15 The material conditional

We have seen how to evaluate arguments whose essential logical materials are the three connectives ‘and’, ‘or’ and ‘not’. We now ask: can our simple truth-table techniques be smoothly extended to deal with arguments involving ‘if’?

15.1 Some arguments involving conditionals

(a) Consider the following elementary arguments:

A If Jack bet on Eclipse, he lost his money. Jack did bet on Eclipse. So Jack lost his money.

B If Jack bet on Eclipse, he lost his money. Jack did not lose his money. So Jack did not bet on Eclipse.

C If Jack bet on Eclipse, he lost his money. So if Jack didn’t lose his money, he didn’t bet on Eclipse.

D If Jack bet on Eclipse, then he lost his money. If Jack lost his money, then he had to walk home. So if Jack bet on Eclipse, he had to walk home.

E Either Jack bet on Eclipse or on Pegasus. If Jack bet on Eclipse, he lost his money. If Jack bet on Pegasus, he lost his money. So, Jack lost his money.

Each of those five arguments involving conditionals looks intuitively valid. Contrast the following four arguments:

F If Jack bet on Eclipse, he lost his money. Jack lost his money. So Jack bet on Eclipse.

G If Jack bet on Eclipse, he lost his money. So if Jack lost his money, then he bet on Eclipse.

H If Jack bet on Eclipse, he lost his money. Jack did not bet on Eclipse. So Jack did not lose his money.

I If Jack bet on Eclipse, he lost his money. So if Jack did not bet on Eclipse, he did not lose his money.

These are plainly fallacious arguments. Suppose foolish Jack lost his money by betting instead on Pegasus, another hopeless horse. Then in each of these four cases, the premisses of the argument can be true and conclusion false.

The validity or invalidity of the arguments **A** to **I** evidently has nothing specifically to do with Jack or betting, but is due to the meaning of the conditional construction introduced by 'if'. It would be good to have a general way of handling such arguments relying on conditionals. On the face of it, 'if' seems to be a binary connective like 'and' and 'or', at least in respect of combining two propositions to form a new one. So the obvious question is: can we straightforwardly extend our techniques for evaluating arguments involving those other connectives to deal with conditionals as well?

15.2 A handful of basic principles governing conditionals

We will regard conditionals *if A, C* and *if A then C* as stylistic variants, so we treat 'then' here as akin to punctuation. (In some conditionals of the second form, perhaps the 'then' does more work: but set any such cases aside.) And we start by introducing some standard terminology used in talking about conditionals:

Given a conditional *if A then C*, we refer to the 'if' clause *A* as the *antecedent* of the conditional, and to the other clause *C* as the *consequent*.

The *converse* of *if A then C* is the conditional *if C then A*.

The *contrapositive* of *if A then C* is the conditional *if not-C then not-A*.

We will also use the same terminology when talking about the formal analogues of ordinary conditionals.

Now, generalizing from **A** and **B**, the following pair of basic principles of conditional reasoning evidently hold good in core cases (as we will see in the next chapter, conditionals come in different flavours, but in this chapter concentrate on conditionals like those in the previous section):

(MP) An inference step of the form *if A then C, A, so C* is valid, and is standardly referred to as *modus ponens*.

(MT) An inference step of the form *if A then C, not-C, so not-A* is valid, and is standardly referred to as *modus tollens*.

Moreover, modus ponens and modus tollens are very *useful* ways of establishing truths. That is to say, we often have cases where *if A then C* and *A* are both true together, so we can infer that *C* is true. And we also have cases where *if A then C* and *not-C* are both true together, so we can infer that *not-A* is true.

Let's note some other fundamental principles which intuitively govern ordinary conditionals. First:

(CF) A conditional *if A then C* can't be true if in fact *A* is true and *C* is false.

That's obvious. But it is worth noting that both (CF) and (MT) are immediate consequences of (MP). Take three propositions of the form (i) *A*, (ii) *if A then C*, (iii) *not-C*. These are inconsistent, since by modus ponens (i) and (ii) imply *C*, which contradicts (iii). Hence, if we keep (ii) and (iii) we have to reject (i), which is the principle (MT). And if we keep (i) and (iii) we have to reject (ii), which is the principle (CF).

Now, the facts (MP), (MT), and (CF) apply not just to simple one-way conditionals *if A then C* but also to two-way biconditionals *A if and only if C* (we'll say more about biconditionals in §15.5). But a distinguishing fact about one-way conditionals is this: unlike biconditionals, they are not usually reversible. As the invalidity of **G** illustrates,

(NC) From a conditional premiss of the form *if A then C*, we can't usually infer the converse conditional *if C then A*.

No doubt, if I have won a Nobel prize, then I am clever. But it doesn't follow that if I am clever, then I have won a Nobel prize.

We should also mention a fifth intuitively appealing general principle about conditionals, which again distinguishes them from biconditionals. This is a version of the principle of *Conditional Proof*:

(CP) If the argument *A*, *B* ∴ *C* is valid, then so is the argument *A* ∴ *if B then C*.

The idea is this: if we want to prove a conditional *if B then C* (given the assumption *A*), then a natural way of doing this is to temporarily suppose that *B* is true and then show that (given *A*) we can infer *C*. For example, by the principle (CP), the validity of **B** implies the validity of **C**. And if **F** were valid, **G** would be valid too; and hence the invalidity of **G** implies the invalidity of **F**.

We will revisit (CP), exploring further examples in §[**ref**]. Meanwhile, we simply note that it comes in more general versions. For we can have more than one assumption like *A* in play. Or we may have no such assumption, in which case the principle becomes this: if the argument *B* ∴ *C* is valid, then the conditional *if B then C* is necessarily true.

Here's one application of that last version which we'll use in a moment. The argument *A and B* ∴ *A* is trivially valid; so *if A and B, then A* is an equally trivial necessary truth.

15.3 Introducing the material conditional

(a) In headline terms, our technique for dealing with arguments involving 'and', 'or' and 'not' has two stages: Step (1) presupposes that the ordinary language connectives are close enough in their core meanings to the corresponding connectives of PL for logically important features of the original argument to get carried over. (If the translation were too loose, a verdict on the PL rendition couldn't be carried back into a verdict on the original argument.) Step (2) then depends on the fact that PL connectives are truth-functional. That's needed to ensure that, given any valuation of the relevant propositional atoms, we can work out whether it makes the premisses of the PL argument true and conclusion false.

Hence, if our approach in the last few chapters is to be straightforwardly extended to cover arguments involving the conditional, we need to be able to render conditional propositions into a formal language with a suitable binary truth-function, and to do this in a way which preserves enough of the core meaning of the conditional.

(b) It is easy to see that there is in fact only one truth-functional conditional-like connective which satisfies inferential principles parallel to the five principles (MP) etc. which apply to ordinary conditionals. Why so?

Suppose we add the symbol '⊃' to a PL-style language in order to represent a connective which is intended to both truth-functional and conditional-like. The issue then is how to fill in its truth-table. With an eye on (CF), we have already completed the second line: the supposed truth-functional conditional has to be false when it has a true antecedent and false consequent. So how does the rest of the table go? We'll give two arguments, slow and fast, for the same completion.

α	γ	$(\alpha \supset \gamma)$
T	T	?
T	F	F
F	T	?
F	F	?

(i) Taking it step by step, note that the first line on the truth-table can't also be completed with 'F'. Otherwise $(\alpha \supset \gamma)$ would never be true when α is true: and in that case we would never be able to use modus ponens with two true premisses $(\alpha \supset \gamma)$ and α to prove that γ.

Similarly, the last line has to be completed with 'T'. Otherwise $(\alpha \supset \gamma)$ would never be true when γ is false: so we'd never be able to use modus tollens with two true premisses $(\alpha \supset \gamma)$ and $\neg\gamma$ to prove that $\neg\alpha$.

Which just leaves one entry for the truth-table to be decided. But if we put 'F' for value of the candidate truth-functional aconditional on the third line, $(\alpha \supset \gamma)$ would have the same truth-table as $(\gamma \supset \alpha)$, making our supposed conditional reversible, contrary to the requirement that a conditional-like connective satisfies the analogue of (NC).

α	γ	$(\alpha \supset \gamma)$
T	T	T
T	F	F
F	T	?
F	F	T

So we get the following:

The only possible candidate for a conditional-like truth-function is the *material conditional*, defined by the truth-table

α	γ	$(\alpha \supset \gamma)$
T	T	T
T	F	F
F	T	T
F	F	T

(ii) For a fast track argument, take the wff '((P ∧ Q) ⊃ P)'. If '⊃' is to be conditional-like, this should be true, come what may (by the remark at the end of the previous section). But the values of the antecedent/consequent of the conditional in this wff can be any of T/T, or F/T, or F/F (depending on the values of 'P' and 'Q'), and the wff needs to evaluate as T in each case. So that forces the same completion of the truth-table for the conditional.

Why 'material conditional' (the name is due to Bertrand Russell)? Let's not worry about that. It's one of those labels that has stuck around, and which is still absolutely standard, long after its exact original connotation has been largely forgotten.

This truth-function is also sometimes known as the *Philonian* conditional, after the ancient Stoic logician Philo, who more-or-less explicitly defined it.

15.4 Ways in which '⊃' is indeed conditional-like

Since PL languages can express all truth-functions, they can already express the material conditional. But it is conventional to add a special symbol for this truth-function, adjusting the syntactic and semantic rules to get richer languages we will call PLC languages. The details are *exactly* as you would expect, as you will see in §15.7, so we do not need to pause over them just yet. So let's continue to use the symbol '⊃' for the material conditional, assuming it is now part of the formal languages we are using.

How well does the material conditional behave as a formal counterpart to the ordinary-language conditional? In this section, we make a start on discussing the issue by confirming that if we translate the various examples in §15.1 using '⊃', and run truth-table tests, we do get the right verdicts. (Remember from §13.5 that we can extend the notion of tautological entailment, and the use of truth-table tests, beyond basic PL languages, so long as we are still dealing with truth-functional connectives.)

So, start with the modus ponens inference

A If Jack bet on Eclipse, he lost his money. Jack did bet on Eclipse. So Jack lost his money.

Rendered into a suitable PLC language with its added material conditional, this straightfowardly goes into

A′ (P ⊃ Q), P ∴ Q

And running a full truth-table test (not exactly hard!), we get

P	Q	(P ⊃ Q)	P	Q
T	T	T	T	T
T	F	F	T	F
F	T	T	F	T
F	F	T	F	F

There are no bad lines, with true premisses and false conclusion. Therefore argument **A′** is tautologically valid. Hence it is plain valid – as was the original version **A**.

We will skip the modus tollens argument **B** – it can be left as a trivial exercise to confirm that the obvious formal translation is tautologically valid too. So let's next look at

C If Jack bet on Eclipse, he lost his money. So if Jack didn't lose his money, he didn't bet on Eclipse.

This can be rendered as

C′ $(P \supset Q) \therefore (\neg Q \supset \neg P)$

Running a truth-table test, this time we get

P	Q	(P ⊃ Q)	(¬Q ⊃ ¬P)
T	T	T	T
T	F	F	F
F	T	T	T
F	F	T	T

So again the inference **C′** is tautologically valid and hence valid. And there's more: the table shows that a material conditional not only tautologically entails its contrapositive, but is tautologically *equivalent* to it. But that still tracks the behaviour of mainstream ordinary-language conditionals, for this is also true:

Straightforward cases of ordinary-language conditionals are equivalent to their contrapositives.

For example, 'If Jack bet on Eclipse, he lost his money' not only entails but is entailed by 'if Jack didn't lose his money, he didn't bet on Eclipse'.

Next example:

D If Jack bet on Eclipse, then Jack lost his money. If Jack lost his money, then Jack had to walk home. So if Jack bet on Eclipse, Jack had to walk home.

This translates into our PLC language as, say,

D′ $(P \supset Q), (Q \supset R) \therefore (P \supset R)$

And running up a truth table using the now familiar shortcuts (evaluating the conclusion first, and then the premisses in turn but only as needed), we get

P	Q	R	(P ⊃ Q)	(Q ⊃ R)	(P ⊃ R)
T	T	T			T
T	T	F	T	F	F
T	F	T			T
T	F	F	F		F
F	T	T			T
F	T	F			T
F	F	T			T
F	F	F			T

There are no bad lines, so the argument is tautologically valid, and therefore plain valid, corresponding again to an intuitively valid vernacular argument.

The next example illustrates another common form of valid inference, often called *argument by cases*:

E Either Jack bet on Eclipse or on Pegasus. If Jack bet on Eclipse, he lost his money. If Jack bet on Pegasus, he lost his money. So, Jack lost his money.

This translates, using the material conditional again, into e.g. the following:

E′ $(\mathsf{P} \lor \mathsf{Q}),\ (\mathsf{P} \supset \mathsf{R}),\ (\mathsf{Q} \supset \mathsf{R})\ \therefore\ \mathsf{R}$

You won't be surprised to learn that this too it tautologically valid. Here's a truth-table to confirm that (again, we evaluate the conclusion first, ignore good lines, and then look at the premisses in order, as needed).

P	Q	R	(P ∨ Q)	(P ⊃ R)	(Q ⊃ R)	R
T	T	T				T
T	T	F	T	F		F
T	F	T				T
T	F	F	T	F		F
F	T	T				T
F	T	F	T	T	F	F
F	F	T				T
F	F	F	F			F

So much, then, for our first five example arguments **A** to **E**: they are both intuitively valid and valid by the truth-table test when translated into suitable PLC languages. Let's turn, then, to the argument

F If Jack bet on Eclipse, he lost his money. Jack lost his money. So Jack bet on Eclipse.

As we noted before, the inference here of the form *If A then C, C, So A* is an instance of a gross fallacy (traditionally called *affirming the consequent*).

Take a translation of **F** using the material conditional, for instance

F′ $(\mathsf{P} \supset \mathsf{Q}),\ \mathsf{Q}\ \therefore\ \mathsf{P}$

Then this inference is tautologically invalid. Just consider the following trivial truth-table:

P	Q	(P ⊃ Q)	Q	P
T	T	T	T	T
T	F	F	F	T
F	T	T	T	F
F	F	T	F	F

And now note that the bad line here evidently corresponds to a possible state of affairs – i.e. Jack's not betting on Eclipse yet losing his money all the same. The possibility of that situation confirms that the original argument is plain invalid. So the tautological invalidity of **F′** directly reveals the invalidity of **F**.

Skipping the similar next case, now consider:

H If Jack bet on Eclipse, he lost his money. Jack did not bet on Eclipse. So Jack did not lose his money.

This inference of the form *If A then C, Not-A, So not-C* commits another horrible fallacy (traditionally called *denying the antecedent*). Translating using the material conditional and running a truth-table test confirms this.

Finally note example **I**. Rendered into a PLC language, this goes into the tautologically invalid

I′ $(\mathsf{P} \supset \mathsf{Q}) \therefore (\neg\mathsf{P} \supset \neg\mathsf{Q})$.

Be very careful, then, to distinguish this sort of fallacious inference from the valid inference from a conditional to its contrapositive as in **C′**!

15.5 'Only if' (and the biconditional)

(a) So far, we have only discussed 'if' conditionals, i.e. conditionals of the form

(1) If *A* then *C*.

(with or without the 'then'). We can also write 'if' conditionals in the form

(1′) *C*, if *A*.

But we will take this, without further ado, to be another merely stylistic variant.

However, there is another basic kind of conditional which needs to be discussed, namely 'only if' conditionals. Consider the pair

(2) *A* only if *C*.
(2′) Only if *C*, *A* (or: only if *C*, then *A*).

Again, we can take these versions to be mere stylistic variants of each other. So the interesting question concerns the relation between (1) *if A then C* and (2) *A only if C*.

Consider the following line of thought:

- Given *if A then C*, this means that *A*'s truth implies *C*, so we will only have *A* if *C* obtains as well – i.e. *A only if C* holds.
- Conversely, given *A only if C*, that means that if *A* is true, then, willynilly, we get *C* as well, i.e. *if A then C*.

So plausibly, (1) and (2) imply each other. In other words,

In many cases, propositions of the form *if A then C* and *A only if C* are equivalent.

To take an example, consider the following pair:

(3) If Einstein's theory is right, space-time is curved..
(4) Einstein's theory is right only if space-time is curved.

These two do indeed seem equivalent. (As often, however, ordinary language has its quirks: can you think of cases where a proposition of the form *if A then C* seems not to be straightforwardly equivalent to the corresponding *A only if C*?)

(b) Concentrate on cases where *A only if C* and *if A then C* do come to the same. Then *both* will have to be rendered into a formal language with only truth-functional connectives in the same way, by using the material conditional. So in particular, *A only if C* will be rendered by the corresponding $(\alpha \supset \gamma)$, where α is the formal translation of *A*, and γ of *C*.

Be careful: translating 'only if' conditionals does seem to get many beginners into a tangle! The basic rule is: first put an 'only if' conditional into the form *A only if C* and then replace the 'only if' with the horseshoe '⊃', to get the corresponding $(\alpha \supset \gamma)$.

The conditional *A if C* is of course translated the other way about, since that is trivially equivalent to *if C then A* and so is to be rendered by the corresponding $(\gamma \supset \alpha)$.

(c) For a simple example, consider the following argument:

J Jack came to the party only if Jo did. Also Jill only came to the party if Jo did. Hence Jo came to the party if either Jack or Jill did.

(Note how the 'only' can get separated in ordinary English from its associated 'if'.)

That argument is evidently valid. What about its correlate in a PLC language? Using 'P' to render the claim Jack came to the party, 'Q' for the claim Jill came to the party, and 'R' for the claim that Jo came to the party, we can transcribe the argument like this:

J′ $(\mathsf{P} \supset \mathsf{R}),\ (\mathsf{Q} \supset \mathsf{R})\ \therefore\ ((\mathsf{P} \lor \mathsf{Q}) \supset \mathsf{R})$

A truth-table shows that this formal version is tautologically valid, as we would hope:

P	Q	R	(P ⊃ R)	(Q ⊃ R)	((P ∨ Q) ⊃ R)
T	T	T			T
T	T	F	F		F
T	F	T			T
T	F	F	F		F
F	T	T			T
F	T	F	T	F	F
F	F	T			T
F	F	F			T

(d) Finally, what about the two-way conditional or *biconditional*, i.e. *A if and only if C*? Plainly, the two-way conditional says more than either of the separate one-way conditionals *A if C* and *A only if C*. Which doesn't stop careless authors sometimes writing *A if C* when they really mean *A if and only if C*. And it doesn't stop careless readers rather often understanding *A if C* as *A if and only if C* even when only the one-way conditional is intended. So take care!

Informally, logicians and logically-inclined philosophers often write the biconditional as simply *A iff C*.

The biconditional *A if and only if C* is the conjunction of *A if C* and *A only if C*. The second conjunct we have just suggested should be translated into our formal language by the relevant $(\alpha \supset \gamma)$; the first conjunct, as we also noted, is to be rendered by $(\gamma \supset \alpha)$. Hence the whole biconditional can be rendered by $((\gamma \supset \alpha) \land (\alpha \supset \gamma))$.

Now, that is a just a little cumbersome. And so, although it certainly isn't necessary, it is quite common to introduce a new symbol '≡' for the two-way truth-functional connective, so that $\alpha \equiv \gamma$ is equivalent to $((\gamma \supset \alpha) \land (\alpha \supset \gamma))$. This is the so-called *material biconditional*; and a simple calculation shows that its truth-table will be as displayed.

α	γ	$(\alpha \equiv \gamma)$
T	T	T
T	F	F
F	T	F
F	F	T

You really should know about this truth-function and its most basic properties, and be able to recognize symbols for it when you see them elsewhere. However, to reduce clutter, we won't officially deploy '≡' in the main text of this book; we will relegate its use to a series of exploratory examples in various of the end-of-chapter Exercises.

15.6 Adopting the material conditional

(a) Examples **A** to **J** show that rendering some simple arguments involving conditionals into a PLC language and then using the truth-table test can yield just the right verdicts about the validity or invalidity of the original arguments. And generalizing from those examples, we see that – as intended – the material conditional indeed satisfies the inferential principles (MP), (MT), (CF) and (NC) from §15.2 (with '⊃' in place of 'if').

The material conditional also satisfies the Conditional Proof principle (CP). Suppose α and β entail γ. Then there is no way of making α and β true and γ false. But that means that there is no way of making α true and $\beta \supset \gamma$ false. So α entails $\beta \supset \gamma$.

In summary, then, the logical behaviour of '⊃' in PLC languages is parallel to that of the vernacular 'if', in at least *some* cases, in *some* key respects.

(b) The obvious next question is: just how close is the relation between ordinary-language conditionals and the material conditional? Does 'if' more generally stand to '⊃' roughly as 'and' stands to '∧'?

This turns out to be a much disputed issue; so let's shelve it until the next chapter. For anyway, even if it turns out that '⊃' is not a close *analysis* of ordinary 'if', we can still adopt it to serve as an easily managed, elegantly simple, *substitute* in formal languages for the messier vernacular conditional. We hereby do so!

In fact, this is exactly how the material conditional was introduced by Frege, the founding father of modern logic, in his *Begriffsschrift*. Frege's aim was to construct a formal language in which mathematical reasoning, in particular, could be represented entirely clearly and unambiguously – and for him, such clarity requires departing from "the peculiarities of ordinary language" as he calls them, while capturing some essential logical content. Choice of notation apart, the central parts of Frege's formal apparatus including the material conditional, together with his basic logical principles (bar one), turn out to be exactly what mathematicians need.

That's why modern mathematicians – who do widely use logical notation for clarificatory purposes – often introduce the material conditional in text books, and then cheerfully say (in a Fregean spirit) that this tidy notion is what *they* are officially going to mean by 'if'. It serves them perfectly in formally regimenting their theories (e.g. in giving axioms for formal arithmetic or set theory). And the rules that the material conditional obeys – like (MP) and (CP) – are just the rules that mathematicians already use in reasoning with conditionals. Much more about this in due course.

This gives us, then, more than enough reason to continue exploring the material conditional. For we will want to investigate what happens when we adopt '⊃' as a 'clean' substitute for the conditional in our formal languages, one which serves the central purposes for which we want conditionals, at least in contexts such as mathematics.

(c) However, if the material conditional perhaps doesn't always capture the core logical content of vernacular conditionals, if we are thinking of it more as a replacement apt for use in some contexts, then we'll have to be correspondingly cautious in using the sort of two-stage procedure for evaluating ordinary-language inferences that we illustrated in §15.4. The idea, recall, is that (1) we render a vernacular argument involving conditionals into an appropriate PLC language, and then (2) we evaluate the corresponding PLC argument by running the truth-table test. But we now have to ask ourselves, case-by-case, whether the rendition of the vernacular argument into PLC at stage (1) really hits off enough of the conditional content of the premisses and conclusion for the verdict on the PLC version to carry back to a verdict on the original argument. That, to be honest, can be a bit tricky! The motto has to be: proceed with caution. (For more, see the Exercises for this and the next chapter.)

15.7 PLC syntax and semantics, officially

The material conditional – like any truth-function – can be expressed using the existing resources of a PL language: both $(\neg\alpha \lor \gamma)$ and $\neg(\alpha \land \neg\gamma)$ have the same truth-table as $(\neg\alpha \supset \gamma)$. So we certainly do not *need* to add a new symbol for the conditional. However, there are trade-offs between various kinds of simplicity here. If we add a symbol for the conditional, we thereby get much nicer-looking translations: the cost will be that – when we develop formal apparatus for dealing with proofs involving these connectives – things get a bit more complicated (e.g. we will need additional inference rules to cover the new connective). On balance, however, it is worth making the trade: so we will continue to add the new connective for the material conditional.

(a) As we said before, we will call the languages which have the four truth-functional connectives built in, including the conditional, PLC languages. Here, then, is a summary of the official syntax of such languages (compare §8.1):

The alphabet of a PLC language is that of a PL language with '⊃' added.

The definition of atomic wffs remains the same as for PL languages.

The rules for forming PLC wffs are now

W1 Any atomic wff counts as a wff.
W2 If A and B are wffs, so is $(A \land B)$.

W3 If A and B are wffs, so is $(A \lor B)$.
W4 If A and B are wffs, so is $(A \supset B)$.
W5 If A is a wff, so is $\neg A$.
W6 Nothing else is a wff.

The old results about the uniqueness of construction trees for wffs of such languages, the definitions of the ideas of main connectives etc., will all still apply.

(b) One comment on the choice of symbolism. Most logicians now use '→' as their preferred formal symbol for the truth-functional conditional. Indeed we will eventually do so. But we prefer to stick for the moment to the old-school symbol for the material conditional truth-function. Then – spoiler alert! – we will later introduce '→' initially to play a different role; it will be a not entirely trivial *result* that in fact '→' and '⊃' come to the same. (By the way, if you prefer to use '→' rather than '⊃' for the one-way conditional, then you will want to use '↔' rather than '≡' for the biconditional.)

(c) As for the semantics of PLC, it is a highly contentious question exactly how close the core meaning of (some) ordinary 'if's really is to the material conditional. To repeat, that's the topic of the next chapter.

However, as far as questions of *valuations* are concerned, those are pleasingly simple. We simply extend the story about valuations for PL languages to cover PLC languages in exactly the way you would expect. We can sum up how '⊃' affects valuations by noting again the truth-table for this new connective, so that we get

The truth tables for the PLC *connectives:*

α	β	$(\alpha \land \beta)$	$(\alpha \lor \beta)$	$(\alpha \supset \beta)$
T	T	T	T	T
T	F	F	T	F
F	T	F	T	T
F	F	F	F	T

α	$\neg\alpha$
T	F
F	T

Tautological validity for arguments in PLC languages is then defined as for arguments in PL languages.

15.8 '⊃' versus '⊨' and '∴'

(a) Recall that in §14.3 we drew a very sharp distinction between the object-language inference marker '∴' and the metalinguistic sign '⊨' for the entailment relation. We now have another sign which we must also distinguish very clearly from both of these, namely the object-language conditional '⊃'. Let's sort things out carefully.

As we stressed before

(1) $(\mathsf{P} \land \mathsf{Q}) \therefore \mathsf{Q}$

is a mini-argument expressed in a PL or PLC language. Someone who asserts (1), assuming that some content has been given to the atoms, is asserting the premiss, asserting the conclusion, and indicating that the second assertion is derived from the first one. Compare the English *Jack is a physicist and Jill is a logician. Hence Jill is a logician*.

By contrast

(2) $(\mathsf{P} \land \mathsf{Q}) \supset \mathsf{Q}$

is a single proposition in a PLC language. Someone who asserts (2) asserts the whole material conditional in that language (whose content will depend on the content of the atoms) but asserts neither the antecedent '$(\mathsf{P} \land \mathsf{Q})$' nor the consequent '$\mathsf{Q}$'. Compare the English *If Jack is a physicist and Jill is a logician, then Jill is a logician*.

By contrast again

(3) $(\mathsf{P} \land \mathsf{Q}) \vDash \mathsf{Q}$

is another single proposition, but this time in our extended English metalanguage, and the PLC wffs here are not used but mentioned. For (3) just abbreviates the English

(3′) '$(\mathsf{P} \land \mathsf{Q})$' tautologically entails 'Q'

Compare *'If Jack is a physicist and Jill is a logician' logically entails 'Jill is a logician'*.

(b) Note too the difference between the grammatical category of the PLC connective '$\supset$' and the English-augmenting sign '$\vDash$'. The horseshoe sign is a *sentential connective*, which needs to be combined with two sentences (wffs) to make a compound sentence. The double turnstile expresses a *relation* between one or more wffs (premisses) and another wff (a conclusion); so the turnstile needs to be combined with expressions referring to wffs to make a sentence.

There's a bad practice of blurring distinctions by talking, not of the 'material conditional', but of 'material implication', and then reading something of the form '$(\alpha \supset \beta)$' as *α materially implies β*. If we read '$\alpha \vDash \beta$' as *α logically implies β*, this makes it sound as if '$(\alpha \supset \beta)$' is the same sort of claim as '$\alpha \vDash \beta$', only rather weaker. But, as we have seen, they are claims of a quite different grammatical kind.

(c) The following important result, however, does link 'logical implication' in the sense of entailment to 'material implication' in the sense of the material conditional:

> Let β, γ be wffs from a PLC language: then $\beta \vDash \gamma$ if and only if $\vDash (\beta \supset \gamma)$.

Or in plain words: the one premiss argument $\beta \therefore \gamma$ is tautologically valid if and only if (iff) the corresponding material conditional $(\beta \supset \gamma)$ is a tautology. Why so?

> By definition, $\beta \vDash \gamma$ iff there is no assignment of values to the atoms which appear in β and γ which makes β true and γ false, i.e. which makes $(\beta \supset \gamma)$ false. Hence $\beta \vDash \gamma$ iff every assignment of values to the relevant atoms makes $(\beta \supset \gamma)$ true, i.e. $\vDash (\beta \supset \gamma)$.

By similar reasoning, we can show that for any side-premiss α,

$\alpha, \beta \vDash \gamma$ if and only if $\alpha \vDash (\beta \supset \gamma)$,

which is, as wanted, a formal analogue of the principle (CP) that we governs (some) ordinary-language conditionals.

15.9 Summary

- The only truth-functional connective that is a candidate for translating the conditional of ordinary discourse is the so-called material conditional. If we represent this connective by '$\supset$', then $(\alpha \supset \gamma)$ is false just when α is true and γ is false, and $(\alpha \supset \gamma)$ is true otherwise.
- Since PL languages are expressively complete, we don't need to add new a symbol to such languages to express the material conditional. But it is convenient to do so. The extended languages, with the symbol '$\supset$' added in the obvious way, will be termed PLC languages.
- At least a range of ordinary-language arguments using conditionals can be rendered into PLC languages using '$\supset$' while preserving their intuitive validity or invalidity. And indeed, the material conditional can serve in some cases as a simple, well-understood, surrogate or substitute for ordinary-language conditionals, e.g. for mathematicians.
- However, the exact relation between the general run of ordinary 'if's and '$\supset$' is highly contentious.
- Many conditionals for the form *A only if C* are equivalent to the corresponding conditional of the form *if A then C*; so both sorts of conditional can be rendered into PLC languages using '$\supset$' equally well (or equally badly!).
- It is important to sharply distinguish '$\supset$' from both '$\therefore$' and '$\vDash$'.

Exercises 15

(To be added)

16 More on conditionals

In the previous chapter, we adopted the material conditional as a surrogate for ordinary conditionals in our formal languages. The question remains: how good a surrogate is it? In *some* key respects, in at least *some* cases, as we have seen, it seems to work very well. But what's the general story? It gets complicated. (The complications, however, won't affect the later development of our logical systems in this book, so don't get bogged down by them! – by all means, just skim through this chapter on a first reading.)

16.1 Types of conditional

(a) We are discussing the biodynamics of kangaroos (as one does). You say:

(1) If kangaroos had no tails, they would topple over.

How do we decide whether you are right? By imagining a possible world very like our actual world, with the same physical laws, and where kangaroos are built much the same except for the lack of tails. We then work out whether the poor beasts in such a possible world would be unbalanced and fall on their noses.

In short, we have to consider not the actual situation (where of course kangaroos *do* have tails), but other possible ways things might have been.

Let's call a conditional that invites this kind of evaluation (i.e. evaluation by thinking not about the world as it is but about other ways things might have been) a *possible-world conditional*.

(b) Here's a memorable pair of examples:

(2) If Oswald didn't shoot Kennedy in Dallas, someone else did.
(3) If Oswald hadn't shot Kennedy in Dallas, someone else would have.

Let's assume that we agree that, in the actual world, Kennedy was definitely shot in Dallas, and we also agree that Oswald did it, acting alone. What should we then think about (2) and (3)?

Evidently, since in the actual world Kennedy was shot, someone must have done it. Hence, if not Oswald, someone else. So we'll take (2) to be true.

But to decide whether (3) is true, we have to consider a non-actual possible world, a world like this one except that Oswald missed and Kennedy wasn't killed. Keeping things as similar as we can to the actual world (as we believe it to be), Oswald would still have been acting alone. There would still be no conspiracy, no back-up marksmen. In such a possible situation, Kennedy would have left Dallas unscathed. Hence, we'll take (3) to be false.

Since they take different truth-values, (2) and (3) must have different contents. Which reinforces the intuition that a 'would have' possible-world conditional like (3) means something different from a simple 'did' conditional like (2).

(c) Now, the logic of possible-world conditionals is beyond the scope of this book, but one thing is clear:

> Possible-world conditionals are not truth-functional, and can not be translated by the material conditional.

For a truth-functional statement has its actual truth-value fixed by the actual truth-values of its constituent sentences. By contrast, the truth-value of a possible-world conditional depends on what happens in other possible scenarios; so its truth-value can't be fixed just by the this-worldly values of the antecedent and consequent. Hence, entirely uncontroversially, the truth-functional rendition can at most be used for conditionals like (2) and not for those like (1) or (3).

(d) What is rather controversial, though, is the range of conditionals that are appropriately 'like (2)' and which are not 'like (1) or (3)'.

Conditionals grammatically like (1) and (3) are conventionally called *subjunctive* conditionals, for supposedly they are couched in the subjunctive mood (though this is often contested). By contrast, (2) – and all the conditionals in examples **A** to **J** in §15.1 above – are conventionally called *indicative* conditionals, being framed in the indicative mood. So does the grammatical subjunctive/indicative distinction mark the distinction between the possible-world conditionals which are definitely not truth-functional and the rest?

Arguably not. For compare (3) with

(4) If Oswald doesn't shoot Kennedy in Dallas, someone else will.

Someone who asserts (4) before the event – believing, perhaps, that a well-planned conspiracy with back-ups is afoot – will after the event assent to (3). Likewise, someone who denies (4) before the event – believing Oswald to be a loner – will deny (3) after the event. Either way, it seems that exactly the *same* type of possible world assessment is relevant for (3) and (4): we need to consider what occurs in a world as like this one as possible given that Oswald doesn't shoot Kennedy there. Plausibly, then, we might well want to treat (3) and (4) as belonging to the *same* general class of possible-world conditionals. But, in traditional grammatical terms, (4) is indicative, not subjunctive. So the supposed distinction between subjunctive and indicative conditionals divides (3) from (4) in an arguably unnatural manner.

(e) Here is a related point. Someone who asserts (2) leaves it open – as far as that claim goes – whether Oswald shot Kennedy; but someone who asserts (3) thereby implies that the antecedent is actually false, and that Oswald didn't miss. For that reason, conditionals like (3) are often called *counterfactual* conditionals. But (4) is not a counterfactual in this sense: someone who asserts (4) again leaves it open whether Oswald will miss. So the counterfactual/non-counterfactual distinction again unnaturally divides (3) from (4).

But we can't pursue such questions about the classification of conditionals any further here. We just note that it isn't obvious exactly where and how to draw the distinction between possible-world conditionals like (1) and (3) which are definitely not apt for a truth-functional treatment, and conditionals like (2) which are still in the running.

(f) From now on, then, let's set aside possible-world conditionals, and concentrate on the remaining 'indicative' conditionals (as we will continue to call them).

We will also set aside, for the moment, conditionals like 'If you heat an iron bar, it expands' or 'If a number is even, then its square is even' (although these are in the indicative mood). For such conditionals are in fact generalizations, equivalent to e.g. 'Take any iron bar, if you heat it, it expands' or 'Every number is such that, if it is even, then its square is even'. So we narrow our focus further, and consider just *singular*, i.e. ungeneralized, indicative conditionals. And let's ask: how well does the material conditional work in capturing the core logical content of at least *these* conditionals?

16.2 Indicative conditionals as material conditionals: for and against

We have seen how the singular indicative conditionals in arguments **A** to **J** from the last chapter can be rendered into PLC using the material conditional in a way that preserves facts about validity and invalidity. This invites an attractively simple generalization:

> (*if* = ⊃) '⊃' stands to 'if' as '∧' stands to 'and' and '∨' stands to 'or'. In other words, the material conditional truth-function captures the core meaning of at least some 'if's, namely those in singular indicative conditionals.

In this section, we first consider a more general argument *for* this position; then we give a brisk argument *against* it.

(a) Consider how an ordinary-language (singular indicative) conditional *if A then C* relates to the corresponding propositions of the forms *either not-A or C* and *it isn't the case that both A and not-C*:

- Suppose *if A then C*. So we either have *not-A*, or we have *A* and hence *C*. So *if A then C* implies *either not-A or C*.

  Conversely, suppose we are given *either not-A or C*. Then if not the first, then the second. So we can infer *if A then C*.

- The claim *if A then C* rules out having *A* true and *C* false. So *if A then C* implies *it isn't the case that both A and not-C*.

  Conversely, suppose we are given that *it isn't the case that both A and not-C*. Then we can infer that if *A* is actually true we can't have *not-C* as well: in other words *if A then C*.

These considerations imply, first, that *if A then C* is equivalent to the truth-functional *not-A or C*. And second, *if A then C* is equivalent to the truth-functional *it isn't the case that both A and not-C*.

Therefore – using α, γ to stand in for formal translations of the English clauses *A*, *C* – we can render the logical content of indicative conditionals *if A then C* by something of the form $(\neg\alpha \lor \gamma)$ and/or by $\neg(\alpha \land \neg\gamma)$. But those are, of course, both equivalent to $(\alpha \supset \gamma)$.

(b) That double-barrelled argument for (*if* = $\supset$) is brisk but rather abstract. So let's amplify one of the steps by giving a realistic illustration of what looks like an equivalence between something of the form *if A then C* and the corresponding *not-A or C* – or, what comes to the same, between *if not-B then C* and the corresponding *B or C*.

Suppose I vividly remember that Jack and Jill threw a party together last Easter, but can't recall whose birthday it was celebrating. On this basis, I believe

(1) Either Jack has a birthday in April or Jill does.

Now, you tell me – a bit hesitantly – that you think that maybe Jack has an October birthday. I agree that you might, for all I know, be right; but I will still infer from (1) that

(2) If Jack was not born in April, then Jill was.

In the context, deriving (2) is surely perfectly acceptable. And note, it doesn't commit me for a moment to supposing that there is any causal or other intrinsic connection between the facts about Jack and Jill's birth months. I just think that at least one of the propositions *Jack was born in April* and *Jill was born in April* happens to be true, and thus if not the first, then the second.

Since (2) unproblematically implies (1), these two therefore seem to be equivalent in the sense of being interdeducible. Hence perfectly acceptable conditionals like (2) are no-more-than-material in what they commit us to about the world.

(c) Of course, when we assert an ordinary conditional, we often *do* think that there is e.g. a causal connection between the matters mentioned in the antecedent and the consequent ('If ice was applied, the swelling went down'). But equally, when we assert a disjunction, it is quite often because we think there is some mechanism ensuring that one disjunct or the other is true ('Either the e-mail was sent straight away or you had a warning message that the e-mail is queued'). However, even if our *ground* for asserting that disjunction is our belief in an appropriate mechanism, what we actually *say* is true just so long as one or other disjunct holds. Likewise, our *ground* for asserting a conditional may be e.g. a belief in some mechanism that ensures that if the antecedent holds the consequent does too. But the 'birthdays' example shows that such a causal mechanism doesn't have to be in place for a conditional to be true.

To continue with that example, suppose my memory of a party is in fact faulty. But despite this, it turns out that while Jack was not born in April, Jill *was* born in April. My claim (2) is proved right, even if I am right by sheer luck. Winning the lottery against the odds is still winning: similarly, getting it right that *if A then C* because *A* is verified and then it turns out by a mere fluke *C* is true too is still getting it right

– no more-than-material-conditional link between antedecent and consequent need be required for truth here.

(d) So far, so favourable to (*if* = ⊃). Now for an argument that goes in the opposite direction. In fact, there are a number of cases where treating 'if's as material conditionals leads to strongly counter-intuitive verdicts on the validity of arguments involving the ordinary conditionals. We'll concentrate on a pivotal example (for more cases, see the Exercises).

Compare, then, the following three inferences (and for vividness, we borrow '⊃' to express the material conditional truth-function in English):

(1) Bacon didn't write *Hamlet*. So, either Bacon didn't write *Hamlet* or he wasn't a playwright.
(2) Bacon didn't write *Hamlet*. So, (Bacon wrote *Hamlet* ⊃ he wasn't a playwright).
(3) Bacon didn't write *Hamlet*. So, if Bacon wrote *Hamlet*, then he wasn't a playwright.

(1) is trivially valid (as is any inference of the form *not-A, so either not-A or C* for inclusive 'or'). By definition, the conclusion of (2) is just another way of writing the conclusion of (1), and hence (2) must be trivially valid too. By contrast, the inference in (3) looks absurd. (3)'s conclusion, it will be said, is quite unacceptable (being the author of *Hamlet* would be enough to make you a playwright, if anything is). So how can the apparently absurd conclusion validly follow from the sensible premiss?

Generalizing, the inference pattern

(M) *not-A*; so *not-A or C*

is unproblematically and trivially reliable (so too, of course, is its formal reflection $\neg\alpha \therefore (\alpha \supset \gamma)$ in a PLC language). On the other hand, inferences of the type

(V) *not-A*; so if *A* then *C*.

strike us as typically quite unacceptable.

Yet (*if* = ⊃) equates the content of the vernacular conditional in (V) and the material conditional truth-function expressed in (M), with the upshot that the two inference patterns should after all be exactly on a par. Which surely is entirely implausible.

16.3 Can we save (*if* = ⊃)?

We have a problem! On the one hand, the last section offers a plausible argument for the view that at least singular indicative conditionals are equivalent to material conditionals. On the other hand, we have also found what looks to be a *very* unwelcome upshot of this view. What to do?

Since (*if* = ⊃) is such an attractively simple theory, let's first see if we can save it from our counterargument. So in this section, we'll consider two possible responses to the apparent absurdity of inferences of the form (V).

(a) The first response simply asks us to think again and revise our initial antipathy to (V):

> Suppose I'm entirely confident that Bacon didn't write *Hamlet*. Then, I'll happily assert 'If Bacon wrote *Hamlet*, then I'm a Dutchman', 'If Bacon wrote *Hamlet*, then pigs can fly', 'If Bacon wrote *Hamlet*, then Donald Trump wrote *Pride and Prejudice*'. In the same spirit, we can also conclude 'If Bacon wrote *Hamlet*, then he wasn't a playright'. Why not?

But the trouble with this bald retort is that 'Dutchman' conditionals – as we might call them – strike us as jokey idiom. To say 'If Bacon wrote *Hamlet*, then I'm a Dutchman' seems too much like going through a *pretence*, playing at asserting a conditional as a dramatic way of inviting you to draw the obvious modus tollens conclusion that Bacon didn't write *Hamlet*. We surely can't safely erect a theory of conditionals on the basis of outliers like that!

(b) The second response is more conciliatory, allows that (V) not only looks absurd at first sight but continues to do so on further reflection, but then tries to reconcile this with (*if* = ⊃):

> It would, we can agree, be uncooperative of me to assert something equivalent to $A \supset C$ when I know that A is false. For I could much more informatively simply tell you, straight out, that *not-A*. That's why you can normally expect me to have grounds for asserting $A \supset C$ other than belief in the falsity of the antecedent.
>
> Similarly, then, for the equivalent ordinary-language *if A then C*. Even though it is really a material conditional, you'll reasonably expect me to have some grounds for asserting it other than belief in *not-A*. Because of that strong expectation, inferences of the form (V) will strike us as odd, even though strictly speaking correct.

Well, it is indeed agreed – we have already noted the point in §7.3(b) – that it would usually be oddly unhelpful of me to assert a bare disjunction when I know e.g. that the first disjunct is definitely true. So yes, in many contexts, it would be odd to assert the disjunction *not-A or C*, or equivalently $A \supset C$, when I know that *not-A*. However, granted all that, nothing strikes us as odd about the inference (M) above. That's because, in the context, the default presumption that I might have grounds for the disjunction *not-A or C* other than belief in *not-A* is *cancelled*, since I'm here quite frankly and explicitly offering *not-A* as my grounds!

But if there its nothing odd about (M), even allowing for what can be expected from co-operative conversationalists, and the ordinary conditional is merely the material conditional again, then there should equally be nothing odd about (V), still allowing for what can be expected from co-operative conversationalists. Yet surely the cases *are* different.

In sum, it still seems that a naive, unqualified, version of identification of (*if* = ⊃) can't explain the difference between the intuitive satisfactoriness of the inferences (M) and

(V). So that strongly suggests that there is *something* more to 'if' than a merely material conditional. But what?

16.4 Credence in conditionals

(a) Let's consider another argument against naively identifying (singular indicative) conditionals with material conditionals:

> Jill is a little late home. Thoughts about various unlikely catastrophes pop unbidden into Jack's mind, such as P: Jill has had an accident. But Jack is good at keeping his worries in check. He in fact thinks it *very* probable that *not-P*.
>
> Now, Jack knows that a disjunction is at least as likely to be true as its first disjunct (put it this way: there are at least as many ways the world might go which make $A \lor B$ true as make A true). So Jack realizes it is also very probable that *not-P or Q*, for any second disjunct at all, including e.g. Q: Jill has been trampled by a herd of elephants. In other words, Jack will give a very high degree of credence to $(P \supset Q)$ because of his very high confidence in the truth of *not-P*.
>
> However, living as they do in a small English city, Jack will think it is very improbable that, if the worst has happened and Jill *has* had an accident, then she has been tramped by a herd of elephants. In other words, Jack will give a very low degree of credence to *if P then Q*.
>
> But how can this be, on the view (*if* = ⊃)? If ordinary conditionals are no more and no less than unadorned material conditionals, Jack should give the same level of credence to *if P then Q* as to $(P \supset Q)$. Since Jack's very different levels of credence are perfectly rational, 'if's aren't '⊃'s.

This line of argument seems to hammer a final nail into the coffin of the view (*if* = ⊃).

But, on the positive side, it also brings into view something important about conditionals. Namely, the degree of acceptability of a conditional *if A then C* normally goes with the level of credence one would have in C, should it indeed turn out that A. And in particular, full confidence in a conditional goes with a preparedness to simply accept C, should it turn out that A. So an unqualified assertion of *if A then C*, expressing such confidence, is tantamount to a promise – if you persuade me that A, then I will agree that C. In short, *in asserting a conditional, I am setting myself up for a modus ponens inference*.

(b) Suppose we somehow build this last idea into our account of the conventional meaning of 'if'. Then we will be able to accommodate some of the key facts about indicative conditionals that we have noted:

(1) Consider the 'Bacon' example from §16.2. Suppose I firmly believe that Bacon did not write *Hamlet*, and that is my sole reason for accepting '(Bacon wrote *Hamlet* ⊃ he wasn't a playwright)'. Then I'm certainly not committing myself to conclude that Bacon wasn't a playwright if, against all the odds, you manage to

persuade me that Bacon *did* write *Hamlet* after all. In that case, I would instead simply take back my assent to '(Bacon wrote *Hamlet* ⊃ he wasn't a playwright)'. In other words, I am *not* setting myself up for a modus ponens inference.

This explains why, in the circumstances, I balk at the conditional 'If Bacon wrote *Hamlet* then he wasn't a playwright'. And we can similarly explain the felt peculiarity of other inferences of the kind (V).

(2) Contrast the 'birthday' example also in §16.2. I don't believe that there is any intrinsic connection between Jack and Jill's birthdays; but still, because of my (supposed) memory of the party, I believe (Jack was not born in April ⊃ Jill was born in April). And if you persuade me that Jack was definitely not born in April, then I will happily infer that Jill was. So in this case, I am setting myself up for a modus ponens inference. That's why I am correspondingly happy to endorse 'If Jack was not born in April, then Jill was'.

(3) Relatedly, ($if = \supset^+$) explains why Dutchman conditionals strike us as unserious. For I say 'If Bacon wrote *Hamlet*, then I am a Dutchman', I am quite patently *not* holding myself ready to make a modus ponens inference. I am conspicuously flouting the usual conventions governing the use of 'if'.

(c) So: how *do* we build into an account of the meaning of 'if' this idea that endorsing a conditional goes with willingness to draw the conclusion of a modus ponens? We could try a direct marriage between the truth-functional theory and the inferential-promise idea, along the following lines:

($if = \supset^+$) A (singular indicative) conditional of the form *if A then C* has the same truth-relevant content as the corresponding $A \supset C$ but the use of 'if' in addition conventionally signals that the user will be prepared to make a modus ponens inference should it turn out that A is true.

So perhaps 'if' stands to '⊃' somewhat as 'but' stands to '∧'. We suggested in §7.2 that *A but B* and *A and B* are true in the same situations, though 'but' is reserved for use when the speaker takes there to be some contrast between the truth of A and the truth of B. The idea now is that an ordinary conditional *if A then C* and the corresponding $A \supset C$ are again true in the same situations, but 'if' is conventionally reserved for use when the speaker is undertaking to endorse C should it turn out that A.

The truth-functional part of this theory can be used to explain why the positive arguments for ($if = \supset$) in §16.2 look as good as they do. The inferential-promise part of this theory explains why one's degree of willingness to endorse a conditional goes, as we want, with one's degree of willingness to endorse C should it turn out that A. However, the trouble with this new theory is that it is difficult to understand how its two halves fit really together.

Let's ask: is the supposed parallel with 'but' a good one? After all, while we might very properly balk at the use of 'but' in e.g. 'Jill is a woman but Jill is a good logician', we can agree that what this unfortunately expressed claim tells us about the world (as opposed to the speaker's attitude to women) is right – Jill *is* a woman, and Jill *is* a good logician. Contrast worried Jack, and the conditional 'If Jill has had an accident, she has been trampled by elephants'. It just doesn't seem right to say that what this ordinary-language conditional tells us about the world (the content of the associated material-conditional, according to the theory) is very probably true, and that dissent from the conditional again just comes from a side constraint on when we should choose 'if' to express that probable truth.

So how *do* we better we build into an account of the meaning of 'if' the idea that one's credence in a conditional should go with the acceptability of its conclusion should the antecedent turn out to be true?

16.5 Let's stick, all the same, to the material conditional!

And with that cliffhanger, we will leave the discussion of vernacular conditionals!

The very extensive philosophical literature on conditionals continues the discussion in a quite inconclusive way, and it certainly won't help us to pursue it any further here. But perhaps we have already said enough to make it plausible that there may be no neat and crisp story to be told about what ordinary 'if's add to '⊃' (for a few more complications, see the Exercises). Rather, the deep disagreements between philosophers of logic show at the very least that the behaviour of conditionals is complex and many-faceted and difficult to account for. You can now see why we delayed tackling the connective 'if' until after we had investigated the easier cases of 'and', 'or' and 'not'!

And all this raises the following question: should we even be *trying* to capture this complex behaviour accurately in a formal language? Changing tack, perhaps we should be welcoming with relief the material conditional as a *replacement* for the vernacular conditional, a replacement which is sufficiently conditional-like to be useful at least for some important formal purposes, while being elegantly clear, perfectly understood. Which, as we said, is exactly the line taken by mathematicians (who want a useful and easily-managed logic for regimenting their arguments) and is the line we recommended in §15.6. We explore this further in coming chapters.

Though note again that if that is our line, then the material conditional comes with a logical health warning. You can't blithely translate ordinary-language 'if's by '⊃'s expecting always to preserve the logic of arguments involving conditionals. Handle with care!

16.6 Summary

- As we saw in the previous chapter, a number of ordinary-language arguments using conditionals can be rendered using '⊃' while preserving their intuitive validity/invalidity. However this at best works only for some 'indicative' conditionals, and not for 'subjunctive/counterfactual' (or 'possible world') conditionals.
- There are serious issues about how far, even in the best cases, '⊃' can capture *all* of the meaning of everyday 'if'. Identifying the two leads to oddities like the acceptability of arguments of the form *Not-A; so if A, then C*.

- Note too that our credence in ordinary conditionals can peel apart from our credence in the corresponding material conditionals.
- However, as in the last chapter, the material conditional can still be recommended as a perfectly clear *replacement* for the ordinary language conditional, with a perfectly determinate logic, suitable for use by mathematicians and others.

Exercises 16

(To be added)