Some New Thoughts on Conditionals

Graham Priest

Philosophy Programmes, CUNY Graduate Center, and the University of Melbourne
priest.graham@gmail.com

1 Introduction

A conditional expresses a connection of some kind between two propositions or states of affairs. The relationship is some kind of dependence; but what, exactly, that is, is, of course, the 64 thousand dollar question. The canonical expression of a conditional in English is of the form: If X then Y.[1] But conditionals can be expressed without using 'if' (were I younger, I would go out rocking every night); and not everything which uses the word 'if' is a conditional (if I may say so, you are looking stunning today). The canonical construction suggests that there is only one relation of conditionality. This may be the natural default assumption but, of course, it may be wrong, as many have supposed.

Indeed, nearly everything about the nature of conditionals is philosophically contentious. The consensus of the 1960s concerning the simple-minded theory of the material conditional has blown apart, leaving no present consensus.

This paper is hardly an attempt to solve all of the many issues concerning conditionals. I doubt that anyone is able to do this. Rather, what I wish to do is put conditionals in a new perspective—one which seems to be relatively simple, natural, and provides a straightforward solution to some standard tangles.

2 Conditionals and Imported Information

For a start, some have argued that conditionals are not truth-apt. This, however, cannot be right. Conditional can occur embedded in contexts which require the embedded sentence to be truth-apt, such as: 'Mary believes that if she goes to the party she will have fun' and 'It is possible that if she goes she will have fun'. Conditionals must, then, have truth conditions. A natural thought is that we evaluate a conditional, 'if A then B' by considering situations in which A is true, and seeing if B is true in these. But which situations? Not all of them. Certain information carries over from the actual world, and must hold in them. Thus, consider the conditional 'If global warming continues at its present pace, sea levels will rise by at least two metres before the end of the century'. We are assuming that in this hypothetical—and hopefully (but increasingly unlikely) counter-factual—situation, the laws of physics, and notably those concerning geo-meteorology, are the same as those of the actual world.

Let us call the information that is carried over *the imported information*. It is to be noted that information is imported in a quite different context: determining what holds in a work of fiction. Given a work of fiction, many of the things that hold, hold because of the explicitly say-so of the author. Thus, in the worlds that realise the Holmes novels, Holmes lives in Baker St, because Doyle tell us so. But it is also true that one can't travel from from London to Edinburgh in an hour, that large doses of arsenic kill people, and so on. Doyle never says

[1] The verb of X may be in indicative or subjunctive moods; that of Y can be in other moods, such as interrogative or imperative. How this is possible has to be part of a full story of conditionals, but I will ignore these other moods in what follows.

these things. They are just imported from the facts about the world—or at least, the world of Britain circa 1900. Now, though conditionals and truth in fiction are different issues, it appears to me that the phenomenon of importation is exactly the same in both of them. If, therefore, one could solve the problem of what information, exactly, is imported into a work of fiction, one would have solved the problem of what is imported into an antecedent conditional situation—and vice versa.

Call this the *importation problem*. If we had a solution to it, we would have gone a long way towards answering the 64 thousand dollar question. I'm afraid that I don't (at least presently—one can always hope!). But even without a precise answer, a couple of things are evident.

First, there would appear to be no reason to suppose that irrelevant matters get imported. Thus, it is true that Graham Priest was born in London in 1948. Yet both of the following would appear to be false: In the Holmes novels, Graham Priest would be born in London in 1948; if Emile Zola had written the Holmes novels, Graham Priest would have been born in London in 1948.

Secondly, and most importantly, what information is imported is context-dependent. Thus, suppose that we are driving on a freeway, and the topic of discussion is high-speed transport. You might say 'if this car were a photon, then some cars would travel at about 3×10^8 m/sec'. What is being imported here is the fact that photons travel with the speed of light. But if the topic of discussion were, instead, a hypothetical physics, you might say 'If this car were a photon, then some photons would travel at about 3 m/sec'. Here, what is being imported is the fact that the car is travelling at 100 km/h. I note that the information that is imported may depend on what those who find themselves in the context concerned know. It is not imported simply because they know it, however—much less believe it to be true. It is imported because it *is* true, and bears on the hypothetical scenario envisaged.

3 A Semantics

One way to make these thoughts formally precise is fairly standard.[2] A propositional language contains the connectives $\wedge$, $\neg$, and $>$. $>$ is the conditional. $\vee$ and $\supset$ may be defined in the usual way. The set of formulas is F. An interpretation is a structure $\langle W, \{R_A : A \in F\}, \nu\rangle$. W is a set of worlds, or situations (hypothetical or otherwise). For every $A \in F$, R_A is a binary relation on W; wR_Aw' may be thought to express the fact that w' is a world at which A is true, and at which all the information imported from w holds. ν is a function which maps every world, w, and every propositional parameter, p, to either 1 or 0; we write this $\nu_w(p) = 1$ (or 0). As I noted, what information imports, and so R_A, depends on the context, c. So the R's may be thought of as dependent on a context parameter, c. However, this plays no role in the formal semantics, so I omit mention of it.

Truth at a world, $\Vdash$, is now defined as follows:

- $w \Vdash p$ iff $\nu_w(p) = 1$
- $w \Vdash \neg A$ iff it is not the case that $w \Vdash A$
- $w \Vdash A \wedge B$ iff $w \Vdash A$ and $w \Vdash B$
- $w \Vdash A > B$ iff for all w' such that wR_Aw', $w' \Vdash B$

[2]See Priest (2008), ch. 4. To handle the semantics of counter-logicals properly, the semantics need to be expanded to include impossible worlds, as in ch. 10 (10.7). However, I ignore this point here.

An inference from premises, Σ, to conclusion, A, is valid, $\Sigma \models A$ iff:

- for any interpretation, and $w \in W$: if $w \Vdash B$ for every $B \in \Sigma$, $w \Vdash A$.

These semantics give the basic conditional logic, C. No constraints are put on the R_As. The intuitive interpretation motives some constraints, however. The first is that:

- if wR_Aw' then $w' \Vdash A$

for w' is one of the worlds where A holds. This verifies $\models A > A$. The second is:

- if $w \Vdash A$ then wR_Aw

for if A is true at w, then whatever information is imported from w, it is true at w; hence, w is one of the worlds that w accesses under R_A. This constraint validates: $A, A > B \models B$.

Thus the logic generated by the intuitive understanding explained is at least as strong as C^+. Whether the understanding motivates other constraints, I leave as an open question.

4 East Gate, West Gate

I want to turn, instead, to the issue of how some of the high-profile examples in the literature are accommodated by the understanding I have described.

The first concerns "Gibbardian standoffs", formulated originally by Gibbard (1981). I take the example as cleaned up by Bennett, who explains the scenario as follows.[3]

> Top Gate holds back water in a lake behind a dam; a channel running down from it splits into two distributaries, one (blocked by East Gate) running eastwards and the other (blocked by West Gate) running westwards. The gates are connected as follows: if east lever is down, opening Top Gate will open East Gate so that the water runs eastwards; and if west lever is down, opening Top Gate will open West Gate so that the water will run westwards. On the rare occasions when both levers are down, Top Gate cannot be opened because the machinery cannot move three gates at once.
>
> Just after the lever-pulling specialist has stopped work, Wesla knows that west lever is down, and thinks 'If Top Gate is open, all the water will run westward'; Esther knows that east lever is down, and thinks 'If Top Gate is open, all the water will run east'.

Both Esther and Wesla seem to speak truth, though they appear to disagree with each other. How is this possible? Moreover Southie, who also knows the connections between the gates and the levers, knows nothing of the current settings. Southie can, however, hear what Esther and Wesla say, and knows them to be reliable. Southie concludes that Top Gate is closed. How so?

Take Esther first. In the context in which she finds herself, the information available to her is that the east lever is down. So this information may import into any hypothetical situation she considers. She considers a scenario in which Top Gate is open, and imports the information that east lever is down. In such situations, the water will flow east. Hence she says: If Top Gate is open, all the water will flow east. The situation with Wesla is exactly the same, except that in the context in which he finds himself, the information available to him is that the west

[3] Bennett (2003), p. 85.

lever is down. Both Esther and Wesla speak truly. Their different contexts deliver different importing information.

Next, consider Southie. One might suppose that Southie reason as follows:

- We know, by testimony, that if Top Gate is open the water will flow east, and if Top Gate is open the water will flow west. It cannot flow both east and west, so Top Gate must be closed.

Such reasoning is incorrect, since the two conditionals are true in different contexts, and cannot be conjoined. In the same way, if Alfie, in New York, says 'It's 04.00h', and Beth in London says 'It's 09.00h', Costa in Melbourne, who hears them both, cannot conclude 'Its 04.00h and 09.00h'. What is going on is this. Southie knows that Esther speaks what she knows to be true. The only way for her to say what she says is that she knows that east lever is down. If that were not the case, she could not import it into her hypothetical situation to conclude that, in it, the water is flowing east. Symmetrically, Southie knows from what Wesla says, that west lever is down. On the basis of this, and knowing the facts about the connections between levers and gates, Southie concludes that Top Gate is closed.

5 Interlude: the English Subjunctive

The second example I want to discuss in the notorious Oswald/Kennedy pair, put forward originally by Adams (1970). But first, an interlude on the English subjunctive. The English subjunctive mood is vestigial, and is also the linguistic analogue of an endangered species. However, to the extent that it is still extant, it works like this.

English has only two tenses: present, *I love*, and past (imperfect), *I loved*. Things which are expressed by grammatical tenses in many other languages are expressed in English by using auxiliary verbs, notably *have* and *will*. So we have future, *I will love*, (past) perfect, *I have loved,* pluperfect, *I had loved*, future perfect, I *will have loved*.

Each of the two tenses has an indicative mood and a subjunctive mood. Take the present tense first. For regular verbs, the present subjunctive is the same as the infinitive, (*to*) *love*. But so is the indicative in all persons, except the third person singular, where one adds an *s*. So the only person in which one can tell the difference is the third person singular: *she loves you* (indicative); *I would that she love me* (subjunctive).

The most irregular verb in English is (*to*) *be*. Here, none of the persons in the indicative is the same as the infinitive (*am*/*are*/*is*, *are*/*are*/*are*). The subjunctive is, however, as in regular verbs: that is, the same as the infinitive. So the difference between indicative and subjunctive shows up in all persons. I *am*, *I be*; *she is*, *she be*; *they are*, *they be*.[4]

Turning to the past tense: in regular verbs, the past subjunctive is the same as the past indicative (and so the same in all persons).[5] However, again, the verb (*to*) *be* is irregular, and the indicative conjugates (was/were/was, were/were/were). The subjunctive in all persons is *were*. So the difference shows up in the first and third persons singular.[6]

6 Oswald and Kennedy

We now turn to the oft' cited pair:

[4] As in the subjunctives: If I /she be allowed to speak my/her mind, it will be a very interesting occasion.
[5] As in: I loved her. I would that she loved me.
[6] As in: He told me she was married. I wish that it were not so.

[1] If Oswald didn't shoot Kennedy, someone else did

[2] If Oswald hadn't shot Kennedy, someone else would have

It is usually claimed that the verb of the antecedent in [1] is indicative, and that that in [2] is subjunctive. It is also claimed that the two have the same antecedent. But the first is true, and the second is false (assuming the results of the Warren commission). Hence, there are two kinds of conditionals: indicative and subjunctive.

Is this so? Consider [1]. We may agree that the verb *did* is a past indicative. To evaluate the conditional, we consider a possible situation in which Oswald didn't shoot Kennedy. We import the information that someone shot Kennedy. So in that situation someone else shot Kennedy. So [1] is true.

What of [2]? For a start, it is not entirely obvious that the verb in the antecedent of [2] is in the subjunctive mood. The main verb is *to have*, which appears in the past tense, *had*. This could be indicative or subjunctive. However, [2] may naturally be rephrased as:

[3] If Oswald were not to have shot Kennedy, somone else would have

The main verb here is the past tense of *to be*; and we have a third person singular. So *were* is the subjunctive.

So far, so good. How do we evaluate [2]? Someone who says this, would appear to be saying exactly the same as someone in the past—just prior to the time of the shooting of Kennedy—who says:[7]

[4] If Oswald does not shoot Kennedy, someone else will.

It would appear, then, that the tense and mood of [2] conspire to take [4], and move its point of evaluation to a past point in time. That is, [2] is the tense of [4]. Generally, 'if A were to have been the case, B would have been the case' is the past tense of 'If A is the case, B will be the case'. Call this the *Backshift Thesis*.

Given the Backshift Thesis, we evaluate [2] as follows. We go back to a time just prior to the time at which Kennedy was shot, and evaluate [4] there. We consider a situation where Oswald does not shoot Kennedy. We import what we know from the Warren commission, that Oswald was acting alone. So in that situation, it is false that someone else will shoot Kennedy. So [4] is false of that time, and [2] is false of now.

The past subjunctive does not, then, deliver a different kind of conditional. The moods and tenses of the verbs in the conditional merely conspire to form the past tense of a conditional.[8] [1] and [2] differ in truth value, since the temporal shift makes it natural to import different information into their antecedents.

7 The Backshift Thesis

One might well doubt the Backshift Thesis. Here is a putative counter-example, put to me by Hartry Field.[9]

Professor X is doing an experiment to detect a mooted particle, the tachyon. He sets up an experimental device, which gives a positive result. He exclaims happily (and truly):

[7] Or with a present subjunctive: 'If Oswald shoot not Kennedy, someone else will.

[8] The *general* behaviour and import of tenses and moods in conditionals is a very tricky subject which, fortunately, we may avoid here.

[9] Hartry and I taught a course on Conditionals in New York in the Fall of 2014. Thanks go to him for many enjoyable and insightful conversations.

[5] If the apparatus is working correctly, we will be justified in believing that there are tachyons.

Later he discovers that the apparatus was not working correctly, and whether there are tachyons is still unknown. The Backshift Thesis says that what [5] expresses at the time, is expressed later by:

[6] If the apparatus had been working correctly, we would have been justified in believing that there are tachyons.

But this is false. Had the apparatus been working correctly, it might or might not have shown a positive result; so we might or might not have been justified in believing that there are tachyons.

However, let us pay careful attention to the information that is imported in each conditional. In its context, the natural understanding of [5] imports information including that the experiment has given positive results. To evaluate it, we consider a world where the apparatus is working correctly, add the information that it gives a positive result, and the existence of a justification follows. However, with the same importation, [6] is also true. Had the apparatus been working correctly, then, given that it had positive results, we would have been justified in believing there to be tachyons.

In its context, the natural understanding of [6] imports information including that it is not known whether or not there are tachyons. So, in some hypothetical situations where the apparatus was working correctly, the results are positive; and in some they are negative. It is not the case in all of them that we have good reason to believe that there are tachyons. But with the same importation, [5] is also false. If the apparatus is working correctly, and we do not know whether or not there are tachyons, it does not follow that we will have good reason to believe that they exist. We just do not know what the outcome of the experiment will be.

[5] and [6] therefore stand or fall together. If we import the information that the results were positive, both stand; if we import the information that the existence of tachyons is unknown, both fall. Granted, it is more *natural* to import different information in the two cases. Be that as it may, the apparent difference between [5] and [6] is not due to the falsity of the Backshift Thesis, but to the change in context which motivates different imported information.

8 Present Subjunctives

I have argued that in the Oswald/Kennedy example, the subjunctive antecedent does not betoken a different kind of conditional. It merely shifts the point of evaluation to the past.

If the mere fact that the verb of an antecedent is in the subjunctive mood delivered a different kind of conditional, one might expect to find this with present subjunctives, just as much as past subjunctives. We do not. There is no significant difference between: 'If Julie goes to the party, she will have fun' and 'if Julie go to the party, she will have fun'. To evaluate both conditionals, we consider situations where Julie goes to the party, we import what we know about what sorts of things will happen at the party, what sort of person Julie is, and see whether she will have fun there. The difference between the two conditionals, if there is one, is that with the subjunctive mood, the speaker expresses more hesitation about whether they expect the antecedent situation to be realised.

Some, however, have claimed to find a difference in conditionals even when the subjunctive is a present subjunctive. Edgington gives the following example:[10]

[10] Edgington (1995), p. 239. I have changed her numbering.

> [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 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:
>
> [7] 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:
>
> [8] If Jones doesn't try to escape tonight, Smith will.

So [7] is false, but [8] is true. But what is making the difference here is not the subjunctive, but the information being imported. In [8] we import the information that one of Jones and Smith will try to escape tonight, so in a situation where Jones does not try to escape, Smith does. But if we import the same information into [7], the result is exactly the same. In [7] the natural imported information is that Smith is not the kind of person to try to escape. So in a scenario where Jones does not try to escape, no one does. But if one imports the same information into [8], it is false for exactly the same reason. Perhaps it is more natural to make different importations in the two cases, but one can hear each conditional in both ways.

9 Conclusion

I summarise the main points of this essay. The truth value of a conditional depends on the information which is imported from the actual situation, which is added to that in the antecedent. (And the information concerns what is true, not what is held to be true.) If in all situations that realise both, the consequent is true, so is the whole conditional. If in some it is false, so is the whole conditional. What information is imported is context-dependent, and may change depending on the interests, knowledge, etc. of those using the conditionals.

The idea explains naturally what is going on in some high-profile examples from the literature—perhaps most notably, where there appears to be a difference between conditionals whose antecedents are indicative and conditionals whose antecedents are subjunctive. Past subjunctives indicate a temporal backshift of the point of evaluation, and so affect the information imported. Present subjunctives have no such effect.

I am well aware that this essay is nothing more than the beginning of a discussion. I am sure, for example, that there are many other examples of conditionals that could profitably be examined, and much that could be learned from them. If I have done enough in this essay to make its central ideas worthy of further investigation, I am content. (That's a conditional.)

References

[1] Adams, E. (1970), 'Subjunctive and Indicative Conditionals', *Foundations of Language* 6: 89–94.

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

[3] Edgington, D. (1995), 'On Conditionals', *Mind* 104: 235-309.

[4] Gibbard, A. (1981), 'Two Recent Theories of Conditionals', pp. 211-47 of Harper, W. L., Stalnaker, R., and Pearce, G. (eds), *Ifs*, Dordrecht: D. Reidel Publishing Co.

[5] Priest, G. (2008), *Introduction to Non-Classical Logic*, 2nd edn, Cambridge: Cambridge University Press.