

Just a little more on the conditional

PETER SMITH

What, in the shortest summary form, is the take-home message from *IFL* about conditionals? There are two entirely uncontentious points:

- (U1) The material conditional is the only truth-function that can possibly be used in giving a formal rendition of (some) conditionals.
- (U2) The material conditional certainly can't be used to render counterfactual/subjunctive conditionals.

But after that, everything becomes a matter of hot controversy! Does the material conditional really reflect the truth-relevant content of *any* conditionals? If so, which ones? I did claim at least this much:

- (G) Some particular conditionals (instances of certain generalized conditionals) can be adequately rendered – as far as their truth-relevant content is concerned – by material conditionals.

Now, I'm not going to be able to dive deeply into the issues and controversies in this Logicbite. But I think it might still be useful to do three things. First, to get us concentrating again on the headline news, let's revisit those two uncontentious points (U1) and (U2) in the company of Tim Button, and (G) in company of Warren Goldfarb. Second, we'll look at an excerpt from Nick Smith's discussion, where he says quite a bit more about a view which I briefly consider in §19.5. Third, I'll quickly review how a range of different text books cope with the tendentious status of the so-called 'paradoxes of material conditional'.



First, then, back to basics on (U1) with Tim Button in *forall χ* , because I like his example (though I will make minor adjustments to his presentation):

When we introduced the characteristic truth table for the material conditional, we did not say anything to justify it. Let's now offer a justification, which follows Dorothy Edgington.¹

Suppose that Lara has drawn some shapes on a piece of paper, and coloured some of them in. We have not seen them, but nevertheless claim:

If any shape is grey, then that shape is also circular.

As it happens (first situation), Lara has drawn the following:



In this case, our claim is surely true. Shapes C and D are not grey, and so can hardly present *counterexamples* to our claim. Shape A *is* grey, but fortunately it is also circular. So our claim has no counterexamples. It must be true.

That means that, in the described circumstances, each of the following *instances* of our general claim must be true too:

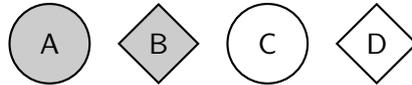
- If A is grey, then it is circular (true antecedent, true consequent)
- If C is grey, then it is circular (false antecedent, true consequent)

¹See the reference at the end of this Logicbite.

- If D is grey, then it is circular (false antecedent, false consequent)

So, if the 'if's here are to be represented by a truth-function, it has to agree with the material conditional in being true on the first, third and fourth line of its truth table.

However, suppose (second situation) Lara had also drawn a fourth shape, thus:



Then our general claim in this situation would have been false. So this claim must also be false:

- If B is grey, then it is circular (true antecedent, false consequent)

Which corresponds to the second line of the truth table for the material conditional.

Thus, from the truth values of our particular conditionals about A, B, C and D in the described situations—which provide us with all possible combinations of truth and falsity in antecedent and consequent—we can read off the truth table for the material conditional.

What this argument shows is that ' \rightarrow ' is the *best* candidate for a truth-functional conditional. Otherwise put, *it is the best conditional that TFL can provide.*

Actually, I think this example in favour of (U1) here shows something more – compare my argument for (G) in §19.3. In the first situation, when the general claim *If any shape is grey, then that shape is also circular* is true, we certainly need the three particular claims of the form *if shape X is grey, X is also circular* to be true for each of the shapes A, C, D; and those particular claims will indeed be true if we take the conditionals here to be material conditionals. But also, the reverse is surely true. In other words, it is enough for the 'if's in the particular claims to behave like material conditionals for the general claim true when it should be (and false when it should be).

So putting it more positively than Button, the bulleted conditionals seem to be examples of ordinary-language conditionals where the 'if's can be read as no-more-than-truth-functional. Here's Goldfarb in his *Deductive Logic* making the same basic point in support of (G) with a different example (recall as you read this that the material condition $P \rightarrow Q$ is equivalent to a negated conjunction $\neg(P \wedge \neg Q)$):

The statement "Every number divisible by four is even" can be rephrased

- (2) No matter what number x may be, if x is divisible by four then x is even.

That is, the statement can be viewed as affirming a bundle of conditionals:

- If 0 is divisible by four then 0 is even
- If 1 is divisible by four then 1 is even
- If 2 is divisible by four then 2 is even

and so on. The interpretation of each of these individual conditionals as material conditionals is just what we need if (2) is to come out true. For among the conditionals in the bundle we find some with true antecedent and true consequent, some with false antecedent and true consequent, and some with false antecedent and false consequent. But we fail to find any with true antecedent and false consequent. Thus each individual conditional, construed as a material conditional, is true. Hence (2) is true, which is exactly what we want. Moreover, each conditional in the bundle amounts to a statement of the form " $\neg(n$ is divisible by four $\wedge n$ is not even)". This bundle of negated conjunctions can then be summed up by

No matter what number x is, it is not the case both that x is divisible by four and x is not even

or, more briefly,

No number is divisible by four yet is not even.

This is a perfectly accurate reformulation of (2).

Generalized conditionals will be treated more fully [later]. They are mentioned here only as an illustration of the central role that the material conditional will play in the analysis of more intricate logical forms.

We will see what Goldfarb means mere in detail later in *IFL*.



Continuing for a moment with Goldfarb, here he is making the uncontentious point (U2):

To be sure, we do not claim that the material conditional is accurate to all uses of “if–then”. In particular, a conditional whose antecedent is in the subjunctive mood cannot be analyzed as a material conditional. Prominent among subjunctive conditionals are the counterfactual ones, for example,

If Robert Kennedy had not been assassinated, he would have become President.

This is called a counterfactual conditional because its antecedent is already assumed to be false. Clearly, then, such conditionals do not behave like material conditionals (nor do we take toward them the everyday attitude of ignoring them once the antecedent is seen to be false). Indeed, they are not truth-functional at all—for, obviously, ordinary usage demands that some counterfactual conditionals with false consequents be true and some with false consequents be false.

Here is Button making the same point more expansively, with a more memorable example:

Consider two sentences:

- (H1) If Hillary Clinton had won the 2016 election, then she would have been the first woman president of the USA.
- (H2) If Hillary Clinton had won the 2016 election, then she would have turned into a helium-filled balloon and floated away into the night sky.

Sentence (H1) is true; sentence (H2) is false, but both have false antecedents and false consequents. (Hillary did not win; she did not become the first woman president of the US; and she did not fill with helium and float away.) So the truth value of the whole sentences are not uniquely determined by the truth value of the parts.

The crucial point is that sentences (H1) and (H2) employ *subjunctive* conditionals, rather than *indicative* conditionals. They ask us to imagine something contrary to fact—after all, Hillary Clinton lost the 2016 election—and then ask us to evaluate what *would* have happened in that case. Such considerations simply cannot be tackled using ‘ \rightarrow ’.

We will say more about the difficulties with conditionals [later]. For now, we will content ourselves with the observation that ‘ \rightarrow ’ is the only candidate for a truth-functional conditional for TFL, but that many English conditionals cannot be represented adequately using ‘ \rightarrow ’. TFL is an intrinsically limited language.

So far, I hope, now so familiar!



We'll now turn to Nick Smith's discussion, where he says more about a view which is ultimately due to Paul Grice and which I briefly discuss in §19.5(b). In the end, I don't endorse that view – though I used to favour it! – but it is very well worth knowing a bit more about the ideas behind this line.

So we start with some of Nick Smith's general description of key Gricean ideas:

One way in which an utterance can be bad (wrong, incorrect) is if the proposition thereby expressed is false. There are, however, many other ways in which utterances can be bad. For example, even though you speak the truth, you may say the wrong thing if your utterance is rude, irrelevant to the discussion at hand, excessively long, unnecessarily complex and hard to follow, phrased in a language that your audience does not speak, too loud, or too soft. In a very influential study of these issues, Paul Grice introduced the Cooperative Principle:

Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.²

In general, other things being equal, utterances not conforming to the Cooperative Principle will be regarded as—to a greater or lesser extent—incorrect. But what exactly is involved in conforming to the Cooperative Principle? To help answer this question, Grice laid out several more specific maxims:

- Maxims of Quantity:
 - Make your contribution as informative as required (for the current purposes of the exchange).
 - Do not make your contribution more informative than required.
- Maxim of Quality:
 - Try to make your contribution one that is true.
 - * Do not say what you believe to be false.
 - * Do not say something for which you lack adequate evidence.
- Maxim of Relation:
 - Be relevant.
- Maxim of Manner:
 - Be perspicuous.
 - * Avoid obscurity of expression.
 - * Avoid ambiguity.
 - * Be brief (avoid unnecessary prolixity).
 - * Be orderly.

In general, other things being equal, utterances not conforming to one or more of these maxims will be regarded as—to a greater or lesser extent—incorrect. For example, suppose that we are all chatting about what we did on the weekend. If I say “I did some stuff,” this will not be regarded as a cooperative contribution—I have not conformed to the first Maxim of Quantity: make your contribution as informative as is required. Likewise, if I launch into a lengthy description beginning with the complete thought process I went through in deciding what to have for breakfast, this will not be regarded as a good contribution—I have not conformed to the second Maxim of Quantity: do not make your contribution more informative than required. . . .

²This is from Paul Grice's hugely influential 1967 William James Lectures ‘Logic and Conversation’, eventually published in a revised form in his 1989 *Studies in the Ways of Words*.

Grice regarded the above maxims as specific instances—that is, specific to the case of conversational exchanges—of more general principles governing any form of rational cooperative behavior. For example, if you are helping me to fix a car and I need four nuts, then I expect you to hand me four, not two or six (cf. the Maxims of Quantity); if you are helping me to make a cake and I need a spoon, then I expect you to hand me a real spoon, not a trick spoon made of rubber (cf. the Maxim of Quality) nor (at this point in the process) an oven cloth (cf. the Maxim of Relation)

A famous sort of example. A driver pauses, winds down the window, to ask me where she can get late-night petrol. I say ‘I believe that the garage a mile down to the left is open 24 hours’. Why does she (rightly) take me not to *know* whether that garage is open? Is it that belief and knowledge are incompatible? Grice would say no: the implication that I don’t know already arises from the fact that if I *did* know, then – as a cooperative speaker, since it takes no more words to say the more helpful thing – I’d say so. So the driver, assuming I am a cooperative speaker – will suppose that I don’t make the stronger, more helpful, claim because I’m not in a position to make it.

We should distinguish, then, (i) the literal content of what we say (and its deductive logical implications) from (ii) what is implied in the context by our saying it in the way we do – what Grice would call *conversational implicatures*. (Exercise, suggested by an example of Nick Smith’s: A friend drops by, as you are sitting down to lunch. “Would you like to join us?” you ask. “I’ve just eaten,” she replies. Why is there a conversational implicature that she doesn’t want to join in the meal?)

Back to conditionals. Nick Smith makes the case for the uncontroversial (U1) and (U2). He then looks at the strong-seeming sort of considerations I give in *IFL* §19.2 for treating simple indicative conditionals as truth-functional. He continues:

So what’s the problem with the proposal that we translate indicative conditionals using \rightarrow ? Well, recall that $\alpha \rightarrow \beta$ is true when α is false or β is true (or both). However, the following conditionals all seem quite wrong, even though the first two have false antecedents and the second two have true consequents:

- (1) If this book is about pop music, it refers to the work of the logician Frege. (false antecedent, true consequent)
- (2) If this book is about pop music, it contains color photos of fruit and vegetables. (false antecedent, false consequent)
- (3) If this book is about logic, its author’s surname starts with ‘S’. (true antecedent, true consequent)
- (4) If this book is about naval architecture, its author’s surname starts with ‘S’. (false antecedent, true consequent)

Therefore, it might seem that we should not translate these propositions into PL using \rightarrow , because then we would translate English conditionals that seem clearly wrong into PL propositions that are straightforwardly true.

What exactly is wrong with the conditionals just given? The problem seems to be that in each case the antecedent has nothing to do with the consequent: believing the antecedent to be true gives us no reason to think that the consequent is true. One might conclude that for an indicative conditional to be true, it is not enough simply for it not to be the case that both the antecedent is true and the consequent is false: there must also be some sort of connection between the two. If we pursue this line of thought, we shall be led to the view that the indicative conditional is not truth functional: whether “if *A* then *B*” is true will depend not just on whether *A* and *B* are true or false but on what they actually say—and in particular, on whether there is the right kind of connection between what *A* says and what *B* says.

The alternative is to defend the view that indicative conditionals have the same truth conditions as material conditionals, by treating the problematic examples in a way that should now be familiar [i.e. by appeal to conversational oddity]: we explain why these conditionals seem wrong in a way that is compatible with their being true. So let us suppose, for the sake of argument, that indicative conditionals have the same truth conditions as material conditionals, and then see what follows from the conversational maxims about the conditions under which indicative conditionals are assertible. For it to be appropriate for one to assert a conditional “if A then B ” (now thought of, for the sake of the argument, as having the same truth conditions as $A \rightarrow B$), one must believe it to be true (Maxim of Quality). That is, one must believe that the actual row is row 1, 3, or 4 in the following truth table:

A	B	$A \rightarrow B$
T	T	T
T	F	F
F	T	T
F	F	T

It must also be the case that one is not confident either way about the truth or falsity of A , and similarly about the truth or falsity of B . To see this, consider the cases in turn. Suppose you are confident that A is false; then (other things being equal) you are in a position to assert $\neg A$ (Maxim of Quality). But saying $\neg A$ is more informative than saying $A \rightarrow B$. . . so you should say $\neg A$ (Maxim of Quantity). Likewise, suppose you believe B to be true; then it would be more informative for you to say B than $A \rightarrow B$. Now suppose you believe A to be true; then, given that you also believe $A \rightarrow B$ to be true, you must (assuming you are minimally rational) believe B to be true; but then, as we have seen, it would be more informative for you to say B than $A \rightarrow B$. . . Likewise, suppose you believe B to be false; then, given that you also believe $A \rightarrow B$ to be true, you must (assuming you are minimally rational) believe A to be false; but then, as we have seen, it would be more informative for you to say $\neg A$ than $A \rightarrow B$. . . Thus, if someone is to utter “if A then B ” correctly, it follows that she is not confident that A is true, nor that it is false, and similarly for B . Yet she must be confident that “if A then B ” is true—that is, that row 2 is not the actual row. The only way she can be in such a position is if there is some sort of connection between A and B : if A 's being true would make it impossible, or very unlikely, for B to be false. Thus, if someone utters “if A then B ,” it will follow from the assumption that she is speaking correctly that she takes there to be some connection between A and B . On this view, then, the existence of such a connection is an implicature: it is not required for the truth of the conditional.

Now let us return to the problematic examples. They seem wrong because there is no connection between their antecedents and the consequents. Thus, we cannot imagine a situation in which any of them would be a good thing to say: if we knew the antecedent was false or the consequent true, it would be more informative to say that straight out, rather than uttering the conditional (Maxim of Quantity); if we did not know either of these things, then we would be in no position to think that the conditional is true (for there is no connection between the antecedent and consequent? hence no reason, if we do not know the truth values of either, to think we are not in a situation where the antecedent is true and the consequent false), and so again we should not say it (Maxim of Quality). We now have an explanation of why the conditionals seem wrong (i.e., we can imagine no situation in which we should want to utter them) that is compatible with their being true.

Which is rather attractively neat. But is it correct? I'll leave it to you to consider the doubts I raise in §19.5(b), and consider whether the sort of view I look at in §19.5(c) might do better. Put it this way: there are – in the abstract – two different mechanisms by which it could be the

case that the material conditional correctly gives the truth-relevant content of indicative ‘if’s, though using the ordinary language ‘if’ *implies* more.

- (i) The extra implication is a *conversational* implicature generated by very general requirements on good practice in rational conversation.
- (ii) The extra implication is *conventionally* generated in the sort of way that the use of ‘but’ conventionally generates some implications over and above what is said by using the colourless ‘and’.

It obviously going to require more general explorations in the philosophy of language to equip us really to adjudicate between such alternatives.



You may well have come across that unhappy phrase ‘the paradoxes of material implication’. Let me stress that there is nothing in the slightest paradoxical about the facts that (i) an inference of the shape $\neg\alpha \therefore (\alpha \rightarrow \beta)$ is valid, and likewise (ii) an inference of the shape $\beta \therefore (\alpha \rightarrow \beta)$ is valid. Those facts are immediate trivial consequences of the definition of the arrow connective, the material conditional, end of story.

Where problems start, of course, is not with those inferences but when we say that (some) ordinary language conditionals have the same truth-conditions as the corresponding material conditionals. Because that of course implies that some inferences of the shape *not-P* \therefore *if P then Q* are valid, and likewise some inferences of the shape *Q* \therefore *if P then Q* are valid. And *those* claims seem puzzling (or, some would say, obviously false). So if there are paradoxes here – and I doubt that ‘paradoxes’ is the right word when we just mean ‘seemingly very implausible claims’! – it is ‘the paradoxes arising from treating (some) ordinary language conditionals as being equivalent to material conditionals’. But that’s less snappy!

OK, that having been said, let’s now have a quick look at how various textbooks respond to the ‘paradoxes’, i.e. respond to the apparent problems in treating (some) ordinary language conditionals as truth-functional. In the three books mentioned so far in this Logicbite, we’ve seen that Nick Smith looks in a kindly way on the Gricean response (though he notes other options in a footnote). While Warren Goldfarb takes something like the pragmatic line of *IFL*, treating the material conditional as a way of rendering the conditional adequate for some purposes (saying that it will emerge in the course of the book when/where this works tolerably well). Tim Button also leaves things unsettled – after giving us another problematic case for the truth-functional rendition of the conditional:

Consider the following sentence:

- 3. It’s not the case that, if God exists, She answers malevolent prayers.

Symbolising this in TFL, we would offer something like ‘ $\neg(G \rightarrow M)$ ’. Now, ‘ $\neg(G \rightarrow M)$ ’ tautologically entails ‘*G*’ (again, check this with a truth table). So if we symbolise sentence 3 in TFL, it seems to entail that God exists. But that’s strange: surely even the atheist can accept sentence 3, without contradicting herself!

(Exercise to think about: how, if at all, could a Gricean like Nick Smith cope with this sort of example?)



What do we find in other text books that we’ve mentioned in other Logicbites? We will just look at a few excerpts.

Lemmon in *Beginning Logic* says this:

It is clear ... that ‘ \rightarrow ’ has logical properties which we should not ordinarily associate with ‘if ... then ...’. This discrepancy is chiefly brought about by the fact that, before we would ordinarily accept ‘if *P* then *Q*’ as true, we should require that *P* and

Q be connected in thought or content, whilst ... no such requirement is imposed on the acceptance of ' $P \rightarrow Q$ '. However, whilst admitting that this discrepancy exists, we may continue safely to adopt ' $P \rightarrow Q$ ' as a rendering of 'if P then Q ' *serviceable for reasoning purposes*, since, as will emerge ..., our rules [of inference for the connectives including ' \rightarrow '] at least have the property that they will never lead us from true assumptions to a false conclusion.

But that's wildly contentious. Take an inference like this:

This iron hasn't been in the fire for an hour. Therefore, if this iron has been in the fire for an hour, I will be able to grasp it without burning myself.

The worry is that this becomes valid, if we render the conditional by ' \rightarrow ' for reasoning purposes – yet it *does* lead from a true premiss (let's suppose) to what looks like a palpably false conclusion. It is no good appealing to some technical result down the road about rules of inference; that's not going to magic away the problem here at the very outset!

Here is Copi in his *Symbolic Logic* getting into similar trouble. He has suggested that many ordinary language conditionals are stronger-than-material.

Although most conditional statements assert more than a merely material implication between antecedent and consequent, we now propose to symbolize any occurrence of 'if-then' by the truth-functional connective ' \supset '. It must be admitted that such symbolizing abstracts from or ignores part of the meaning of most conditional statements. But the proposal can be justified on the grounds that the validity of valid arguments involving conditionals is preserved when the conditionals are regarded as asserting material implications only, as will be established in the following section.

Leaving aside the question of whether Copi later establishes what he says he does, note that it can't be enough. Suppose we have an ordinary language argument Arg involving conditionals, and render it using the material conditional to give a formalized argument Arg . Then even if we can show that when Arg is valid, then Arg will be valid, that tells us nothing about what we can deduce from translating Arg , running a truth-table on Arg and getting the verdict 'valid' on the formal argument – the test result, as far as anything Copi has told us determines, is compatible with Arg being horribly invalid. The truth table test could allow for myriad false positives as far as the original $Args$ are concerned.



Teller in his *A Modern Logic Primer* says, in the usual kind of way, that the material conditional is the best option for a truth-functional rendition of the conditional. But he adds

In many cases, the truth or falsity of an English 'If ... then ...' sentence depends on a nonlogical connection between the truth and falsity of the sentences which one puts in the blanks. The connection is often causal, temporal, or both. Consider the claim that 'If you stub your toe, then it will hurt.' Not only does assertion of this sentence claim that there is some causal connection between stubbing your toe and its hurting, this assertion also claims that the pain will come after the stubbing. However, sentence logic is insensitive to such connections. Sentence logic is a theory only of truth functions, of connectives which are defined entirely in terms of the truth and falsity of the component sentences. So no connective defined in sentence logic can give us a good transcription of the English 'If ... then ...' in all its uses.

We could well worry about the details of this. First, we are thinking about renditions of *singular* conditionals (not more or less disguised generalizations), and we might suspect that a claim like 'If you stub your toe, then it will hurt' is in fact to be read as a generalization, so is beside the present point. But set that aside. Another difficulty is that, yes, your *grounds* for making the stubbed-toe claim might your belief that there is a causal connection between stubbing toes and feeling pain. But so what? It doesn't follow that this that part of what is strictly speaking *asserted* by your conditional claim.

Another case. I bet you \$100 that if Jack takes this novel medicine, he will get better. I'm confident because I think this novel medicine is causally efficacious (my grounds for my claim are a causal belief). You are doubtful, so bet against me. Now, Jack takes the medicine. He indeed gets better. Then, surely, you have to pay up. As it happens, it is a fluke. My grounds weren't good ones – but fortunately Jack would have got better anyway. But no matter. I didn't explicitly say that if Jack takes this novel medicine, then *as a causal result* he will get better. I more cautiously only said that in the case he takes the medicine, he gets better. And by a fluke that was what happened.

The authors of *The Logic Book* make the usual sorts of points around and about (U1) and (U2), but also take a similar line to Teller, assigning some 'if's causal content:

When an English conditional is based on a scientific law, paraphrasing that conditional as a material conditional can be problematic. An example is

If this rod is made of metal, then it will expand when heated.

A simple law of physics lies behind this claim: all metals expand when heated, and the conditional is in effect claiming that if the rod in question is made of metal then heating it will cause it to expand. A paraphrase of this causal claim as a material conditional does not capture this causal connection. The failure to capture such causal connections may or may not be acceptable, depending on the context and on what questions we are asking about the sentence or set of sentences being paraphrased.

But again it is one thing to say that the quoted conditional is *based* on a causal law, and it is another thing to say that it is itself *making* a causal claim – and what are we supposed to make of the idea that the quoted conditional is 'in effect' making a causal claim?

Here's Hurley in a *Concise Introduction to Logic*, wandering in the same neck of the woods:

The material conditional is a kind of conditional statement whose truth value depends purely on the truth or falsity of the antecedent and consequent and not on any inferential connection between antecedent and consequent. Since many conditional statements in ordinary language express such an inferential connection, when the horseshoe operator is used to translate them, part of their meaning is left out. For example, compare the following two statements:

If Shakespeare wrote Hamlet, then the sun rises in the east.
If ice is lighter than water, then ice floats in water.

The first statement expresses no inferential connection between antecedent and consequent, so using the horseshoe operator to translate it results in no loss of meaning. However, the second statement does express such a connection. The fact that ice is lighter than water is the reason why it floats. Accordingly, when the horseshoe operator is used to translate the second statement, this special meaning is lost.

What entitles Hurley to say that one conditional expresses an inferential connection and the other doesn't?

It might indeed be the case that the grounds someone has for asserting the two conditionals might be of different sorts. Imagine we are in some sort of quiz or puzzle situation, having to work out from clues the truth or falsity of some ten quite unrelated claims. The clues may lead us to the position where we know that at least one of P and Q is true, and so we can assert 'if not- P then Q ' without making any assumption that there is an internal connection between the topics of the antecedent and the consequent. So this would be a paradigm case where the conditional "expresses no inferential connection between antecedent and consequent" and we might imagine the first conditional above would be like that. While our confidence in the second conditional is no doubt based on causal knowledge about how the world works. But the fact that the two conditionals might naturally be asserted on different sorts of grounds, doesn't mean that

they express different sorts of things – it doesn't mean that the different *grounds* (there being an underlying “connection” or not) change what is actually *said* by the two uses of ‘if’.

But at least Hurley (who has more to say) makes an effort here! Simpson in his *Essentials of Symbolic Logic* insouciantly remarks

In most cases, translating English ‘if . . . then’ sentences into symbols is not difficult.

The key symbol in question is the horseshoe with its truth-functional definition. There isn't even a whisper about subjunctive or counterfactual conditionals, let alone other issues. I suppose we can forgive authors for not wanting to tangle with the problems about conditionals . . .but sweeping them under the carpet is really not an option.



There is a great deal more to be said. For pointers to further debates, you can't do better than read the *Stanford Encyclopaedia* article by Dorothy Edgington, ‘Conditionals’, 2014 (<http://plato.stanford.edu/entries/conditionals/>). You'll see why I suggested we should back off from any ambition to *faithfully capture* in a formal setting the core content of (a wide family of cases of) ordinary ‘if’s. Perhaps we should instead regard the material conditional as a suitable *substitute* for the vernacular conditional, one which is sufficiently conditional-like to be happily adopted for some important formal purposes, while being elegantly clear, perfectly understood, and very easy to work with. That's the line I read into e.g. Goldfarb, and which is very explicitly taken in *IFL*.