Draft – 1/22/2006

The Conditionals of Deliberation + Whither Middle Knowledge

Keith DeRose
Yale University

[Paper in progress: Currently a mess in several places, and incomplete]

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 matter-of-factly set up a paper, before we're 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.[1]

That's from a paper from more than 25 years ago, but the extremely widespread assumption it expresses remains in full force today: The conditionals of deliberation are counterfactuals.

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!

---

[1]. "Counterfactuals and Two Kinds of Expected Utility," in W.L. Harper, R. Stalnaker, and G. Pearce, ed., Ifs (D. Reidel Publishing Company, 1978): pp. 153-190; p. 153.

And by "conditionals of deliberation," I mean here conditionals that play the role in deliberation that Gibbard and Harper describe above. We'll 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's meant in it by "indicative conditionals." For our purposes (as for most 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 believe 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).

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's more realistic to suppose that if one libertarian free action is good, then a really primo world would have many such actions, but we'll keep the issue simple by supposing that you desire only one such good action.) So you decide to create such a creature — let's call her Eve — and put her in one of the exactly right possible situations (we'll call this situation S1) where she's free to choose between two courses of action, one good and one bad, and where her performing the good action is exactly what's 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, since either of these two courses of action are inconsistent with Eve's exercising libertarian freedom.[2] 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's fairly widely agreed that you cannot here utilize any <u>simple foreknowledge</u> you might have of such non-conditional propositions as <u>Eve will sin (in S1)</u>. 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 favorably 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,[3] then, though it looks as if you've thereby avoided the trouble, it also looks like you didn't after all know that Eve would sin, since it turns out not to be true that Eve sins, and you cannot have known what wasn't 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 <u>Eve will sin in S1</u>, but is rather "middle knowledge" of certain <u>conditionals</u>. But which conditionals? Ignoring here how the participants

[2]. Expl. of libertarian freedom? — Value, compatibilism, etc.

[3]. Physics stuff, too

in the middle knowledge debate would respond,[4] here is the natural answer to this question: What would really help you in your Divine predicament is knowledge of something like

(C) If I put Eve into situation S1, she will sin,

or, less personally,

(D) If Eve is put into situation S1, she will sin.

Suppose you know those conditionals. Then you'll know not to put Eve into S1, 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 S2 that is such that you foresee that Eve will do the right thing <u>if</u> she's put into S2, and you therefore put her into S2 rather than S1 (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 <u>complement</u>[5] of (C),

(Cc) If I put Eve into S1, she will not sin,

and (D)'s complement, then you know that you are free and clear to put Eve into S1, 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 God can have knowledge of the relevant conditionals – indeed, that is the hot controversy over whether God has "middle

---

[4] See…

[5]. I will call (A) and (Ac), pairs of conditionals sharing the same antecedent, but having opposite consequents -- conditionals of the forms <u>A → C</u> and <u>A → ~C</u> -- "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.)

knowledge." But at least these seem to be the conditionals that it <u>would</u> 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 $5,000 for her car, she will accept.

Of course, not being God, divinely certain knowledge is not in the cards for us. <u>Maybe</u> even knowledge simpliciter 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 $5,000 for her car is 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 gas,
>
> (G) If I start walking to my meeting only 5 minutes before it starts, I will be late,
>
> (H) If the house is painted, it will look much better.

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 insofar 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 isn't completely up to me) its antecedent true. And to the extent I believe one of these conditionals and want its consequent not to be true, I have reason to try to avoid the truth of its antecedent. So 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 paper.

What's more, all of (C) – (H) appear to be indicative conditionals. (And, beyond initial appearances, there are strong reasons for thinking their meaning is such that they should be grouped with (A), as we'll see below in sections _-_.) So it looks like we already have good reason to think that at least some conditionals of deliberation are indicatives.

3. Straightforward Vs. "Were"d-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 it is 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 would probably counsel God to consider this "were"-d up version of (C), rather than (C) itself, in deliberation:

> (Cw) If I were to put Eve into situation S1, 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"-d up counterparts:

> (Ew) If I were to offer Eve $5,000 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 gas,
>
> (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 painted, it would look much better.

These “were”-d 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 “were”s and “would”s inserted in them here and there, they seem somehow most subjunctive than what we’ll call their “straightforward” cousins, the likes of (C) – (H).

But, in part because they are future-directed, neither of these types of future-directed conditionals (henceforth “FDCs”) – neither the “straightforward” nor the “‘were’d 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 are 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” vs. “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”d-Up FDCs

It’s of course possible that the two types of FDCs we’re dealing with belong in different camps. Indeed, some who think they understand what “moods” of verbs amount to and think that these moods are a good indicator of which semantic “camp” conditionals belong might quickly classify our straightforward conditionals as indicatives and the “were”d-up conditionals as subjunctives. Since conditionals of both types are “conditionals of deliberation” as we’re 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”d-up FDC at least doesn’t <u>seem</u> to be much like the relation between (A) and (B). The difference between (A) and (B), as we noted, is quite sharp, and could easily lead one with certain views about Kennedy’s assassination to accept one and reject the other. By contrast, when we compare, for instance, these two conditionals that we’ve already considered:

(C) If I put Eve into situation S1, 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). As William Lycan has observed, there don't seem to be "Adams pairs" of future-directed conditionals.[6] What's more, 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 accept one while rejecting the other. It is very 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 feel of a contradiction. And it's even odd to combine an assertion of either of these conditionals with a question about the acceptability of the other. It's 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?

---

[6] See Lycan, <u>Real Conditionals</u> (Oxford UP, 2001), pp. 162-166, for discussion.

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

When one considers some straightforward FDC and the corresponding "were"d-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"d-up version also somehow signals that its antecedent is improbable, where the type of signaling in question is such that the conditional isn't rendered false if the antecedent is not actually improbable. This suggests itself because many of the situations where one would use the "were"d-up FDC are ones where one believes the antecedent is improbable. But when one more carefully considers the situations in which one might opt for the "were"d-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's quite likely that it is (or will be) false.

But so far this is all just a preliminary look. As we will see, the differences between a straightforward FDC and its "were"d-up analogue can go deeper than 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"d-up counterparts. This close relation at least seems to be quite different from the sharp contrast between the likes of (A) and (B). A main desideratum of an account between the relation of these two types of FDCs is that it make sense of this sense of a very close relation between them. Ultimately, in section _, I will propose a hypothesis as to the relation between these two types of FDCs.

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). This will be done in this and in the following section.

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.[7]

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. (<u>Conditionals</u>, p. 12)

Jackson does not argue that the conditional in the above quotation is assertable to degree 0.5. That is just an observation Jackson makes. I find this quite puzzling. To the extent that I can intuit the degree to which the conditional is assertable, I would give it a value much lower than 0.5.[8] (Forced to assign a number, I'd go for something like 0.06.) It's a fair coin. So I have <u>no idea</u> which of its two sides it will land on if I toss it. I would have to say that I'm 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 doesn't seem a close call: Neither conditional seems anywhere close to half-way assertable. It's tempting to say that I'm in no position <u>at all</u> to assert either conditional, which might tempt one to give them both a flat 0 on the 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 shouldn't just be given a 0. Still, 0.5 seems way too high.

---

[7] Jackson, *Conditionals*, p. 11

[8] assert***i***bility

Indeed, I suspect the only way someone would reach the conclusion that Jackson's conditional has an assertability of 0.5 is if one was already convinced that its assertability was equal to the relevant conditional probability, which one knew to be 0.5. (But then one shouldn't be seeking to support Jackson's hypothesis by citing that assertability value as an observation that matches the theory's prediction.)

So I don't have much sympathy for Jackson's hypothesis. Still, if one is inclined to think that the asssertability 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 <u>is</u> 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."[9] 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. Lewis's hypothesis is

---

[9] Lewis, "Probabilities of Conditionals and Conditional Probability," Philosophical Review 85 (1976): 297-315, at p. 297.

[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. ("Probabilities of Conditionals and Conditional Probability," 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 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."

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 didn't 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 favored account of the assertability conditions of indicative conditionals is a version of the Ramsey test.[12] I'll 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.[13]

While this seems on the right track, 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'm in no position to assert. 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

---

Both Jackson and Lewis take themselves to be following E.M. Adams in their hypotheses about the assertability conditions of indicative conditionals. But if I'm 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'm 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 hasn't 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.

[12] F.P. Ramsey, "General Propositions and Causality," in R.B. Braithwaite, ed., The Foundations of Mathematics and Other Logical Essays (London: K. Paul, Trench, Trubner & Co. Ltd., 1931); reprinted in D.H. Mellor, ed., Foundations: Essays in Philosophy, Logic, Mathematics, and Economics (London: Humanities Press, 1978).

[13] Lycan, Real Conditionals, p. 48.

of that set that I would not be in a position to assert. Then it seems that I'm not in a position to assert the conditional A>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" – is 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 iff 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 iff P's probability for you is sufficiently close to 1, then you'll 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 accept the knowledge account of assertion, on which one is positioned to assert what one knows.[14] This suggests a version of the Ramsey test on which we ask whether , if A were added, as a certainty, to one's belief set, that would allow one to come 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>C iff 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>C is assertable for one; if not, not. Here we leave open the further matter of what it takes generally to be in a position to assert some non-conditional proposition. That's the version of the Ramsey test account that I advocate. It seems to correctly articulate the assertability conditions of paradigmatic indicate 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 indicatives.

[14] For discussion and defense, see my "Knowledge, Assertion, and Context," Philosophical Review 111 (2002): 167-203.

6. Straightforward FDCs are Indicatives: The Paradox of Indicative Conditionals

"Indicative" conditionals like (A) display a truly remarkable property: They are subject to what Frank Jackson has dubbed[15] the "Paradox of Indicative Conditionals." (Being subject to such a paradox is one of the chief ways that indicative conditionals are ill-behaved.) While it's widely recognized that indicatives like (A) have this property, I'm 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 we used in the previous section. There we used the conditions under which the sentences in question seem assertable. Another genus of semantic markers 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 the above name, the Paradox of Indicative Conditionals was nicely set up by Robert Stalnaker,[16] using as his example the paradigmatically indicative conditional,

($\sim I \rightarrow J$) If the Butler didn't do it, the gardener did.

The Paradox consists in two apparent facts about ($\sim I \rightarrow J$); it is a <u>remarkable</u> 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, ($\sim I \rightarrow J$) 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,

($I \vee J \therefore \sim I \rightarrow J$) Either the butler did it or the gardener did it. Therefore, if the Butler didn't do it, the gardener did,

[15]. Ref. I'm at least unaware of anyone using this terminology before Jackson.

[16]. Ref. to "Indicative Conditionals".

they would certainly seem to be performing a perfectly valid inference. However, the strong intuition that ($I \vee J \therefore \sim I \rightarrow J$) is valid clashes with a second strong intuition, namely, that ($\sim I \rightarrow J$) is <u>not</u> entailed by the opposite of its antedent,

(I) The butler did it.

The reasoning,

($I \therefore \sim I \rightarrow J$) 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 ($I \vee J \therefore \sim I \rightarrow J$) is valid or that ($I \therefore \sim I \rightarrow J$) is invalid — must be wrong. Given that (I) entails ($I \vee J$), and given the transitivity of entailment, it just can't be that ($\sim I \rightarrow J$) is entailed by the "weaker" ($I \vee J$) but fails to be entailed by the "stronger" (I).

This suggests a test: If a conditional, $A \rightarrow C$, has the remarkable property of being subject to the "Paradox of Indicative Conditionals" — that is, if it gives the strong appearance of being entailed by *~A or C* but also seems not to be entailed by *~A* — then it should be classified with the indicatives. Note that we are using highly suspect intuitions in applying this test, but also that we are not in any way relying on our intuitions being correct. Indeed, whenever a conditional does elicit the two intuitions that indicate it should be classified with the indicatives, we know that at least one of those intuitions must be wrong.[17] We are using how inferences involving conditionals strike us as a classifying device, even where we know that at least some of the intuitions are misleading.

[17]. Well, at least one must be wrong if validity is understood in the usual way — as the impossibility of the premise being true while the conclusion is false. Those who don't think indicative conditionals have truth conditions will often propose other relations between premises and conclusions to stand in for validity, as understood above, and some such relations will be such that they really do hold for $\sim A$ *or* $C \therefore A \dashrightarrow C$, but not for $\sim A \therefore A \dashrightarrow C$.

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 ($I \therefore \sim I \rightarrow J$).

On the other hand, as we would expect, the "subjunctive" conditional (B) is not subject to the paradox, for (B) does not seem to be entailed by (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,

in contrast to ($K \vee L \therefore A$), does not seem valid, so the subjunctive, (B), is not subject to the Paradox.

When we apply this test to straightforward future-directed conditionals, we find that these are subject to the Paradox.

(H) If the house is painted, it will look much better,

for example, is subject to the paradox, for

($M \vee N \therefore H$) 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 seem unexciting but valid,[18] while

($M \therefore H$) The house won't be painted. Therefore, if the house is painted, it will look much better

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

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

What of the wered-up FDCs? Sometimes they seem to have the assertability conditions of Indicatives; sometimes not. And sometimes they seem subject to the Paradox; sometimes not. We still have to address the relation between straightforward and "were"d-up FDCs.

7. Sly Pete and the Problem of Bad Advice

We will now consider two worries that one might have about the supposition that indicative conditionals being the conditionals of deliberation. Both are nicely illustrated by Allan Gibbard's tale of the riverboat gambler, Sly Pete, which we will slightly modify for our current purposes.[19] The story of Sly Pete will also prepare us to give an account of the relation between straightforward and "were"d-up FDCs.

---

[18] x

[19]. Gibbard (1980), pp. 226-229, 231-234.

Sly Pete is playing a new card game called *Risk It!* against Gullible Gus.[20] 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 the final round — in this case, Pete — gets to decide whether he wants to "take the risk." If he decides not to take the risk, then he simply keeps the money he has won before this final round – in this case, $1,000. If he decides to take the risk, then if his is the higher card, his winnings are doubled – in this case, to $2,000. But if he decides to take the risk, and his card is the lower one, he walks away with nothing.

In our first version of the story, your henchman Sigmund (the signaler) has seen what card Gus is holding, has signaled to Pete that Gus's card is 83, and has received Pete's return sign confirming that he got the message, and knows that Gus is holding 83. Sigmund doesn't know what card Pete is holding, and so doesn't know which player holds the higher card, but because he knows that Pete knows what both cards are, and because he knows that Pete is not stupid enough to "take the risk" if card is the lower one, he knows that, and is able to report to you that:

(O) If Pete takes the risk, 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 takes the risk, Pete cannot use this conditional that Sigmund knows in <u>his</u> deliberation about whether or not to take the risk. If Pete overhears (perhaps through the use of a

[20]. In Gibbard's story, Pete is playing Poker. Some readers, however, don't 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 henchman each hand you a note, and you are unable to tell which note came from which person. I've changed that to accommodate the different philosophical lessons I'm looking to draw from the story.

listening device) Sigmund reporting to you that “If Pete takes the risk, he will win,” it would be disastrous for Pete to reason as follows: “Well, Sigmund seems to know that I’ll win if I take the risk, so I will take the risk.” This is a case where using a straightforward FDC as a conditional of deliberation leads to trouble: Where the conditional constitutes bad advice when used in deliberation.

There are other cases where it seems that indicatives constitute bad advice if used as conditionals of deliberation. I won’t try to specify this range of cases exactly, but many of the cases are the types of situations which motivate what’s known these days as “causal decision theory” – a program I’m at least roughly on board with. 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 result in her dying before the age of 30, she has strong grounds to accept that, very probably,

(P) If I go to seminary, I will die before the age of 30.

Yet, as most can sense, this, plus her desire not to die young, provides her with no reason to choose against the seminary, for she is either has the genetic condition in question or she doesn’t, 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 conditionals seems to indicate a course of action that I at least accept as irrational – though I suppose that judgment 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 [4], and are persuaded only by the former.)…Our decision theory is that of Gibbard and Harper [1], or something similar.[21]

I am a committed two-boxer, like Lewis. So if letting indicatives be our guide led to choosing one box in Newcomb's problem, I'd 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's your henchman Snoopy (the snooper), rather than Sigmund, who is on the scene. Snoopy doesn't 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 Snoopy is able to help <u>you</u>, 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 — 55, let's say —, he knows that, and is able to report to you that:

(Oc) If Pete takes the risk, 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 signaling and reporting of (O), as in version 1, and leaves the scene, and then Snoopy does his snooping and reporting of (Oc), as in 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 has reported to you — quite appropriately, knowingly, and truthfully, it seems — that "If Pete takes

[21] David Lewis, "Why Ain'cha Rich?", *Noûs* 15 (1981): 377-380, at p. 377. The "Gibbard and Harper piece Lewis refers to is the article from which the quotation at the start of this paper was taken.

the risk, he will win." As in version 2, Snoopy knows that Pete holds the lower card, and so has reported to you — again, quite appropriately, knowingly, and truthfully, it seems — that "If Pete takes the risk, he will not win." Are we to suppose that both of these reports <u>are</u> true, and that you know both that Pete will win if he takes the risk and that Pete will not win if he takes the risk? This would appear to be a violation of the "Law" of Conditional Non-Contradiction — the Law that $A \rightarrow C$ and $A \rightarrow \sim C$ can't both be true.[22]

There are excellent reasons, roughly of the type that Gibbard gives,[23] for thinking that both reports are true — or at least that neither is false. Because they are competent speakers using the relevant assertions in an appropriate manner, we shouldn't charge either Sigmund's or Snoopy's claim with falsehood unless there's some relevant fact which they are getting wrong, and their mistake about this relevant fact explains why they are making a false assertion. But neither henchman is making any mistake about any underlying matter. To be sure, each is <u>ignorant</u> about an important fact: Snoopy doesn't realize that Pete knows what Gus's card is, and Sigmund doesn't know that Pete is holding the higher 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's Sigmund's report of (O) that seems the more likely candidate for being false (though perhaps reasonable), let's work this out in his case. Pete in fact holds the lower card, and Sigmund is 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 takes the risk. However, while Sigmund doesn't know that Pete holds the lower card, he does recognize the substantial possibility that that's the case. In fact, from Sigmund's point of view, the probability that Pete's card is lower than Gus's is .83. (Recall that Sigmund knows that Gus holds card 83, but doesn't know which

---

22. Why not just say that this would <u>be</u> a violation of the Law? Some would try to preserve the Law, while retaining the truth of both reports, by appealing to extreme context-sensitivity: only <u>If Pete risks he will win</u> is Sigmund-true; only <u>If Pete risks he will not win</u> is Snoopy-true.

23. See Gibbard (1980), the bottom paragraph on p. 231. Gibbard is arguing for the non-falsehood of slightly different, past-directed conditionals. He relies on the point that neither henchman -- in Gibbard's telling, they're named Zack and Jack -- is making any relevant mistake, but does not argue that the relevant facts of which they're ignorant are incapable of rendering their statement false.

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 of view is probably false. But Sigmund does appropriately assert that Pete will win if he takes the risk. So the fact that Pete holds the lower card must not render Sigmund's claim false. But, then, what does? Nothing — there are no good candidates. Likewise for Snoopy and his ignorance of the fact that Pete knows what Gus's card is. It's 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. (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're clever enough, upon receiving both henchmen's reports, is precisely what a material conditional reading of indicative conditionals would indicate: $A \rightarrow C$; $A \rightarrow \sim C$; therefore, $\sim A$ — Pete will not take the risk!)

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 good 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 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 take the risk; Snoopy's conditional would counsel him not to.

Which should Pete heed?

9. The Solution to the Problems: Deliberationally Useless Conditionals

Not a hard question, actually: Of course, Pete should listen to Snoopy's (Oc) and not take the risk. (Oc), not (O), is, we will say, deliberatively 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. (Oc), by contrast, is deliberationally useless. Under what conditions is a straightforward FDC deliberationally useless? We will address that question in the next section.

But for now, what's vital for us to note is that, as normal, competent speakers, we demonstrate an awareness of the fact that some conditionals, while perhaps useful for other purposes, are not deliberatively useful, for we won't 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 take the risk, he could change (Oc) to the second person and tell Pete, "If you take the risk, you'll 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 deliberatively useful.

Let me emphasize that in saying that (O) is not deliberatively useful, I don't 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 deliberatively useful, I mean narrowly that it is not useful for the agent involved in deciding whether 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 "deliberatively useful," Sigmund won't inform the deliberating Pete, who is a deliberating agent with respect to (O), of (O).

Some might be 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 make its consequent true can be shown by these cases. Suppose you tell Sigmund that you're considering calling Pete on his cell phone, to tell him whether or not to take the risk, and it's in that connection that you're wondering whether Pete will win if he takes the risk. 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). On the other hand, suppose that in a new variant of our story, Sigmund doesn't give Pete any signals, but rather hands Pete a note that says what Gus's card is. Pete doesn't yet know what Gus's card is, but he does know that he will know what card Gus holds before he had to decide whether to take the risk. 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 won't take the risk unless he has the higher card, so he knows that he'll win if he takes the risk. 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. (Gullible Gus, true to his name, doesn't object at all to Pete taking phone calls during the game.) Pete can now assure her, by telling her he's certain of (O). So, here's a case where Pete, though he's the agent involved in (O), can assert the deliberationally useless (O). But of course, he can't use it in a context of deliberation: In deciding whether to take the risk, he can't reason to himself: "Well, as I know full well, I'll win if I take the risk. So I should take the risk.")

The observation that it's a component of linguistic competence not to use a deliberationally useless conditional in the context of deliberation is vital because it provides the solution 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 won't assert/use them in contexts of deliberation when they are deliberationally useless. We don't 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 give conflicting advice, it's only the good advice of the deliberationally useful conditional that we give or take.

10. When FDCs are Deliberationally Useless: Backtracking Grounds

When are FDCs deliberationally useless? An investigation of a lot of cases suggests this answer: FDCs are deliberationally useless when they are based on backtracking grounds. We'll briefly look at just a three of these cases, two of which we've already encountered here, starting the conditionals involved in our Sly Pete stories.

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 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 reasons forward in time, asking what will happen if the antecedent of the conditional is made true, given how he thinks the world is prior to the time of the antecedent. (Since he knows Pete holds the lower card, adding the supposition that Pete takes the risk leads to the conclusion that Pete loses.) By 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 tell-tale 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 doesn't 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 – .17. But he when he supposes the antecedent of (O), he reasons backward in the temporal and causal order of things, and conditionally revises his view of what's happening before the time of the antecedent, and then reasons forward in time, using his conditionally revised view of the relevant state of things 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) likewise display this backtracking pattern of reasoning. After provisionally making the supposition that she chooses going to seminary, she then reaches backward in the causal and temporal order to conditionally revise her view of what her genetic condition is, and then conditionally reasons forward to her untimely death.

And the one-boxer's reasoning for

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

seems to display that same backtracking pattern. Having supposed he will choose just one box, he reasons backward in the temporal and causal order to conditionally determine how much money was put in the box, and then forward to his winning a fabulous fortune.

In the cases Sigmund's (O), Sopie's (P), and the one-boxer's (Q), backtracking reasoning is used, and in each case the conditional is deliberationally useless. (I take this to be pretty clear in all three cases, and obvious in the first two.) What's more, I believe a look at many more cases shows that it's conditionals based on backtracking grounds that are blocked from being used in deliberation 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 bracktracking grounds does not render a conditional deliberationally useless, so long as those grounds are not needed for the agent to know the conditional. It's <u>dependence upon</u>, and not the mere <u>presence of</u>, backtracking grounds that render a conditional deliberationally useless. To the extent that one's knowledge of a conditional depends on backtracking grounds, that conditional is deliberationally 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 don't know the nature of your informant's grounds, you don't know whether the conditional can be properly used in deliberation.

[Note in Lewis quotation about Newcomb's problem that those who go by counterfactuals also have to exclude backtracking as well. This is not a relative disadvantage.]

11. A Hypothesis as to the Relation between Straightforward and "Were"d-Up FDCs

Consider again our friend Sigmund from the Sly Pete story. Sigmund has signaled 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 take the risk only if Pete holds the higher card, he informed you, as we recall, that:

(O) If Pete takes the risk, 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've not yet reached the time at which Pete has announced whether he will take the risk: Suppose that there is mandatory ½-hour waiting period between when the cards are distributed in the final round of the game, and when the leader announces whether he's taking the risk.) So long as he's not considering whether to intervene 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 can Sigmund similarly use and assert the "were"d-up version of (O),

(Ow) If Pete were to take the risk, he would win ?

[Section incomplete]

12. Conclusion: The Role of FDCs in Deliberation

X

OK: That will have to do for now. But I'll quickly say where I think this leads to in the middle knowledge debate.

Whither Middle Knowledge?

What is middle knowledge? Or, more precisely, since it's controversial whether it exists: What would middle knowledge be? In the literature, there seem to be two very different answers to this question that are run together because it's commonly supposed they amount to the same thing. In either case, middle knowledge is knowledge God might have of certain conditionals. But which conditionals? Here's where we get our two different answers.

First – and surely the more important, I'd have to think – is this: Middle knowledge is the knowledge of conditionals that would be useful to God in deliberation. More specifically, it is the knowledge of conditionals the having of which would allow God to exercise Molinistic control over the world — i.e., would allow to know (with Divine certainty) what creatures will freely do if God puts them in situations in which they possess libertarian freedom, so that God does not have to take any real risks and yet His avoiding of risks never makes him pass up on opportunity that, for all He knows, might have paid off.[24] Which conditionals could play this role? Our examples from section 2,

(C) If I put Eve into situation S1, she will sin

(Cc) If I put Even into situation S1, she will not sin,

would seem to be objects of middle knowledge in this sense: If God knows (with Divine certainty) which (if either) of these is true, He seems ready to exercise Molinistic control with respect to the situation in question. If its (Cc) that He knows is true, God knows that He can put Eve into situation S1, and get the result He wants, while taking no real risk that Eve will mess everything up. On the other hand, if it's (C) that He knows is true, He knows to give up on the whole Eve-in-S1 idea, and to try something else instead. In this case, God doesn't get the outcome He most wants (Eve doing right in S1), though that outcome is a genuine metaphysical

[24]. Molinistic control — And you exercised Molinistic control, for, though you didn't get what you what you wanted (Eve doing right in S1), you knew with certainty that this was something you could not get, so you didn't give up any real chance of getting what you wanted.

possibility, but He is still exercising Molinistic control, since He is not passing up an opportunity that, for all He knows, might have paid off. (A way of thinking of Molinistic control is this: God's "menu" of possible outcomes doesn't contain all the metaphysical possibilities, but God does know precisely what is and what is not on His menu, and, even in cases of indeterminism, is able to get precisely the item on the menu that He chooses, without taking any risks.)

But, secondly, you'll be given <u>examples</u> of the kind of conditionals the knowing of which would constitute middle knowledge. The ur-examples of conditionals served up for this purpose are Alvin Plantinga's

> (R) If Curley had been offered $20,000, he would have accepted the bribe[25]

and Robert Adams's

> (S) If President Kennedy had not been shot, he would have bombed North Viet Nam.[26]

And these are notably different from the conditionals one might naturally and innocently reach for as conditionals that could play the role in Divine deliberation that the objects of middle knowledge are supposed to play – like our (C) and (Cc). Of the differences that leap right out, one of the most relevant to our current concerns is this:[27] Our examples are future-directed, while

---

[25]. Ref.

[26]. See Plantinga (1974), p. 174 and Adams (1977), p. 109.

[27]. Another obvious difference: My (B) employs some philosopher-speak in its reference to a "situation S1," while (G) and (H), by avoiding such formalism, have the advantage of being sentences normal speakers of English might actually use! In this case, a little philosopher-speak is justified, for in speaking of the mysterious "situation S1," we may suppose — as I do suppose — that this is an exactly specified situation. There are presumably many significantly different situations that President Kennedy may have faced had he not been shot, and whether or not he would have bombed North Viet Nam may have crucially depended on the exact nature of the situation he faced as he made his choice. (And similar points would hold for Curley.) Deniers of middle knowledge don't mean to be merely claiming that neither (H) nor its complement

> (Hc) If President Kennedy had not been shot, he would not have bombed North Viet Nam

is true merely because there is no determinate fact of the matter as to exactly which situation Kennedy would have faced had he not been shot, and he would have bombed in some of the situations he might have faced, and not in others. They rather mean to be making the stronger

(R) and (S) are backward-looking. To get conditionals like our examples, we'd have to change the ur-examples to

(Rf) If Curley is offered \$20,000, he will accept the bribe

and

(Sf) If President Kennedy is not shot, he will bomb North Viet Nam.

Now, because these sentences are about possible events that are now, as we speak of them, in our past, it's awkward to change the examples to the future-directed (Rf) and (Sf). But, if we're trying to discuss conditionals the knowledge of which would actually be useful to God in exercising providential control, the change is surely one for the better. For if God was to exercise Molinistic control over the relevant events, wouldn't He have to know beforehand what He would then put — if He spoke English! — in terms of (Rf) and (Sf) before deciding whether or not to allow Curley and Kennedy to be in the relevant situations? Rather than being useful for exercising Molinistic control, (R) and (S) look like they're in place only when it's too late to avoid disaster: They appear to be useless examples of "Monday morning quarterbacking," as E.W. Adams aptly pointed out (back in the days when quarterbacks at least sometimes called their own plays).[28] To initial appearances, despite their status as the ur-examples of what middle knowledge would be, (R) and (S) appear to be even so much as relevant to Molinistic control only insofar as they are somehow past-tense versions of (Rf) and (Sf) — insofar as their relation to (Rf) and (Sf) is that they are true now (when it's too late), if and only if (Rf) and (Sf) were true beforehand.

Is it plausible to suppose that that (R) and (S) <u>are</u> past-tense versions of (Rf) and (Sf)? (R) and (S) are paradigmatic subjunctives, while (Rf) and (Sf) are straightforward FDCs, which we've seen, are conditionals of deliberation, and should be classified as indicatives. It seems

---

claim that, with respect to the various particular situations which Kennedy would or might have faced had he not been shot, there's no truth of the matter as to what he would have done in each of those situations, supposing these are situations in which he possesses libertarian freedom with respect to the action in question. (See Plantinga,…xx.) To make our discussion relevant to the real debate over middle knowledge, we will pretend that the antecedents of the conditionals we're dealing with exactly specify situations that free agents might face. We will, then, pretend that (G) and (H) do not vaguely refer to innumerably many significantly different situations Kennedy or Curley might have faced, but rather specify precise situations in their antecedents, as (B) was designed to do.

[28]. ?? E.W. Adams (1975), p. 133.

highly implausible that conditionals on opposite sides of the great divide among conditionals could be in any relation such as one being true at one time iff the other is true at another. At the very least, the assumption that the two explanations of what middle knowledge is amount to the same thing looks very problematic. So the question, "Does God have middle knowledge?" splits into two:

> 1. Does God know the conditionals that are typically cited as objects of middle knowledge – past-direct paradigmatic subjunctives like (R) and (S)?
>
> 2. Does God know the conditionals the knowledge of which would enable Him to exercise Molinistic control in cases of indeterminism?

For the record, my answers are: Yes to 1, no to 2. But some explanation is called for, only a little of which I'll provide here.

<u>Question 1.</u> Plantinga and Adams both discuss applying the standard possible world semantics to the object of middle knowledge, and the result is a mess. But they're operating under the constraint that middle knowledge has to be useful to God in exercising Molinistic control. Now that we've seen the mistake of assuming that there's only one question here, and have rightly divided our questions into two, we ask whether God can know the likes of (R) and (S), freed from the assumption that such knowledge will be useful to God in exercising providential control. I answer by taking the possible worlds semantics very seriously: (R) means that in the closest worlds in which Curley was offered $20,000, he accepted the bribe, and (S) means that in the closest worlds in which Kennedy is not shot, he bombed North Viet Nam. Can God know such things?

Well, as I've already intimated in giving my answer to question 1, I don't see why not. But a couple of caveats are needed. Most importantly, subjunctives like (R) and (S) and highly context-sensitive, and can take on many different contents, and in some contexts (including perhaps abstract philosophical discussions), it may be problematically unclear exactly what do mean. In particular, there are different ways of measuring closeness to the actual world, and a subjunctive may be true on some, but not on other, ways of measuring closeness. In saying God has MK-1, we of course recognize that God knows about all this messiness. He knows, for each way of understanding the exact content of the conditional in question, whether it's true or not on that understanding, and He knows, to the extent that there is a fact of the matter, what is the correct way of understanding the conditional in any context in which it comes up.

Another caveat: In at least some uses (and I suspect, even typically)…
[In progress…]

Question 2.

Very briefly (since I'm running very short on time): To exercise Molinistic control, what God has to know (with Divine certainty) is the likes of (C) and (Cc). Here's how that story will end up, according to me: MK-2 stands or falls with simple foreknowledge. God knows (Cc), for example, iff, having determined to put Eve in S1, He then has simple foreknowledge of the fact that she will not sin. Adams seems to hold out the possibility that one can still accept simple foreknowledge, even having denied middle knowledge.[29] This is a vain hope, according to me. Some defenders of middle knowledge have claimed that middle knowledge stands or falls together with simple foreknowledge. I'm agreeing with that, where by middle knowledge we mean MK-2. But they urge this as part of an attack against the deniers of middle knowledge: "Middle knowledge stands or falls with simple foreknowledge. But, <u>of course</u>, simple foreknowledge stands. We don't want to deny that, do we? So middle knowledge stands as well." But, while I agree with them about the two (at least where middle knowledge is understood as MK-2) standing or falling together, I say they both fall. (God is still omniscient, knowing all truths, because I'm an "Aristotelian" about future contingents, and so there are no facts to be known about the future in the relevant cases.) But this would take us into the many wonders of Open Theism…

[29] ref