FDCs have the assertability conditions of indicative conditionals faces new counter-examples in addition to the type of counter-example already discussed for straightforward FDCs. Recall that even where the assertability conditions of indicative conditionals are met, straightforward FDCs go unassertable in contexts of deliberation when they are deliberately useless. In the case of ‘were’ed-up FDCs, it does not matter whether the context is one of deliberation: Where a ‘were’ed-up FDC is deliberately useless (where one is relying on backtracking grounds to meet the assertability conditions of indicative conditionals with respect to a ‘were’ed-up FDC), it is wrong to assert that ‘were’ed-up FDC, even if the context is not one of deliberation.³¹ We have already considered such a case: Sigmund, having signalled to Pete what Gus’s card was

³¹ The prohibition on asserting an ‘were’ed-up FDC when it is deliberately useless may not be absolute. For instance, such assertions may seem more-or-less acceptable where the speaker explicitly makes it clear that his grounds are of the backtracking variety. Thus, in the example we are currently considering, Sigmund could correctly say to a friend, ‘If Pete were to play, that would be because he had the higher card. So, if he were to play, he would win.’ Or so it sometimes seems to me — though this seems a delicate call. What is fairly clear is that ‘were’ed-up FDCs have a greater tendency than their straightforward counterparts to indicate deliberational usefulness. This is what is shown by the fact that in examples like the one currently under consideration (and Edgington’s example in note 30, above), in which the speaker is not in a context of deliberation, his grounds render the FDC deliberately useless, but his grounds are not made explicit, the straightforward FDC seems correct but its ‘were’ed-up analogue seems wrong. Beyond this, the details are negotiable.

and having received the confirmation sign from Pete, and then having left the scene, is not in a context of deliberation (recalling again that we are supposing Sigmund is not contemplating intervening further in Pete's game). Sigmund satisfies the distinctive assertability conditions of indicative conditionals with respect to (O) and (Ow), but, as I have already agreed with Gibbard, (Ow) seems unassertable for Sigmund in this situation. More generally, where the assertability conditions characteristic of indicative conditionals are met with respect to a 'were'ed-up FDC, but that conditional is deliberately useless, then that 'were'ed-up FDC is unassertable.³² Otherwise, 'were'ed-up FDCs seem to have the assertability conditions of indicative conditionals.³³

³² Similarly, in Edgington's example (in note 30, above), at least as one would most naturally imagine the case, (3c) and (3b) are deliberately useless: given the types of grounds Edgington gives for it, (3b) does not provide one trying to lure Smith into attempting an escape with a reason to see to it that Jones not try to escape in order to get Smith to try. To apply our deeper level of analysis from section 10, it also seems the speaker's grounds in Edgington's case would have to be of the backtracking variety. The speaker seems to have very strong reasons for thinking a plan is afoot which will certainly be attempted and by which either Smith or Jones will attempt to escape, and has good grounds for thinking that it is Jones who is involved — though the speaker's evidence that one of them is involved in the plot is strong enough that the supposition that one will make the attempt can survive learning that it will not be Jones. When the speaker then provisionally makes the assumption that (3b)'s antecedent is true (that Jones will not try to escape), this causes her first to revise her view of the current situation (before the (future) time of the antecedent) from thinking it is probably Jones that is involved in the plan to thinking it is Smith who is involved, before she reasons forward in time and in the causal order to Smith's attempted escape.) But since we also do not naturally imagine these deliberately useless conditionals as being used in contexts of deliberation, we find the straightforward (3b) to be correct, despite its deliberational uselessness, while the 'were'ed-up (3c) seems wrong. Note, however, that if we imagine that it *is* being used in a context of deliberation — for example, by someone who is involved in trying to get Smith to attempt an escape — (3b) too becomes unassertable, so long as we suppose that the speaker only has the types of grounds for it specified by Edgington. Note also that if you imagine the speaker's grounds to be different, so as to render these conditionals deliberately useful — suppose, for instance, that the speaker knows that Smith will be watching Jones, and Smith will attempt an escape if Jones does not — then both (3c) and (3b) are assertable, whether or not they are used in contexts of deliberation.

³³ Some also find that 'were'ed-up FDCs also pass the other test we have discussed for grouping conditionals with the indicatives: that they are subject to the 'Paradox of Indicative Conditionals'. However, I must report that I have had very mixed results in trying this test on respondents. Half of the test is easily passed: inferences like

The butler will do it; therefore, if the butler were not to do it, the gardener would do it do produce the (fairly clear) intuitive appearance of being invalid. It is the other half of the test that is problematic — whether the likes of

The butler will do it or the gardener will do it; therefore, if the butler were not to do it, the gardener would do it

Why would this be? On the standard view, on which ‘were’ed-up FDCs are subjunctives, but straightforward FDCs are indicatives, it is very surprising that the assertability conditions of these two types of FDCs would be in this way so closely related to each other, while these two types are on opposite sides of the great divide between conditionals. (One can imagine how a backer of the standard view might begin to strain to account for this fact, but it seems to me such an attempt would be quite strained indeed.)

What these observations instead suggest, and what I propose, is that ‘were’ed-up FDCs are souped-up indicative conditionals: they have the same meaning as their straightforward cousins, which we have already argued are indicative conditionals, except that ‘were’ed-up FDCs have a couple of additional components to their meaning. First, as we briefly discussed back in section 4, it seems that ‘were’ing up a conditional can serve the function of calling attention to the possibility that its antecedent is (or will be) false. This, however, does not seem to have the status of a warranted assertability condition of ‘were’ed-up FDCs: if you were to assert a ‘were’ed-up FDC when there is no good reason to call attention to the possibility of the falsity of its antecedent, this does not make your assertion seem wrong or unwarranted. But the second, and more serious, additional component of these conditionals’ meaning does concern a warranted assertability condition: *Whether or not the context is one of deliberation*, one is not in a position to assert a ‘were’ed-up FDC if it is deliberationally useless, or, alternatively, if one is depending on backtracking grounds to meet the assertability conditions of indicative conditionals with respect to it.³⁴ This contrasts with straightforward FDCs,

produce the intuitive appearance of validity. Some find that they do, and so, at least according to their intuitions, ‘were’ed-up FDCs pass this test for indicativeness. Others, however, have the opposite intuition, and most people I have asked find that whatever intuition they have to be less than clear and overpowering. On the hypothesis I am about to put forward concerning the relation between ‘were’ed-up and straightforward FDCs, these mixed results are unsurprising (even given that straightforward FDCs seem to pass the test more clearly), given that the aspect of the meaning of ‘were’ed-up FDCs that goes beyond the meaning of their straightforward counterparts (what the ‘were’-d up FDCs somehow signal about the nature of one’s grounds) is not something guaranteed to be the case by the disjunction that is the premiss of the problematic inference.

³⁴ Again, it is dependence upon, and not just possession of, backtracking grounds that kills assertability, as our modified version of the Sly Pete example from note 28, above, illustrates — when we realize that, in that modified example (in which Pete holds the higher card and Snoopy knows this to be so, but in which Snoopy also still possesses his backtracking grounds for (O)), Snoopy is in a position to assert (Ow) whether or not he is in a context of deliberation.

concerning which I claim that deliberational uselessness blocks assertion only in contexts of deliberation. Thus, these types of FDCs come apart in cases of non-deliberation when they are based on grounds that render them deliberationally useless. In cases of non-deliberation, a 'were'ed-up FDC, $A_w \rightarrow C_w$, is unassertable, even where its straightforward FDC counterpart, $A \rightarrow C$, is assertable, if one is relying on backtracking grounds to be in a position to assert $A \rightarrow C$.

This account seems to me just right in capturing the closeness of 'were'ed-up FDCs to their corresponding straightforward cousins. (No surprise there: I was led to this account primarily as a way to capture just that.) Since backtracking grounds are usually not prominent in the situations in which we are evaluating conditionals, the difference between these two types of FDCs is usually unimportant. Where backtracking grounds *are* indispensable, this account explains why the 'were'ed-up FDCs go unassertable, even in contexts of non-deliberation.

Why would we have a device for souping up FDCs in just that way? It seems that the role FDCs play in deliberation (even in our special, narrow use of 'deliberation') is one of their most important functions. We have observed how it is wrong to assert deliberationally useless straightforward FDCs in contexts of deliberation. But, especially as conditionals are passed as information from one speaker to another, it is easy to lose track of what kind of grounds they are based on, and thus whether they are deliberationally useful. Additionally, it is not always clear to speakers which hearers are deliberating agents (in our narrow sense of that term) with respect to which conditionals. Thus, it does not seem to me at all surprising that our language would develop a device for clearly marking out conditionals as based on the right sorts of grounds to be deliberationally useful.

Finally, while I hope it is clear why, on my views, straightforward and 'were'ed-up FDCs can both function as conditionals of deliberation, it is worth quickly noting that the difference I posit between their meanings also makes it unsurprising that the 'were'ed-up FDCs would have been the FDCs that philosophers would identify as the conditionals of deliberation: in more clearly and consistently screening out the situations in which conditionals become deliberationally useless, 'were'ed-up FDCs are particularly well-suited to use in deliberation.

13. Into the swamp!

In this article, I have remained neutral about the semantics of indicative conditionals.³⁵ As I mentioned in section 1, this is largely because that is a controversial matter. I should emphasize that this is no ordinary, run-of-the-mill philosophical controversy. This is a swamp. The main theories of the meanings of these strange creatures are all over the map — indeed, they are about as far apart as accounts of the meanings of bits of natural language can get — with each of these vastly divergent alternatives claiming substantial allegiance. One can gather something of the nature of this swamp from this very apt closing remark from an encyclopaedia article on ‘Indicative Conditionals’, written by one of the swamp’s most expert navigators, Frank Jackson:

Warning: almost everything about indicative conditionals is controversial, including whether they are well labelled by the term ‘indicative’, and some even deny the validity of *modus ponens*! (Jackson 1998)

While not wading through any of the swamp, I have argued that FDCs of the two varieties we have looked at belong in the swamp, and addressed some general worries about things from that swamp playing a certain key role in deliberation. In addition, I have tried to explain some of the key features of the behaviour of our two types of FDCs by taking the behaviour of paradigmatic indicatives as a starting point, and understanding the workings of our two types of FDCs as variations on that starting point: starting with the knowledge that indicative conditionals mean something or other such that they display such-and-such behaviour (most notably, that they have certain characteristic assertability conditions), we can best understand our two categories of FDCs as being indicative conditionals with added assertability conditions important to their being able to play a key role in deliberation. That seems the most promising way to explain the

³⁵ This neutrality is perhaps most obviously displayed in the paragraph towards the end of section 8 that begins with the words ‘And if indicative conditionals are *not* the sort of things that can be true or false ...’ However, the reader may have noticed that I will write, for instance, of our knowledge or lack of knowledge, of certain indicative conditionals, and may wonder whether I thereby betray a commitment to some propositional understanding of indicative conditionals: an understanding according to which an indicative conditional takes propositions (the antecedent and consequent) as input and generates as output a new proposition (that can be either true or false). In fact, however, I myself do not think indicative conditionals are in that way propositional (see DeRose and Grandy 1999, esp. Sect. 2) But 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. It is that belief that is betrayed by my own rather free use of such talk.

behaviour of these FDCs—behaviour which is in most cases eerily similar to that of paradigmatic indicatives.

Looking forward, the next thing to do would be to enter the swamp, and work out a general account of the meanings of indicative conditionals into which the conclusions of this article could be smoothly plugged. Of course, much work has already been done on the meaning of indicative conditionals. I will close with a couple of brief remarks about the bearing of my conclusions here on that work.

First, of course, if my conclusions here are correct, there is more for theories of indicative conditionals to account for than is usually thought. In particular, ‘were’ed-up FDCs, though usually classified as subjunctive conditionals, would instead fall in the scope of conditionals covered by a theory of indicative conditionals. But I also place another burden on theories of indicative conditionals: they must make sense of the use of indicative conditionals in deliberation. (But I have also addressed a couple of worries one might have about indicatives playing that role that backers of various theories about indicative conditionals can adapt and make use of.) It is tempting to dip just a toe or two into the swamp, and take a quick glance at how this might play out within the contexts of a couple of the leading theories of indicative conditionals, but I find that the topic is too huge and complicated for me to provide a meaningful but compact peek at it here.

Second, if I am right that indicative conditionals are the conditionals of deliberation, that makes indicative conditionals, and the task of understanding them, far more important than they would otherwise seem to be.³⁶

References

- Adams, E. W. 1975: *The Logic of Conditionals*. Dordrecht: D. Reidel.
 Adams, Robert M. 1977: ‘Middle Knowledge and the Problem of Evil’. *American Philosophical Quarterly*, 14, pp. 109–17.
 Braithwaite, R. B. (ed.) 1931: *The Foundations of Mathematics and Other Logical Essays*. London: K. Paul, Trench, Trubner & Co.

³⁶ The ideas expressed in this article (or, in some cases, ancestors of them) were tried out in talks I gave to philosophy departments at the University of Michigan; the University of Arizona; Massachusetts Institute of Technology; the University of Texas, Austin; the University of Glasgow; the University of Dundee; and Princeton University; and also at the 2006 Philosophy of Religion conference at the University of Missouri, Columbia. I thank the audiences at these talks for very helpful discussion. Thanks to the editor at *Mind* and to anonymous referees for some extremely helpful comments and suggestions.

- Cross, Charles B. 2002: 'Doesn't-Will and Didn't-Did'. *Australasian Journal of Philosophy*, Vol. 80, pp. 101–6.
- DeRose, Keith 2002: 'Assertion, Knowledge, and Context'. *Philosophical Review*, 111, pp. 167–203.
- DeRose, Keith and Richard Grandy 1999: 'Conditional Assertions and "Biscuit" Conditionals'. *Nous*, 33, pp. 405–20.
- Dudman, V. H. 1994: 'Against the Indicative'. *Australasian Journal of Philosophy*, 72, pp. 17–26.
- 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.
- Gibbard, Allan and William L. Harper 1978: 'Counterfactuals and Two Kinds of Expected Utility'. In Harper, Stalnaker, and Pearce 1981, pp. 153–90. Originally published in Hooker, Leach, and McClennen 1978.
- Harper, W. L., R. Stalnaker, and G. Pearce (eds) 1981: *Ifs: Conditionals, Belief, Decision, Chance, and Time*. Dordrecht: D. Reidel.
- Hooker, C. A., J. J. Leach, and E. F. McClennen (eds) 1978: *Foundations and Applications of Decision Theory*, Vol. I. Dordrecht: D. Reidel.
- Craig, Edward (ed.) 1998: *Routledge Encyclopedia of Philosophy*. London: Routledge.
- Jackson, Frank 1987: *Conditionals*. Oxford: Basil Blackwell.
- 1998: 'Indicative Conditionals'. *Routledge Encyclopedia of Philosophy* [online encyclopedia], <<http://www.rep.routledge.com/article/X017>>, accessed 20 August 2009.
- Lewis, David 1973: *Counterfactuals*. Oxford: Basil Blackwell.
- 1976: 'Probabilities of Conditionals and Conditional Probability'. *Philosophical Review*, 85, pp. 297–315.
- 1979: 'Counterfactual Dependence and Time's Arrow'. *Nous*, 13, pp. 455–76.
- 1981: 'Why Ain'cha Rich?' *Nous*, 15, pp. 377–80.
- Lycan, William 2001: *Real Conditionals*. Oxford: Oxford University Press.
- Nozick, Robert 1969: 'Newcomb's Problem and Two Principles of Choice'. In Rescher 1969, pp. 114–46.
- Plantinga, Alvin 1974: *The Nature of Necessity*. Oxford: Clarendon Press.
- Ramsey, F. P. 1931: 'General Propositions and Causality'. In Braithwaite 1931, pp. 237–55.

- Rescher, N. (ed.) 1969: *Essays in Honor of Carl G. Hempel*. Dordrecht: D. Reidel.
- Stalnaker, Robert 1975: 'Indicative Conditionals'. *Philosophia*, 5, pp. 269–86, Reprinted in Stalnaker 1999, pp. 63–77.
- 1999: *Context and Content*. Oxford: Oxford University Press.