Some New Thoughts on Conditionals

August 6, 2016

Abstract

The paper describes a new way of thinking about conditionals, in terms of information transfer between worlds. This way of looking at things provides an answer to some of the standard problems concerning conditionals, and undercuts the claim that indicative and subjunctive conditionals are distinct.

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.

[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.

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

2.1 Truth Conditions

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.[2] 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) counterfactual—situation, the laws of physics, and notably those concerning geo-meteorology, are the same as those of the actual world.[3] Let us call the information that is carried over *the imported information*.

2.2 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,

[2]Or more generally, when evaluating the conditional $A > B$ at world w, information is carried over from w.

[3]The imported information may or may not itself contain conditionals. Clearly, there is a potential for an infinite regress here, though whether one ever eventuates is moot. Even if it does, however, I see no reason why such a regress should be vicious.

because Doyle tells 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 these things. They are just imported from the facts about the world—or at least, the world of Britain circa 1900. Exactly what information of this kind, however, is imported? Now, though conditionals and truth in fiction are different issues, it appears to me that the phenomenon of importation is very similar in the two cases. If, therefore, one could solve either problem, one would have gone a long way towards solving the other.[4]

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.

2.3 Context

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. There seems to be no reason, then, why all pieces of information consistent with the antecedent import.[5]

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

[4]One difference between the two cases (pointed out to me by Franz Berto) is that in the case of fiction, though the information imported is normally true, it may not be. Suppose, for example, that there is some scientific claim, C, that was generally believed to be true in the late 19th Century, but which is actually false, and that in a Doyle story Holmes presupposes this to successfully solve a case. Doyle, however, does not state C explicitly, since he assumes that his readers all know it. Now, there may well be some post-modern interpretations of the story where Holmes just got lucky. However, the natural interpretation of the story is one in which C holds.

[5]Indeed, when logically consistent antecedents come into the picture (a topic that I will not pursue in this essay), consistency with the antecedent cannot even by a *necessary* condition.

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.

A similar example is given by Goodman.[6] Essentially, it concerns the pair:

- If I were Julius Caesar, I wouldn't be alive in the 21st Century.
- If Julius Caesar were I, he would be alive in the 21st Century.

Both conditionals can be heard as true, but different information is imported in each case. In the first, that Julius Caesar lived in the 1st Century BCE; in the second that I am alive in the 21st Century CE. Note that the antecedents in the two examples are logically equivalent. The order of the terms in the identity simply suggests what information it is that is to be imported.

2.4 East Gate, West Gate

By way of illustrating these ideas, let me explain how they resolve one of the problem areas of conditionals. This concerns "Gibbard Standoffs". These were formulated originally by Gibbard (1981). I take an example as cleaned up by Bennett, who explains the scenario as follows.[7]

> 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.

[6] Goodman (1947), p. 115.
[7] Bennett (2003), p. 85.

> 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 the truth, though they appear to disagree with each other. How is this possible? Moreover Southie, knows nothing of the settings of the levers, but can 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 lever is down. Both Esther and Wesla speak truly. Their different contexts deliver different importing information.

Next, consider Southie. One might suppose that Southie reasons 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. One cannot pool the information that it is 04.00h (in New York) and 09.00h (in London) to conclude that it's 04.00h and 09.00h (anywhere).

What is going on is this. Southie knows that both conditionals are true, relative to their context. So there is information, ι_E and ι_W such that in any world where Top gate is open and ι_E holds, the water flows east, and any world in which Top Gate is open and ι_W holds, the water will flow west. But the actual world is a world where both ι_E and ι_W hold. So if Top gate were actually open, the water would flow east and west, which is impossible. So Top Gate must be closed.

3 Matters Semantics

3.1 A Semantics

One way to make these ideas precise is with a formal semantics. One way do this in a fairly standard.[8]

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', $w' \Vdash B$

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 $w' \Vdash A$

[8] 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.

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.[9]

3.2 Material Validity

While we are on formal matters. Let me comment on another. The first concerns the question of how it is we can reason with conditional inferences that are formally invalid.

There are well known counter-examples to various conditional inferences:

- Transitivity, $A > B, B > C \vdash A > C$. If Hoover had been born in Russia, he would have been a communist. If Hoover had been a communist, he would have been a traitor. Hence, if Hoover had been born in Russia, he would have been a traitor.

- Antecedent Strengthening, $A > C \vdash (A \wedge B) > C$. If you jump off a tall building, you will die. Hence, if you jump off of a tall building and you are wearing a safely harness, you will die.

- Contraposition, $A > B \vdash \neg B > \neg A$. If you take the car, it will not break down *en route*. If the car breaks down *en route*, you don't take it.

[9] An interesting question in this context is as follows. Consider a conditional with an embedded conditional, such as $A \rightarrow (B \rightarrow C)$. Is the information imported in evaluating the outer conditional the same as that imported in evaluating the inner conditional? If it is, this will verify the following condition: if $w_1R_Aw_2$ and $w_2R_Bw_3$ then $w_1R_Bw_3$. For if R_A imports the information ι and R_B imports the information ι (or more), then w_3 is a world where B is true and all of ι is imported. Hence, $w_1R_Bw_3$. Nothing said in this essay settles this matter.

And indeed, these inferences are formally invalid in the above semantics. In what follows, I will discuss only the first of these. Analogous comments apply to the other two.

A salient fact about Transitivity is that we use it to reason, and apparently perfectly correctly, much of the time. Thus, we reason:

- If I am in Paris, I am in the France. If I am in France, I am in Europe. Hence, if I am in Paris, I am in Europe.

This is perfectly good. How can this be if the argument is invalid?[10]

Note that in the Hoover example, there is a crucial difference between the information imported in the conclusion and the information imported into one of the premises. In particular, the second premise imports the information that Hoover was an American. Whatever information is imported into the conclusion, this is certainly not part of it. By contrast, the information imported in each of the three conditionals into the Paris example is exactly the same: the facts of European geography—or at least, different parts of the same body of information. And if the information imported into the conclusion is simply whatever is imported into the premises, the argument is truth-preserving. For consider the inference $A > B, B > C \vdash A > C$. And let the information imported in the two premises be ι_1 and ι_2. Let us evaluate the conclusion where the information imported is $\iota_1 \wedge \iota_2$. We go to the worlds in which A is true and $\iota_1 \wedge \iota_2$ is realised. Since ι_1 is realised, B is true there; and since ι_2 is realised, C is true there, as required.[11]

The inference, then, though not formally valid, is truth preserving because of collateral considerations. We might say, borrowing a term from medieval logic, that it is materially valid.

[10]The main idea in what follows can be found essentially in Whittaker (2016), though it appears in a somewhat different form there.

[11]For the example concerning Antecedent Strengthening, part of the information imported in the premise is that nothing breaks the fall. If the conclusion is evaluated with the same information, it is still true. (The safety harness must have broken.) The example concerning Contraposition is slightly less straightforward. Part of the information imported into the premise is that the car is reliable. If we import that information into the conclusion, then, because of the antecedent, we end up at an impossible world. (If the car breaks down, it was not reliable.) On the present semantics, this makes the conclusion vacuousuly true. I note that if one deploys a semantics that allows for impossible worlds, Contraposition may fail for quite independent reasons: truth preservation forward does not guarantee falsity preservtion backward.

4 Indicative and Subjunctive

4.1 The Oswald and Kennedy Pair

I now want to turn to the question of so called indicative and subjunctive conditionals. A very standard view is to the effect that these are two different kinds of conditionals. This, I think, is false. It is worth getting straight on what, exactly, the English subjunctive is, but this would constitute something of a digression here, so I put the matter in an appendix.

The difference between the two conditionals is usually claimed to be established by the like of the notorious Oswald/Kennedy pair, put forward originally by Adams (1970). These are as follows.

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

[2] If Oswald were not to have shot Kennedy, someone else would have

Here, it is claimed, we have two sentences with the same antecedent, though the mood of the first is indicative, and the mood of the second is subjunctive. Since the first is true and the second is false, there are two kinds of conditionals. Is this so?

Consider [1]. 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.

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:[12]

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

It would appear, then, that the tense and mood of [2] conspire to take [3], and move its point of evaluation to a past point in time. That is, [2] is the past tense of [3]. 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 [3].

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

We import what we know from the Warren commission, namely that Oswald was acting alone. So in that situation, it is false that someone else will shoot Kennedy. So [3] 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.[13] [1] and [2] differ in truth value, since the temporal shift makes it natural to import different information into their antecedents.

4.2 The Backshift Thesis

One might well doubt the Backshift Thesis. Here is a putative counter-example, put to me by Hartry Field.[14]

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):

[4] 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 [4] expresses at the time, is expressed later by:

[5] If the apparatus were to have 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 [4] 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

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

[14] 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.

justification follows. However, with the same importation, [5] 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 [5] 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, [4] 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.

[4] and [5] 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 [4] and [5] is not due to the falsity of the Backshift Thesis, but to the change in context which motivates different imported information.

4.3 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', or more colloquially, 'if Julie were to go to the party, she would 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:[15]

> [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:
>
> [6] 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:
>
> [7] If Jones doesn't try to escape tonight, Smith will.

So [6] is false, but [7] is true. But what is making the difference here is not the subjunctive, but the information being imported. In [7] 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 [6], the result is exactly the same. In [6] 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 [7], 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.

A somewhat different example to the same end is given by Rott:[16]

> Suppose that one Sunday night you approach a small town of which you know that it has exactly two snackbars. Just before entering the town you meet a man eating a hamburger. You have good reason to accept the following indicative conditional:
>
> [8] If snackbar A is closed, then snackbar B is open.
>
> Suppose now that after entering the town, you see that A is in fact open. If the difference between indicative and subjunctive conditionals lay only in the acceptance status of the antecedent,

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

[16] Rott (201+); quotation reformatted. He takes the pair to distinguish between what he calls *epistemic* (indicative) and *ontological* (subjunctive) conditionals. The distinction is from Linström and Rabinowicz (1995), who note that the distinction may not necessarily line up with the grammar.

> we could change the grammatical mood and keep the conditional. But would we really accept the corresponding subjunctive conditional
>
> [9] If snackbar A were closed, then snackbar B would be open.
>
> It seems clear to me that it is not justfiied to accept this conditional. The holders of the two snackbars may well decide on their opening hours entirely independently, so there is no reason to believe that A's being closed makes it any more probable that B is open.

Again, the difference is due to what is being imported, as Rott, in fact, makes clear. The obvious reading of [9] imports the information that the owners of the two bars may be acting independently. With this importation, [8], equally, is false. On the other hand, the natural interpretation of [8] imports the information that at least one of the two snackbars in open. But if one imports just this information into [9] (in the context before one enters the town, and so is in no posion to import further information), it is equally true.

5 Three Objections

Let me finish with three objections to the above account. The first goes as follows.[17] In certain circumstances, the information imported in a conditional need not actually be true. Thus, consider the conditional: if I were a hobbit I would have hairy feet. One can naturally hear this as true. If so, one has imported the information that all hobbits have hairy feet. That's not really true. One thing one might say here is to insist that the conditional *is* true, vacuously. So are lots of others, of course, such as that all hobbits are 10 feet tall. Yet one would not be inclined to say that if I were a hobbit I would be 10 feet tall. The truth that all hobbits are 10 feet tall does not naturally import. Why not? Simply because we are using what holds in Tolkein's world as an appropriate filter.[18]

[17] Many thanks to Damian Melamedoff for point.

[18] It might appear that this strategy can not be applied when the imported information is not a conditional, such as 'If I were to see Bilbo, I would see a hobbit', where the information imported is that Bilbo is a hobbit. However, the imported information can be thought of as the conditional: $\forall x$(if x =Bilbo then x is a hobbit).

Alternatively one might grant that the claim that all hobbits have hairy feet is not really true, but argue as follows. First, note that, in other contexts, one might import different information. Consider: if I were a hobbit, some hobbits would not have hairy feet. This is true when one imports the information that I do not have hairy feet. So what sort of context would it be in which I consider the conditional in question to be true? It would be the sort of context where I am thinking of myself as inhabiting the world as described by Tolkin. So one might more accurately think of the conditional as: if I were a hobbit in a world that realises the Tolkein story, I would have hairy feet (in that world). (That is, where t is a Tolkein world: $t \Vdash$ I am a hobbit $> t \Vdash$ I have hairy feet. The imported information is then that, in Tolkein's world all hobbits have hairy feet ($t \Vdash$ All hobbits have hair feet.) And this *is* true.

Considerations of context also help to resolve a second objection.[19] Suppose that you are considering buying a lottery ticket. The winner is selected randomly, and the odds are a million to one against you. You utter the conditional 'if I buy the ticket I will win'. The conditional seems false, but you buy it, and in fact you win. You say 'I told you so'. What is happening here? When the conditional is uttered, we consider all the worlds in which you buy the ticket, importing the information about the lottery. In some of the worlds that realise these facts, you win; but in the vast majority you don't. So the conditional is, in fact, false. But when you say, ' I told you so', you are importing the fact that you *did* win. When this is imported, then, of course, in any world where you bought the ticket is a word where you won. So with this importation, the conditional is true.

A third objection[20] goes as follows. I hold out a pen, p. The following conditional would seem to be true: if I drop p, it will fall. But the conditional: 'if I drop p and (it is either attached to a helium balloon or it is not) it will fall', is false. But the extra conjunct in the antecedent is a logical truth, so it can make no difference. If one does, indeed, hear the second conditional is false, this is, I think, simply because the extra conjunct changes the information naturally taken to be imported. In the first it is imported that p is not attached to a helium balloon. In the second, no such information is imported, so there a worlds where p falls, and worlds where it does not.[21]

[19]Thanks go to Monique Whittaker for the example.

[20]Thanks go to Larry Horn here.

[21]The point applies, more generally, to so called 'Sobel Sequences' (Sobel (1970)), where the addition of successive conjuncts to the antecedent changes the information naturally

6 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.)[22]

7 Appendix: the English Subjunctive

If one is going to discuss indicative and subjunctive conditionals, it is worth getting straight exactly what the English subjunctive is. In this appendix, I lay the matter out. 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

taken to be imported.

[22]Versions of this paper were read to the Department of Philosophy at the University of Toronto, December 2015, the 20th Amsterdam Colloquium, University of Amsterdam, December 2015, and the 4th Colombian Congress on Logic, Epistemology and the Philosophy of Science, Unversidad de los Andes, February 2016. I am very grateful to those present for helpful questions and observations.

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*.[23]

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).[24] However, again, the verb (*to*) *be* is irregular, and the past 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.[25]

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] Goodman, N. (1947), 'The Problem of Counterfactual Conditionals', *Journal of Philosophy* 44: 113-28.

[23] As in the subjunctives: If I/she be allowed to speak my/her mind, it will be a very interesting occasion.

[24] As in: I loved her. I would that she loved me.

[25] As in: He told me she was married. I wish that it were not so.

[6] Linström S., and Rabinowicz, W. (1995), 'The Ramsey Test Revisited', pp. 157-92 of Crocco, G., Fariñas del Cerro, L., and Herzig, A., *Conditionals and Artificial Intelligence*, Oxford: Oxford University Press.

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

[8] Rott, H. (201+), 'Moody Conditionals: Hamburgers, Switches, and the Tragic Death of an American President', http://www.illc.uva.nl/j50/contribs/rott/rott.pdf

[9] Sobel, J. H. (1970), 'Utilitarianisms: Simple and General', *Inquiry* 13: 394–449.

[10] Whittaker, M. (2016), *The C*3 *Conditional*, PhD Thesis, City University of New York.