

## Interlude: ‘Truth trees’ vs ‘natural deduction’

(a) What have we done since the last Interlude?

We first explained the essential ‘divide and rule’ approach to logic. The general idea is that we sidestep having to deal with the quirks of ordinary language by rendering vernacular arguments into well-behaved artificial languages, and then we investigate the validity or otherwise of the resulting formalized arguments.

We then considered how to formalize arguments involving just ‘and’, ‘or’ and ‘not’ as their essential logical materials. This is a toy example; but discussing it has enabled us to introduce a whole raft of key logical ideas in a simple context.

We defined the *syntax* of our toy PL languages. We learnt how wffs of such languages are built up (in ways that can be displayed using construction trees). We also defined the notion of scope, the notion of a main connective, and related ideas.

Next, *semantics*. We discussed the interpretation of PL languages: we assign truth-relevant contents (senses) to atoms, and read the connectives as expressing bare conjunction, inclusive disjunction, and negation – understanding these in a way that makes the connectives truth-functional.

Given the truth-functional interpretation of the connectives, it follows that the truth-value of a wff is determined by a valuation of its atoms.

We defined a tautology as a wff true on every combinatorially possible valuation of its atoms, while a contradiction is false on every such valuation.

Along the way, we touched on issues about the distinction between variables  $A, B, \dots$  and ‘ $\alpha$ ’, ‘ $\beta$ ’,  $\dots$  added to the metalanguage and the likes of ‘P’, ‘Q’,  $\dots$  in our artificial object-languages. We also considered quotation conventions.

This got us to the position where we could take the pivotal step, and define the notion of tautological validity:

A PL argument is tautologically valid if every combinatorially possible valuation of the relevant atoms which makes the premisses true makes the conclusion true.

We discussed the relation between this notion of validity and the informal notion of logical validity that we defined at the outset. We can plausibly argue that an inference is tautologically valid just if it is valid (in the informal sense) in virtue of the meaning of the logical connectives. At any rate, the notion of tautological validity gives us a cleanly defined rational reconstruction for the notion of logical validity as applied to a certain class of inferences.

We introduced *one* way of establishing tautological validity, namely by a brute-force truth-table test. We can in this way decide, quite mechanically, which PL inferences are tautologically valid.

Finally, we considered the prospects for dealing with arguments also involving ‘if’ in the same way in which we have been dealing with PL arguments:

We met the so-called material conditional, which is the only possible conditional-like truth-function, and we added this to our formal apparatus to get PLC languages. However, the material conditional is a surrogate for the ‘if’ in (at best) some indicative conditionals, and there is a serious issue about how far it captures the meaning of ‘if’ even in these cases.

So: where do we go next?

(b) Take for example the following arguments:

- (1) Felix is a cat. All cats are scary. So, Felix is scary.
- (2) The baron hates anyone who loves Alma. Fritz loves Alma. So, the baron hates Fritz.
- (3) Some philosophers admire every logician. No philosopher admires any charlatan. So, no logician is a charlatan.

These are evidently logically valid. But if we expose no more structure than the way in which whole propositional clauses are related by connectives (if any), then each of these arguments can only be assigned the utterly unreliable inferential form  $A, B, \text{ so } C$ .

We therefore have to dig deeper inside the premisses and conclusions of these arguments. What is crucial in each case is plainly the presence of expressions of generality – *quantifiers*, as linguists call them – such as ‘all’, ‘any’, ‘some’, ‘no’. So we will need a theory of how such quantified propositions work, and how they logically interact with each other and with simpler propositions. The resulting *quantification theory* (also known as ‘first-order logic’ or ‘predicate calculus’) is, as we will come to see, vastly more powerful than propositional logic.

However – unlike classical propositional logic – we *can’t* in general use a brute-force calculation akin to a truth-table test to decide whether a quantificational inference is valid or not. We need some other approach to establishing validity here. Two options (from a number) are ‘semantic tableaux’ or *truth trees*, and *natural deduction* proofs. We can gesture to what these involve by returning to propositional logic for a moment.

(c) As we have just reminded ourselves, the basic idea of tautological validity is quite beautifully clear and straightforward. But how do we establish that a PL inference actually *is* valid in this sense?

There are various ways of doing this. So far, however, we have only met the most obvious, brute-force, method – namely, trudge through all the different valuations of the atoms occurring in the inference, and wait to see if a ‘bad line’ with true premisses and false conclusion turns up! This certainly isn’t pretty. And as examples get more complicated, applying this method will quickly become a horrendous task.

So one obvious question is: can we do better? Here’s a rather silly but very telling example to think about. Take the argument

**A**       $\neg P, \neg P', \neg P'', \neg P''', \dots, \neg P'''''' \therefore Q$

where there are *fifty* similar premisses, each the negation of a different atom. The argument is of course tautologically invalid. But suppose we set off mindlessly to construct a truth table (with all the 'P's listed on the left before the 'Q'). Then we have to write down a  $2^{51}$  line truth table. And it won't be until the very last line, which assigns F to all the 'P's and to 'Q', that we get to the lone 'bad' valuation making all the premisses true and the conclusion false. Even if we could write down a complete line of the table per second, it would take *half a million years* before we got to that decisive final line!

So how do we know – straight off, without waiting many millennia! – that argument **A** is tautologically invalid? Well, instead of working through all the combinatorially possible valuations of its atoms looking for a 'bad' line, we can approach things from the other end. In other words, we suppose that there *is* a bad valuation, i.e. one where the premisses are true and the negation of the conclusion is true too; then we 'work backwards' to see what such a valuation would have to look like. And in the case of **A**, the job is *very* easily done. We see immediately that we just need to make all the 'P's false and 'Q' false too; the premisses will then all be true and conclusion false.

Now let's generalize this idea of 'working backwards'. Suppose we want to test an inference for tautological validity. We suppose that there *is* a bad valuation, and then we systematically unravel the implications of this supposition. We will either find a bad valuation (showing the inference is invalid), or we will get embroiled in contradictions every way we turn (showing that there is no bad valuation after all, hence that the inference is valid). This strategy can be developed into the truth-tree method for testing for validity – 'trees' come into the picture because we often have to set out pathways through alternative options. This method is certainly more elegant, and often much faster, than running a brute truth-table test. And, importantly, the use of truth trees can be extended in a straightforward way to establish the validity of valid quantificational arguments (though we can no longer apply it purely mechanically).

(d) Going now in a different direction, recall that, in Chapters 4 and 5, we talked about *two* informal ways of demonstrating the validity/invalidity of inferences – namely, showing validity by giving a step-by-step proof, and showing invalidity by finding a counterexample. Turning to inferences in PL/PLC languages, we have seen that looking for a counterexample to show tautological invalidity becomes looking for a bad line on a truth table (see §15.2, which also noted that, in this special case, because we can do systematic searches, searching and *failing* to find a counterexample establishes validity). We can also think of constructing a tree test as another systematic way of looking for a counterexample/bad line, by supposing there is one and working backwards to discover what it will look like if it exists.

However, we can surely also seek to construct more direct step-by-step proofs of validity inside our formalized PL/PLC languages. Consider, for example, the argument

**B**       $((P \wedge Q) \wedge R), (S \wedge P') \therefore (Q \wedge S)$

We can run up a thirty-two line truth table to validate this little argument (trying to find a bad line and failing). We can alternatively do a tree test, once we know how. But it is *much* more natural to reflect that, quite trivially, the following two logical laws hold:

- ( $\wedge$ E) From  $(\alpha \wedge \beta)$  we can infer  $\alpha$  and we can infer  $\beta$  (we can, as it were, eliminate a conjunction and derive either conjunct).
- ( $\wedge$ I) From  $\alpha$  and  $\beta$  we can infer  $(\alpha \wedge \beta)$  (we can, as it were, introduce a conjunction by deriving it from its conjuncts).

We can then use these two rules to construct a little proof as follows:

<b>B'</b>	(1) $((P \wedge Q) \wedge R)$	(premiss)
	(2) $(S \wedge P')$	(premiss)
	(3) $(P \wedge Q)$	(from 1 by $\wedge$ E)
	(4) $Q$	(from 3 by $\wedge$ E)
	(5) $S$	(from 2 by $\wedge$ E)
	(6) $(Q \wedge S)$	(from 4, 5 by $\wedge$ I)

Each step on the way is a valid tautological entailment, hence so too is the leap from the initial premisses of **B** to the target final conclusion. Which is what we wanted to show.

We can now aim to put together more rules of inference like ( $\wedge$ I) and ( $\wedge$ E) to govern the other connectives too, and then explain how to construct natural-looking step-by-step PL/PLC deductions using these rules. In a phrase, we can develop a *natural deduction* proof system for propositional logic. Again, this can be extended in a smooth way to deal with quantificational arguments.

(e) In summary, then, it turns out that – when we move on from propositional logic and start to tackle the crucial case of arguments which turn on the use of quantifiers – we can't in general use a brute-force calculation to determine the validity of an inference. We *can*, however, use truth trees to establish the validity of quantifier inferences. Equally, we can use natural deduction proofs.

Truth trees for quantifier arguments are relatively easy to understand and easy to use. Natural deduction is, well, rather natural! There is much to be said for knowing about both approaches at a fairly early stage in your logical education. However, there is equally much to be said for not overloading a first logic course or text. So which shall we focus on? Truth trees or natural deduction proofs for quantifier arguments?

It is no good pretending that the decision is an easy one. And even if we decide to take the second option, that's not the end of our problems, as there are significantly different flavours of natural deduction, and people disagree (vigorously) about which one to introduce to beginners! We have to choose. For better or worse, the main focus of this book will be on truth trees. And it is easier to understand truth trees for quantifier arguments if we first meet trees in the tame and now familiar context of propositional logic. Which explains why our next three chapters are about trees for PL/PLC inferences (though you are welcome to take chapters in a different order, skip forward to the next Interlude and only return to propositional trees on a need-to-know basis, after you have read Chapter 27).

In the final chapters of the book, however, we *will* give just a taster of natural deduction – enough to introduce a number of key ideas you should know about for philosophical purposes, and to provide a route into some more advanced formal logic texts.

## 18 Introducing PL/PLC truth trees

We now know what it is for PL and PLC inferences to be tautologically valid, and we have seen how tautological validity relates to logical validity. We have also met *one* way of determining whether an inference is tautologically valid, namely by a brute-force truth-table test. We next explore another way of testing for tautological validity – one that is more elegant, usually faster, and considerably less boring!

In this chapter, then, we introduce *truth trees* by talking slowly through a number of examples. We streamline this new method of testing in the next chapter, though still proceeding relatively informally. We give an official summary account in Chapter 20.

### 18.1 Three very quick preliminaries

- (1) In this chapter and the following two, take wffs in all our examples to belong to some interpreted PL/PLC language or other. And take ‘valid’/‘invalid’ always to mean *tautologically* valid/invalid.
- (2) It will be useful to have a word for those wffs which are either atoms or negations of atoms. The conventional but slightly odd term is ‘literal’. Let’s use instead the more memorable ‘simple’. Other wffs are therefore ‘complex’.
- (3) Most significantly, in this chapter we will usually prefer to say of a wff, not that it is *false* on a given valuation of its atoms, but rather that its *negation is true*. For example, instead of saying that a (tautologically) invalid inference is one where there’s a valuation which makes the premisses true and conclusion false, we will now prefer to say that *an inference is (tautologically) invalid just if there is a relevant valuation which makes the premisses and the negation of the conclusion all true*.

Obviously nothing can turn on this alternative way of putting things; in our two-valued logical framework, a wff is false if and only if its negation is true. But there is a reason for this apparently quirky preference for writing things of the form  $\neg\alpha := T$  rather than  $\alpha := F$ . We reveal all in the next chapter.

### 18.2 ‘Working backwards’

- (a) In the preceding Interlude, we considered a daft inference with fifty premisses – all negated atoms, all different – and an unrelated conclusion:

$$\text{A} \quad \neg P, \neg P', \neg P'', \neg P''', \dots, \neg P'''''' \therefore Q$$

As we noted, it would be quite crazy to try to write down a full truth table to test the status of this argument (which would take millennia). It is enough to remark that **A** is invalid if and only if there is a valuation of the fifty-one atoms which makes the premisses and negated conclusion all true – i.e. a valuation such that

$$\neg P := T, \neg P' := T, \neg P'' := T, \neg P''' := T, \dots, \neg P'''\dots'' := T, \neg Q := T.$$

Then we merely have to note that we can immediately read off just such a valuation, namely

$$P := F, P' := F, P'' := F, P''' := F, \dots, P'''\dots'' := F, Q := F.$$

Which settles it. **A** is invalid.

(b) Take next the argument

$$\mathbf{B} \quad (P \wedge \neg Q), (R \wedge \neg S) \therefore (Q \vee S).$$

This too is obviously invalid, as we can confirm with a sixteen-line truth table. But we can also argue as follows, this time laying out our reasoning step by step.

The inference **B** is invalid if and only if there is a valuation of the four atoms which makes the premisses true and the negation of the conclusion true too, so

- (1)  $(P \wedge \neg Q) := T$
- (2)  $(R \wedge \neg S) := T$
- (3)  $\neg(Q \vee S) := T$

This time, we can't instantly read off a valuation which makes (1) to (3) hold. But note that by the truth table for conjunction, (1) holds if and only if these claims hold too:

- (4)  $P := T$
- (5)  $\neg Q := T$

Similarly, (2) holds if and only if these claims do too:

- (6)  $R := T$
- (7)  $\neg S := T$

Finally, (3) holds if and only if the plain disjunction ' $(Q \vee S)$ ' is false, which requires that both disjuncts are false. Avoiding talk of falsehood, this means that the negation of each disjunct has to be true:

- (8)  $\neg Q := T$
- (9)  $\neg S := T$

So, putting everything together, we can ensure that (1) to (3) hold – i.e. we can make the premisses and negated conclusion of **B** all true – by taking the valuation where  $P := T$ ,  $\neg Q := T$  (i.e.  $Q := F$ ),  $R := T$ , and  $\neg S := T$  (i.e.  $S := F$ ). Hence **B** is indeed invalid.

(c) For a contrasting third example, consider

$$\mathbf{C} \quad (P \wedge \neg Q) \therefore \neg(Q \wedge R).$$

Again, the inference is invalid just if there is a valuation of the three atoms here which makes the premiss and negated conclusion both true, i.e. which satisfies

## §18.2 'Working backwards'

169

- (1)  $(P \wedge \neg Q) := T$   
 (2)  $\neg\neg(Q \wedge R) := T$

For (1) to obtain, we need

- (3)  $P := T$   
 (4)  $\neg Q := T$

By the truth table for negation applied twice, for (2) to obtain, we need

- (5)  $(Q \wedge R) := T$

and hence we need both

- (6)  $Q := T$   
 (7)  $R := T$

But hold on! We have now hit a contradiction. *No valuation can make  $\neg Q$  and  $Q$  simultaneously true together.* So there can not possibly be a valuation satisfying both (1) and (2) and making the premiss of **C** true and conclusion false.

Which shows that **C** is valid.

(d) In these three examples we have been – as we put it before – *working backwards*.

On the brute-force truth-table method, in order to decide whether a given argument is tautologically valid, we start by laying out all the combinatorially possible valuations for the relevant atoms. Then we plod forwards through each valuation, calculating the truth values of the premisses and conclusion as we go, looking for a 'bad' line. But now – as already sketched in the preceding Interlude – we are starting on the same decision task from the other end. We suppose that there *is* a bad valuation, i.e. one where the premisses are true and the negation of the conclusion is true too. Then we try to unravel the implications of this supposition, working backwards towards uncovering a valuation of the relevant atoms that is indeed bad.

If we *can* uncover a bad valuation, then obviously the argument being tested is *not* valid. On the other hand, if we get inextricably entangled in contradiction as we work backwards, then (by an informal *reductio ad absurdum* inference) that shows that there is no bad valuation after all, and the argument in question is *valid*.

This is, in headline terms, the appealing strategy that we are exploring in this chapter. Though as we will see in the next section, things don't in general go quite as simply as our first easy examples might suggest.

(e) Before moving on, one quick comment about our example **C**. Note that at line (2) we gave the negation of the conclusion ' $\neg(Q \wedge R)$ ' as ' $\neg\neg(Q \wedge R)$ ' with its extra prefixed negation sign. Then at line (5) we stripped away the double negation again by using the obvious rule that tells us that  $\neg\neg\alpha := T$  holds just when  $\alpha := T$ . Why go through this little two-step dance? Why not immediately jump to treating the negation of something of the form  $\neg\alpha$  as plain  $\alpha$ ?

Because, in more advanced work (beyond the scope of this book), you will eventually encounter important logical systems with a non-standard understanding of how negation works. In particular, you will meet so-called intuitionistic logic where  $\neg\neg\alpha$  is in general not equivalent to  $\alpha$  (we touch on this in §??). It is therefore *not* a good idea to get

into the habit of simply ignoring double negations, even if that is permitted in standard two-valued logic.

For us, then, negating a PL or PLC wff is *always* a matter of *adding* a prefixed negation sign, as already noted in §9.1(d).

### 18.3 Why trees?

(a) Here is an obviously invalid little argument:

$$\mathbf{D} \quad (P \vee Q) \therefore P.$$

What happens if we try working backwards to expose a valuation which makes the premiss true and negated conclusion true too?

We start, as before, with the assumption that there *is* a bad valuation, one which makes the following hold:

$$\begin{array}{ll} (1) & (P \vee Q) := T \\ (2) & \neg P := T \end{array}$$

Given the truth table for disjunction, (1) tells us that either (3a)  $P := T$  or (3b)  $Q := T$  or perhaps both – and either alternative suffices to verify the disjunction. *But (1) doesn't tell us which alternative holds.* We will have to consider the two cases in turn.

In fact, the alternative (3a) where  $P := T$  need not detail us long, for it is immediately inconsistent with (2), i.e. with  $\neg P := T$ .

The alternative (3b) where  $Q := T$  fares better; there is no contradiction to be found this time. The resulting valuation on which  $\neg P := T$  (i.e.  $P := F$ ),  $Q := T$  makes the premiss of **D** true and the negated conclusion true too. As we knew all along, then, the inference **D** is indeed invalid.

(b) We want a perspicuous way of setting out reasoning like this when our line of argument forks, and we have to explore alternatives. The obvious device is to draw forking branches in our chains of reasoning. We then get (upside down, trunk-at-the-top) *truth trees*, and can display our reasoning about **D** like this:

$$\begin{array}{lll} (1) & (P \vee Q) := T & \text{(premiss assumed true)} \\ (2) & \neg P := T & \text{(negated conclusion assumed true)} \\ & \swarrow \quad \searrow & \\ (3) & P := T \quad Q := T & \text{(alternatives from 1)} \\ & * & \text{(contradiction on left branch!)} \end{array}$$

This mode of presentation should be self-explanatory. The main novelty compared with our informal reasoning a moment ago is that, instead of considering options sequentially, we consider them side-by-side; we split our reasoning into left and right branches as we consider the alternative ways of making (1) come out true.

The other novelty is that we now begin to use ‘\*’ as an *absurdity sign*, signalling that we have hit a contradiction on a path down the truth tree from the trunk-at-the-top. In other words, ‘\*’ at the end of a path indicates that it features a pair of wffs  $\alpha$  and  $\neg\alpha$  which are both assigned T (note that the wff  $\alpha$  here doesn't have to be an atom).

§18.4 Two more examples

171

Our miniature downward branching tree for argument **D**, then, has two paths through it running from top to bottom, as we explore alternative potential ways of making true the initial claims at the top of its trunk. The left-hand path is, as we will say, *closed* with an absurdity marker; but the other path remains *open*. And from the open path we can read off a perfectly consistent assignment of values to the relevant atoms which satisfies (1) and (2), making **D** invalid.

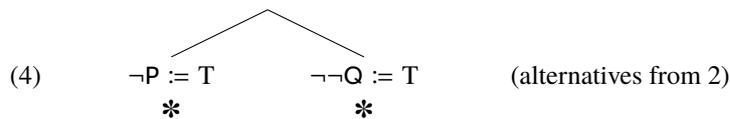
(c) Take next the following argument:

$$\mathbf{E} \quad P, \neg(P \wedge \neg Q) \therefore Q.$$

This is invalid if and only if we can consistently make the premisses and negated conclusion all true, as follows:

- (1)  $P := T$  (premiss assumed true)
- (2)  $\neg(P \wedge \neg Q) := T$  (premiss assumed true)
- (3)  $\neg Q := T$  (negated conclusion assumed true)

Now, a conjunction is false when at least one conjunct is false. Equivalently, a negated conjunction is true just when at least one negated conjunct is true. In other words, given  $\neg(\alpha \wedge \beta) := T$  we must have at least one of  $\neg\alpha := T$  or  $\neg\beta := T$ . This is the principle we can apply in moving from (2) to the alternatives at (4):



This time *both* alternatives immediately produce contradictory assignments of truth values ( $P := T$  and  $\neg P := T$  on the left-branching path from top to bottom;  $\neg Q := T$  and  $\neg\neg Q := T$  on the right-branching path). So both paths get closed off with the absurdity sign. Which shows that neither of the potential ways of satisfying conditions (1) to (3) is consistent. Therefore the inference **E** has to be valid.

18.4 Two more examples

(a) For our next example, take the argument

$$\mathbf{F} \quad (P \rightarrow Q), (Q \rightarrow R) \therefore (P \rightarrow R).$$

We showed that this is valid by a brute-force test in §16.4. We'll now try to get the same result by 'working backwards': so we assume the argument *isn't* valid, and aim for absurdities. We start, then, by assuming that there is a valuation satisfying

- (1)  $(P \rightarrow Q) := T$  (premiss assumed true)
- (2)  $(Q \rightarrow R) := T$  (premiss assumed true)
- (3)  $\neg(P \rightarrow R) := T$  (negated conclusion assumed true)

Now, a material conditional is true, remember, if either its antecedent is false or consequent is true. So, avoiding talk of falsity, we have the following unpacking rule:

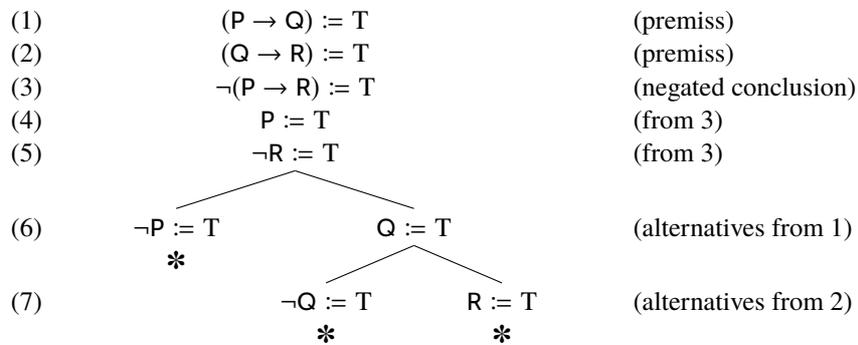
Given  $(\alpha \rightarrow \beta) := T$ , we have to consider the alternatives  $\neg\alpha := T$  or  $\beta := T$ .

And what are the implications of assuming a negated material conditional is true, i.e. that the conditional is false? This requires the conditional's antecedent to be true and conclusion false. Again avoiding talk of falsity, the corresponding unpacking rule is as follows:

Given  $\neg(\alpha \rightarrow \beta) := T$ , then we must have both  $\alpha := T$  and  $\neg\beta := T$ .

So which unpacking rule – to continue with that vivid metaphor – shall we apply first in constructing our tree? Let's start by applying the non-branching rule to (3); this delays having to consider branching possibilities and keeps our tree tidier. As we will confirm in §20.2(d), the order in which we tackle complex wffs can't make any difference to the eventual verdict on the argument which we are testing with a tree.

Following that plan, here's our completed truth tree:



First we draw out the implications of (3). Then at (6) we consider the implications of (1), exploring alternatives. The alternative on the left-hand path, though, doesn't give us a consistent way of ensuring (1) and (3) hold: so we can close off that path.

The right-hand path stays open at (6). But we haven't yet drawn out the implications of (2). This requires us to consider alternatives again, leading to our second fork in the tree. And either way, we immediately hit a contradiction with what's further up the path.

Which all goes to show that there is no consistent way of working backwards to make (1), (2) and (3) all correct. Hence, as we want, **F** is shown to be valid.

(b) An aside. Suppose we are given  $(\alpha \rightarrow \beta) := T$ . It follows that one of these *three* options must hold:  $\alpha := T, \beta := T$  or else  $\neg\alpha := T, \beta := T$  or else  $\neg\alpha := T, \neg\beta := T$  (corresponding to the three lines on the truth table where the material conditional is true). So why not adopt an unpacking rule for  $(\alpha \rightarrow \beta) := T$ , with three-way branching?

Aesthetics! It is much neater to keep our trees binary if we can, with only two branches starting at any point where we have to explore alternatives. So this is what we will do.

(c) We take just one more example for now, in order to highlight two crucial points about constructing branching trees like these. So consider:

$$\mathbf{G} \quad ((P \wedge Q) \vee R) \therefore \neg(\neg P \vee \neg R).$$

Is this valid? We look to see whether there is a valuation which makes the premiss and the negated conclusion both true, i.e. a valuation which gives us

§18.4 Two more examples

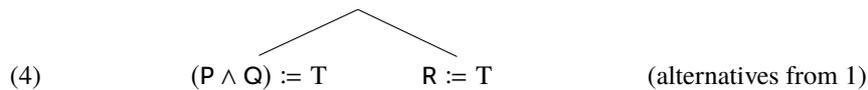
173

- (1)  $((P \wedge Q) \vee R) := T$  (premiss)
- (2)  $\neg\neg(\neg P \vee \neg R) := T$  (negated conclusion)

Trivially, (2) holds just when we have

- (3)  $(\neg P \vee \neg R) := T$  (from 2)

To unpack the implications of (1) we now need to consider alternative scenarios:



For ‘ $(P \wedge Q)$ ’ to be true, both the conjuncts have to be true. So we can continue the left branch as follows:

- (5)  $P := T$  (from 4)
- (6)  $Q := T$  (from 4)

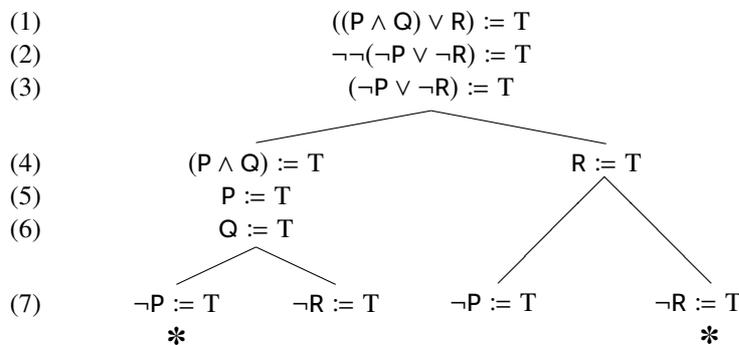
But note – and this illustrates the first obvious point we need to emphasize – these additional conditions ‘ $P := T$ ’ and ‘ $Q := T$ ’ are added to the left branch only. For it is only on the scenario described on the path going left at (4) that we are assuming ‘ $(P \wedge Q) := T$ ’ holds; and so it is only while further elaborating *this* scenario that we are committed to the implications of that assumption.

Generalizing, we have this obvious but crucial rule:

Suppose a statement  $\alpha := T$  appears on a tree, with  $\alpha$  complex. When we unpack the (perhaps alternative) implications of this statement, we add the results *only* to the open paths through the tree which actually contain the statement!

(d) Continuing with our example, we haven’t yet unpacked (3) ‘ $(\neg P \vee \neg R) := T$ ’. But note that this assumption sits both on the path going left from (1) to (6) and also on the path going right from (1) to (4). Since (3) is in play whichever route down the tree we follow, we need to go on to consider its implications either way. And to keep track of when we have fully unpacked the implications of assigning truth to some wff, the simplest policy is to add the relevant implications to all relevant paths at the same time. This is the policy we adopt.

So the next step in continuing our tree produces this:



At line (7), then, we have added the alternatives generated by (3) to *both* open paths containing (3).

Generalizing gives us the second key point we need to emphasize:

Suppose a statement  $\alpha := T$  appears on a tree, with  $\alpha$  complex. When we unpack the (perhaps alternative) implications of this statement, we will add the results to *all* the open paths through the tree which actually contain the statement.

(e) Now to conclude work on our example (continuing to spell things out in detail). Two of the four extended paths generated by unpacking (3) to get to line (7) can be closed off. On the other two paths, there aren't any complex wffs still to unpack, so there are no more lurking contradictions to be exposed – therefore *those* two paths cannot be made to close. Let's take them in turn.

The left-hand open path contains the assignment of values  $P := T$ ,  $Q := T$ ,  $\neg R := T$  (i.e.  $R := F$ ). This valuation gives us one way of making (1) and (2) correct, i.e. one way of making **G**'s premiss true and its negated conclusion true too. That's enough already to show that **G** is *not* valid.

But to continue, the right-hand open path contains the assignment  $\neg P := T$  (i.e.  $P := F$ ), and  $R := T$ . Those values also make (1) and (2) correct, however we settle the value of 'Q' (why?). Hence, depending on the choice we then make for 'Q', this gives us two more ways of making **G**'s premiss true and its negated conclusion true too.

So, in total, our tree reveals *three* different bad valuations of the three relevant atoms making the argument's premiss true and conclusion false. Check that this agrees with the verdict of a brute-force truth table.

## 18.5 Summary

Let's pause for a brisk half-time summary about the tree test for tautological validity. (We will be doing more work to elaborate and really nail down the following claims – but hopefully these headlines should already look plausible.)

We can test an inference for tautological validity by assuming that it is invalid, i.e. by assuming there is a 'bad' valuation which makes its premisses true and negated conclusion true too, and then working backwards, trying to uncover a consistent valuation which warrants that assumption.

In working backwards, we move at each stage from the assignment of truth to some complex wff to corresponding assignments of truth to some shorter subformula(s) – though we may have to consider alternatives. When we have to consider alternatives, our reasoning is naturally set out in the form of a downward branching tree.

If working backwards, perhaps through alternatives, always leads to contradictions, then we know that the original assumption that the inference is invalid has to be incorrect. In other words, if all paths through the tree setting out our working can be closed off with absurdity signs, then the argument is *valid*.

Exercises 18

175

Conversely, if at least one alternative path through the tree stays open, i.e. involves no contradiction (even when we have unpacked all the implications of assignments of truth to complex wffs on the path), then the argument is indeed *invalid*.

Exercises 18

(To be added)

## 19 Unsigned trees

In the last chapter, our ‘working backwards’ arguments to show tautological validity/invalidity are regimented as *signed* trees (we will speak of ‘trees’ even when, as in our very first examples in §18.2, we actually have only a bare trunk without any branching). The trees are ‘signed’ in the sense that the wffs that appear on the trees are all *explicitly* assigned a truth value.

However, in our particular version of signed trees, the wffs are all assigned T – that’s because, whenever we might have been tempted to say a wff was false, we instead said that its negation is true. The pay-off from this little bit of trickery is that we can now move to so-called *unsigned* truth trees, where wffs appear ‘naked’. The move couldn’t be simpler! *We just delete ‘:= T’ wherever it appears on one of our signed trees.*

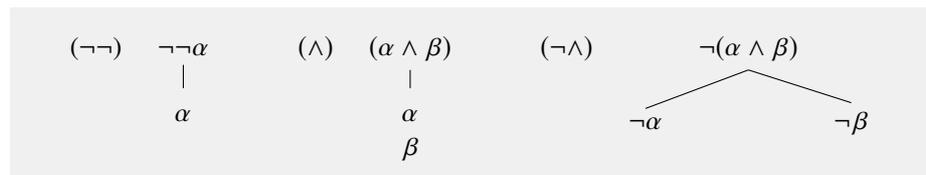
### 19.1 Rules for negation and conjunction

(a) The principles that we have been using for unpacking the implications of complex PL/PLC wffs assigned T on a signed tree can now become rules for dealing with ‘naked’ wffs on an unsigned tree. We have already met all the needed rules in their signed form. But in this chapter, we will (re-)present the rules as rules for unsigned trees, and in a more systematic order.

Sticking with the ‘unpacking’ metaphor, let’s call the wff that we are unpacking at any stage the *parent* wff, and then call the shorter wffs we add its *children*. As before, we require our unpacking rules to ensure the following:

*Parents-to-children* When we apply a non-branching rule, if the parent wff is true (on some supposed valuation) then each new added child wff is true too. When we apply a branching rule, if the parent wff is true then, on at least one branch of each added fork, the new child wff is true too.

Let’s start then with the rules for unpacking complex wffs of one of the following three shapes:  $\neg\neg\alpha$ ,  $(\alpha \wedge \beta)$ ,  $\neg(\alpha \wedge \beta)$ . In diagrammatic form, the rules are:



§19.1 Rules for negation and conjunction

We read these diagrams as you would now expect. But to laboriously spell things out:

The rule ( $\neg\neg$ ) tells us that, given a particular wff of the form  $\neg\neg\alpha$  on the tree (invisibly assigned T), we can unpack the double negation by adding  $\alpha$  (invisibly assigned T) to the tree. Add where? As explained in §18.4(c) and (d), we add a child  $\alpha$  at the foot of every open path containing the parent wff we are unpacking.

The rule ( $\wedge$ ) tells us that, given a parent wff of the form  $(\alpha \wedge \beta)$  (invisibly assigned T), we unpack the conjunction by adding both children  $\alpha$  and  $\beta$  (each invisibly assigned T), again at the foot of every open path containing the parent wff.

The rule ( $\neg\wedge$ ) tells us that given a parent wff of the form  $\neg(\alpha \wedge \beta)$ , then we unpack this by adding a fork with a branch to the child  $\neg\alpha$  and a branch to the child  $\neg\beta$ , again at the foot of every open path containing the parent wff.

Each rule obviously satisfies Parents-to-children.

(b) Let's construct, then, an unsigned truth tree to test the argument

$$\mathbf{A} \quad ((P \wedge Q) \wedge R) \therefore (Q \wedge (R \wedge P)).$$

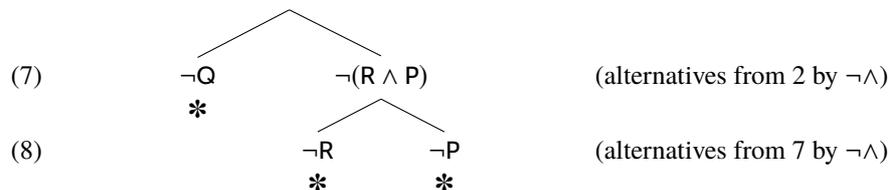
We start by assuming that there is valuation which makes the premiss true and negation of the conclusion true, so now write down those wffs unsigned at the top of the trunk:

- (1)  $((P \wedge Q) \wedge R)$  (premiss)
- (2)  $\neg(Q \wedge (R \wedge P))$  (negated conclusion)

We follow the tactic of keeping our tree tidy by applying a non-branching rule before a branching rule when we have a choice. So we continue by adding more wffs which must also be true (on the same supposed valuation):

- (3)  $(P \wedge Q)$  (from 1 by  $\wedge$ )
- (4)  $R$  (from 1 by  $\wedge$ )
- (5)  $P$  (from 3 by  $\wedge$ )
- (6)  $Q$  (from 3 by  $\wedge$ )

Next, we apply our branching rule twice to extract alternatives from (2):

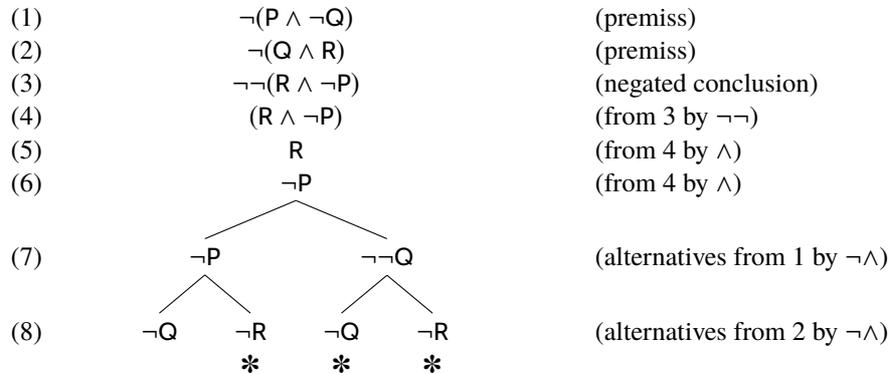


Each path leads to contradiction, since it contains a wff and its negation (both invisibly assigned T). Hence the assumption that **A**'s premiss and negated conclusion are both true leads at every turn to absurdity. Hence **A** is valid. Just as you will have hoped!

(c) Next, consider the argument

$$\mathbf{B} \quad \neg(P \wedge \neg Q), \neg(Q \wedge R) \therefore \neg(R \wedge \neg P).$$

You know now how a truth tree will start. And here's one finished tree:



We have once more applied non-branching rules first, starting with ( $\neg\neg$ ) to get line (4), and then applying ( $\wedge$ ) to get the next two lines. Then we applied our branching rule ( $\neg\wedge$ ) to extract the implications of (1) and (2). (Experiment: what happens if we unpack (1), (2) and (3) in different orders: which order leads to the neatest tree, and why?)

This tree is now finished in the sense that no complex wff on the open path remains to be dealt with. – we have already fully drawn out the implications of each of (1) to (4). And we are left with an open path where there is no more unpacking to be done. So there can be no lurking contradictions yet to be exposed. This shows that **B** is invalid.

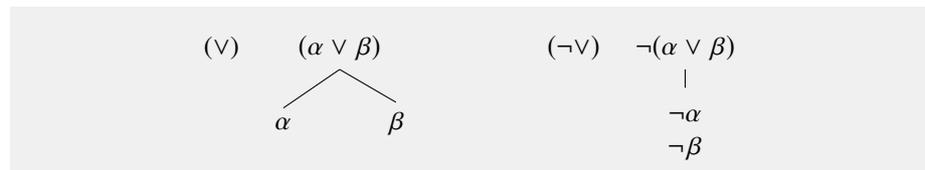
Why so? Well, by inspection you will see that our three rules so far not only satisfy the Parents-to-children principle, but also a kind of converse:

*Children-to-parents* Whenever we apply a rule, whether it is branching or non-branching, if a parent’s one or two resulting children on a particular path through the tree are true (on some supposed valuation), so too is the parent wff.

Therefore, as we go up some particular path on a tree constructed using our rules so far, if all the children we encounter on that path are true, so too are their parents (and their parents’ parents, etc.). So in our current example, if we take the valuation making the simple wffs ‘ $\neg P$ ’, ‘ $\neg Q$ ’, and ‘ $R$ ’ on the open path all true, then this will make the parents of these wffs true, and those parents will make *their* parents (if any) true, thereby making the complex wffs (1), (2) and (3) all true together. Which confirms that **B** is invalid.

## 19.2 Rules for disjunction

(a) Next, here are the two rules for unpacking disjunctions, naked and negated:



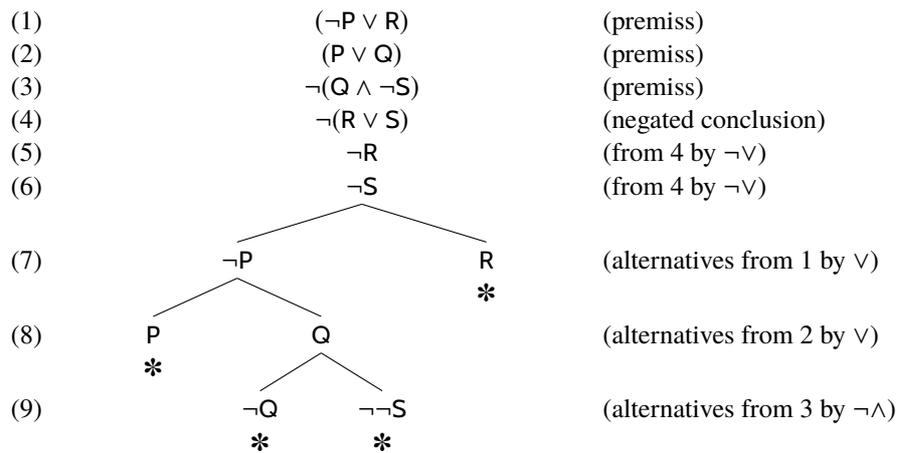
The branching rule ( $\vee$ ) is of course to be interpreted like the branching rule ( $\neg\wedge$ ); and the non-branching rule ( $\neg\vee$ ) is to be interpreted like the non-branching rule ( $\wedge$ ). And our rules obviously continue to satisfy Parents-to-children and Children-to-parents.

§19.3 Rules for the material conditional

We can now apply these disjunction rules to construct an unsigned truth tree to test e.g. the following for validity:

$$\mathbf{C} \quad (\neg P \vee R), (P \vee Q), \neg(Q \wedge \neg S) \therefore (R \vee S).$$

And here is one possible result (if we follow the usual tactic of applying non-branching rules first when we can, to keep the tree tidy):



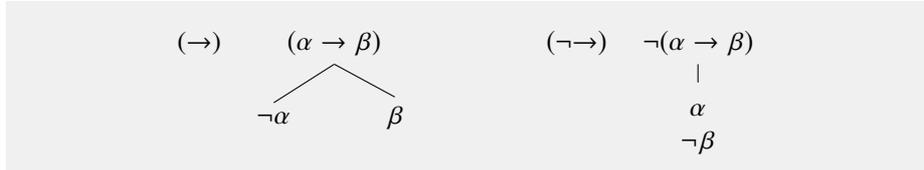
We have assumed that some valuation makes the premisses and negated conclusion of **C** all true, and worked backwards to try to uncover a valuation of simple wffs which makes this the case. But we hit contradiction every way we turn. So there is no such valuation. Hence the argument **C** is valid.

(b) We have continued to put line numbers on the left of our trees and put some minimal commentary on the right. But these are optional extras. Strip them away, and the uncommented tree which shows argument **C** is valid is very quick and neat.

We have in fact seen this very same argument before in §14.5(b). It was there labelled **E** and we ran a truth-table test on it. So we are now in a position to compare the amount of work involved in the two tests for tautological validity applied to the same simple argument. And we see that in the time it takes to write down just the frame of the truth-table test (writing the premisses and conclusion across the top, and the catalogue of sixteen possible valuations of the atoms down the left side), i.e. before we have done *any* real work filling in the table by evaluating premisses and conclusions, the corresponding (uncommented) tree test for validity will be finished. That's not at all untypical.

19.3 Rules for the material conditional

If we are working in a PLC language, we will need two additional rules for dealing with material conditionals, naked and negated. We've seen the rules in their signed form before in §18.4(a), so we simply need to restate them in their unsigned form. Interpret the following diagrams for rules just like the previous cases, and note that the two Parent/Children principles hold again:



(a) Let's work through an example exploiting these rules. But now we speed things up by dropping the explicit running commentary on the right. So consider

**D**  $((P \rightarrow Q) \rightarrow (R \rightarrow \neg S)), (\neg R \rightarrow (Q \wedge P')), (\neg P \rightarrow P') \therefore (S \rightarrow P')$ .

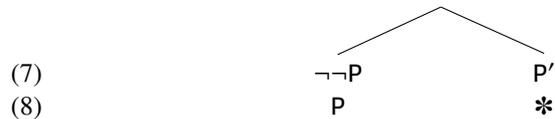
The inference is invalid, then, if there is a valuation which makes the following all true:

- (1)  $((P \rightarrow Q) \rightarrow (R \rightarrow \neg S))$
- (2)  $(\neg R \rightarrow (Q \wedge P'))$
- (3)  $(\neg P \rightarrow P')$
- (4)  $\neg(S \rightarrow P')$

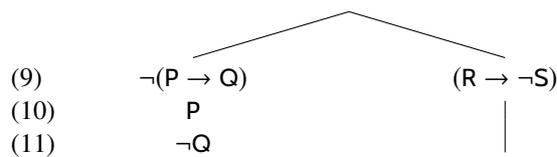
Applying the tactic of applying non-branching rules first, we start 'working backwards' by using rule  $(\neg \rightarrow)$  on (4). So these too must be true on the supposed valuation:

- (5)  $S$
- (6)  $\neg P'$

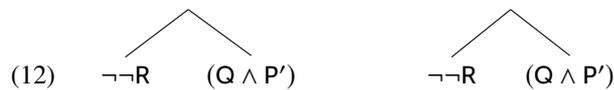
Let's deal with (3) next, because a moment's thought tells us that one resulting branch will immediately close. Then unpacking the wff on the other branch, we get:



We look next at the implications of assuming (1) is true on our supposed valuation, and then apply our non-branching rule  $(\neg \rightarrow)$  on the left:



Which leaves just the wff at (2) and the wff on the right-hand branch at (9) to unpack. Let's deal with the first, and add the results of applying the rule  $(\rightarrow)$  to both open paths:



Which yields four open paths through the tree, and on each path there remain complex wffs still to be unpacked. Take the left-most path first. We can continue it thus:

- (13)  $R$

And now there are no complex wffs on this left-most path still to be unpacked; so there can be no lurking contradiction to be found, at least on that path. So let's make the

## §19.4 'Checking off'

181

simple wffs on that path all true, by putting  $P := T$ ,  $Q := F$ ,  $R := T$ ,  $S := T$ ,  $P' := F$ . By repeated applications of the Children-to-Parents principle, this valuation will make the more complex wffs found as we go up that path all true too. In particular, this valuation makes (1) to (4) at the top of the path true, showing that argument **D** is invalid.

If – as is usually the case – all we care about is a verdict on validity, we can now stop work on the tree. Finding one bad line on a truth table is enough to show the argument being tested is invalid. It is exactly similar here: finding one open branch on which every complex wff has been unpacked is enough to show that the argument being tested is invalid.

(b) We have now met seven unpacking rules for dealing with complex PLC wffs on an unsigned tree. Do we need any more?

No. Any complex wff – i.e. any wff other than an atom or negated atom – falls into one of two classes. Either (i) it has one of the three binary connectives as its main connective. Or else (ii) it begins with a negation sign and then the main connective of the rest of the wff is either a binary connective or negation again. So (i) covers three kinds of wff, and (ii) covers another four. We need seven rules for dealing with complex wffs of these seven kinds. Which is exactly what we now have.

## 19.4 'Checking off'

(a) We should now highlight an obvious but crucial point.

Suppose that we have a currently open path as we work through a truth tree starting with the premisses and negation of the conclusion of an argument. This doesn't yet settle anything, if there is still a complex wff on the path that hasn't yet been unpacked.

Why? Because that wff's implications could – for all we yet know – generate a contradiction, closing the path. Hence, before we can conclude anything about an argument from a currently open path on its tree, we do have to check that *every* complex wff on the path really has been taken account of.

To make this check easier, it is very useful to have a device for signalling quite explicitly when a complex wff has been unpacked. The standard convention is to *check off* a wff at the point when we applying an unpacking rule to it.

(b) Now, suppose we have already applied the appropriate rule, adding children to every open path containing the parent wff we are unpacking. Then it would be pointlessly redundant to apply the rule to the very same wff again. It would be worse than redundant to get 'stuck in a loop', repeatedly re-applying the same rule to the same wff, time without end. Henceforth, then, we adopt the following rule:

When the appropriate unpacking rule has been applied to a complex wff, we *check off* the wff (with a '✓'). A wff, once checked off, is then treated as 'used up', i.e. as no longer available to have one of the unpacking rules (re)applied to it.

(c) Let's have a mini-example of checkmarks in use. So consider again

**E**       $((P \wedge Q) \vee R) \therefore \neg(\neg P \vee \neg R).$

This was argument **G** in the last chapter, where we used a signed tree to test it. Now we use an unsigned tree, checking off complex wffs as we go. So start the tree with

- (1)  $((P \wedge Q) \vee R)$
- (2)  $\neg\neg(\neg P \vee \neg R)$

Applying the unpacking rule ( $\neg\neg$ ), this becomes

- (1)  $((P \wedge Q) \vee R)$
- (2)  $\neg\neg(\neg P \vee \neg R) \checkmark$
- (3)  $(\neg P \vee \neg R)$

And then applying the unpacking rule ( $\vee$ ) to the first wff we get

- (1)  $((P \wedge Q) \vee R) \checkmark$
- (2)  $\neg\neg(\neg P \vee \neg R) \checkmark$
- (3)  $(\neg P \vee \neg R)$
- (4)  $(P \wedge Q) \quad R$

Next, using the non-branching rule ( $\wedge$ ) on the left branch, the tree becomes

- (1)  $((P \wedge Q) \vee R) \checkmark$
- (2)  $\neg\neg(\neg P \vee \neg R) \checkmark$
- (3)  $(\neg P \vee \neg R)$
- (4)  $(P \wedge Q) \checkmark \quad R$
- (5)  $P$
- (6)  $Q$

Now, in this simple little example – even if we had not been explicitly checking wffs off as we unpack their implications – we would hardly be likely to overlook the fact that we haven't yet dealt with the third wff! But the use of check marks makes this quite unmissable. So we continue:

- (1)  $((P \wedge Q) \vee R) \checkmark$
  - (2)  $\neg\neg(\neg P \vee \neg R) \checkmark$
  - (3)  $(\neg P \vee \neg R) \checkmark$
  - (4)  $(P \wedge Q) \checkmark \quad R$
  - (5)  $P$
  - (6)  $Q$
  - (7)  $\neg P \quad \neg R \quad \neg P \quad \neg R$
- \* \*

And *now* we are done. There are two open paths through the tree on which every complex wff is explicitly checked off. Therefore, as we already know, **E** is invalid.

19.5 Trees for tautologies

(a) We can test whether an argument  $\alpha_1, \alpha_2, \dots, \alpha_n \therefore \gamma$  is tautologically valid by putting all the  $\alpha$ s and the negated conclusion  $\neg\gamma$  at the top of a truth tree, and then using the tree-building rules. The argument is tautologically valid if and only if every path through a tree headed by the  $\alpha$ s and  $\neg\gamma$  closes.

A tautology, we remarked in §15.3(c), can be thought of as the conclusion of a tautologically valid argument with no premisses. This immediately tells us how to use a truth tree to determine whether a wff  $\gamma$  is a tautology. We put its negation  $\neg\gamma$  at the top of a truth tree, and use the tree-building rules as before. Then  $\gamma$  is a tautology if and only if every path through a tree headed by  $\neg\gamma$  closes.

(b) For a first example, take the wff

$$\mathbf{F} \quad ((P \wedge (Q \vee R)) \vee (\neg P \wedge (\neg Q \vee \neg R))).$$

Is this a tautology? To apply the tree test, we start with the wff's negation (quietly assigned the value T):

$$(1) \quad \neg((P \wedge (Q \vee R)) \vee (\neg P \wedge (\neg Q \vee \neg R)))$$

We then work backwards to determine whether this assumption in fact leads to contradiction. This wff has the form  $\neg(\alpha \vee \beta)$ , where  $\alpha$  is ' $(P \wedge (Q \vee R))$ ' and  $\beta$  is ' $(\neg P \wedge (\neg Q \vee \neg R))$ '. So we apply the non-branching rule ( $\neg\vee$ ) to (1) to get two derived wffs  $\neg\alpha$  and  $\neg\beta$  at lines (2) and (3). We then continue in the obvious way by unpacking (2) to get as far as the following:

$$\begin{array}{l} (1) \quad \neg((P \wedge (Q \vee R)) \vee (\neg P \wedge (\neg Q \vee \neg R))) \checkmark \\ (2) \quad \neg(P \wedge (Q \vee R)) \checkmark \\ (3) \quad \neg(\neg P \wedge (\neg Q \vee \neg R)) \\ \quad \swarrow \quad \searrow \\ (4) \quad \neg P \quad \quad \quad \neg(Q \vee R) \checkmark \\ (5) \quad \quad \quad \quad \quad \neg Q \\ (6) \quad \quad \quad \quad \quad \neg R \end{array}$$

At this stage, the only complex wff to unpack is (3); so let's apply the rule ( $\neg\wedge$ ), adding the results to both open branches:

$$(7) \quad \begin{array}{cc} \neg\neg P & \neg(\neg Q \vee \neg R) \\ * & \end{array} \quad \begin{array}{cc} \neg\neg P & \neg(\neg Q \vee \neg R) \end{array}$$

Which leaves us with three branches which so far remain open. Tackling the middle one first (hardly difficult!) we can add

$$(8) \quad \quad \quad P$$

Remembering that (3) and now ' $\neg\neg P$ ' at (7) have been checked off, there are no remaining complex wffs on *this* path waiting to be unpacked. Hence there can be no lurking contradictions on this path – even if we continue work elsewhere on the tree, *this* path

must stay open. Consider, then, the valuation which makes its simple wffs  $P$ ,  $\neg Q$  and  $\neg R$  all true; this will make the more complex wffs further up this path true, hence will make (1) – the negation of  $\mathbf{F}$  – true. Hence our valuation makes  $\mathbf{F}$  false. Therefore it is not a tautology.

(c) One last example in this chapter. Consider the string of symbols

$$\mathbf{G} \quad (\neg(\neg(P \wedge Q) \wedge \neg(P \wedge R)) \vee \neg(P \wedge (Q \vee R))).$$

Is this string even a wff? Well, that can be confirmed by the trick of changing the shape of some matching pairs of brackets, thus:

$$[\neg\{\neg(P \wedge Q) \wedge \neg(P \wedge R)\} \vee \neg\{P \wedge (Q \vee R)\}].$$

Then we see this is a wff of the form  $[\alpha \vee \beta]$  with  $\alpha$  the wff ‘ $\neg\{\neg(P \wedge Q) \wedge \neg(P \wedge R)\}$ ’, and  $\beta$  the wff ‘ $\neg\{P \wedge (Q \vee R)\}$ ’. So is it a tautology?

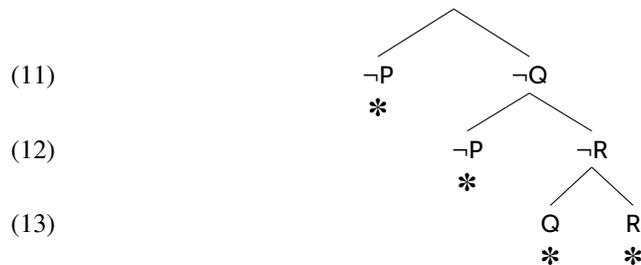
We start our tree test with the negation of  $\mathbf{G}$ , which will therefore be of the form  $\neg[\alpha \vee \beta]$  with the same  $\alpha$  and  $\beta$ . So we can immediately apply the rule  $(\neg\vee)$ , to get

$$\begin{array}{l} (1) \quad \neg[\neg\{\neg(P \wedge Q) \wedge \neg(P \wedge R)\} \vee \neg\{P \wedge (Q \vee R)\}] \checkmark \\ (2) \quad \neg\neg\{\neg(P \wedge Q) \wedge \neg(P \wedge R)\} \\ (3) \quad \neg\neg\{P \wedge (Q \vee R)\} \end{array}$$

Continue in the obvious way, by applying the rule  $(\neg\neg)$  to both (2) and (3), and check them off. Then apply the rule  $(\wedge)$  twice to get

$$\begin{array}{l} (1) \quad \neg[\neg\{\neg(P \wedge Q) \wedge \neg(P \wedge R)\} \vee \neg\{P \wedge (Q \vee R)\}] \checkmark \\ (2) \quad \neg\neg\{\neg(P \wedge Q) \wedge \neg(P \wedge R)\} \checkmark \\ (3) \quad \neg\neg\{P \wedge (Q \vee R)\} \checkmark \\ (4) \quad \{\neg(P \wedge Q) \wedge \neg(P \wedge R)\} \checkmark \\ (5) \quad \{P \wedge (Q \vee R)\} \checkmark \\ (6) \quad \neg(P \wedge Q) \\ (7) \quad \neg(P \wedge R) \\ (8) \quad P \\ (9) \quad (Q \vee R) \end{array}$$

As this reminds us again, one lovely feature of working with trees is that, as we apply the unpacking rules, the wffs that we add to the tree get simpler and simpler. We are now just left with the last four very much more tractable wffs to deal with (as everything earlier has been checked off). It is a very simple matter to confirm that the tree now quickly closes. Unpacking (6), (7) and (9) in turn gives us:



§19.6 Two sorts of trees: a reality check

185

Hence the assumption that we can make the negation of  $G$  true leads to contradiction;  $G$  is indeed a tautology.

## 19.6 Two sorts of trees: a reality check

Let's finish this chapter by very briefly emphasizing a key difference between the construction/parse trees we met before and our truth trees – ignoring the relatively trivial matter of which way up we draw them.

Reading up a construction tree, an expression before a branch point is a wff if and only if *both* expressions above the branch point are wffs. Reading down a truth tree, a wff before a branch point is true if and only if *at least one* of wffs below the branch point is true.

So we can put it this way: there is a good sense in which branchings on construction trees are *conjunctive*, while branchings on truth trees are *disjunctive*. You need to be alert to this important distinction whenever you encounter trees of one sort or another in logic – are we conjoining simpler claims into more complex ones or exploring disjunctive options?

## 19.7 Summary

In the previous chapter, we saw how to test an argument for tautological validity by 'working backwards'. We assume we can make the premiss and negated conclusion of the argument all true, and try to determine what the assumed bad valuation would look like. In the general case, this will necessitate considering alternative scenarios, and the resulting exploration can naturally be laid out in the form of a downward-branching tree.

By very simple tricks, we ensured that every wff on such a tree was assigned the value T. But then, if every wff is assigned the same value, we need not bother to explicitly write down the value – resulting in this chapter in our preferred *unsigned* trees.

Still somewhat informally, we have stated and explored the rules for constructing an unsigned tree using 'unpacking' rules taking us from more complex wffs on the tree to simpler wffs, and 'closing off' paths through the tree if they contain a contradictory pair of wffs. (We give an official account in the next chapter.)

An argument is tautologically valid if every path through a properly constructed tree starting with the premisses and negated conclusion ends in contradiction and is closed. The argument is not valid if a corresponding tree has a branch that stays open even when every complex wff on it is unpacked.

We also noted that we can similarly use a tree to test a wff to determine whether it is a tautology. A wff is a tautology if and only if every path through a properly constructed tree starting with its negation ends in contradiction.

186

Unsigned trees

Exercises 19

(To be added)

## 20 Truth trees further explored

The last two chapters introduced truth trees in a relatively relaxed way. This chapter starts by telling the same story again more crisply, in order to fix ideas and to provide a resource for future revision. We then say just a little more about the practice of using truth trees to test arguments for tautological validity. (We will not keep repeating the point made in §19.5, that trees can also be used to test wffs for being tautologies.)

### 20.1 Reminders about trees and our terminology

As is now familiar, truth trees will be downward-branching *binary* trees, meaning that a branch never forks more than two ways. We can take the general idea of a downward-branching binary tree to be intuitively clear. However, we should perhaps gather together some terminology, old and new. First, for binary trees in general:

- (1) Positions on a tree are conventionally called *nodes*.
- (2) Going down a binary tree, each node is followed by zero, one, or two *successor nodes*. A node is conventionally thought of as being joined to its successors by *edges*; but, for visual economy, in the case of truth trees, we have only drawn edges from a node to its successors when there is more than one successor – in other words, when the tree branches. (Compare our parse trees in §9.2 where, at least in a minimal way, we depicted all the edges.)
- (3) Starting from the top node of the tree, the nodes and edges going downwards until the first branch point, if there is one, form the *trunk* of the tree.
- (4) Start again from the top node and track down some route along edges through the tree to an *end-node* without successors, i.e. to the bottom tip of some branch. The nodes and edges form a (full) *path* through the tree.
- (5) On a tree (as opposed to other kinds of ‘graphs’ made of nodes and edges) the paths downwards from the top node fan out without ever rejoining at a lower level: as in the case of a respectable family tree, incestuous reunions are banned. In an obvious sense, then, a tree is *loop-free*; this entails that there is only one path from the top node down to a given end-node.

Now for some initial terminology specific to our truth trees:

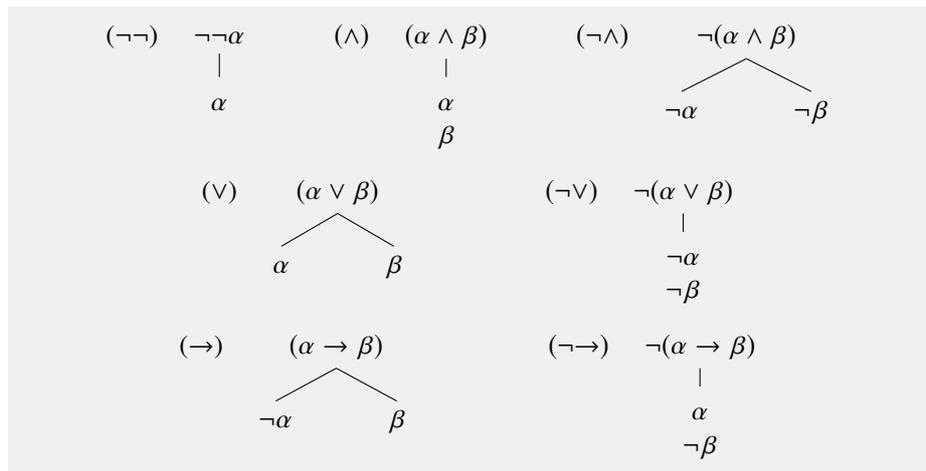
- (6) On a truth tree, the nodes are occupied, labelled, or decorated by wffs (either from a PL or from a PLC language). If a wff  $\alpha$  occupies a node on a path, we will say more briefly that the path *contains*  $\alpha$ , or that  $\alpha$  is *on* the path.

- (7) If a path contains, for some wff  $\alpha$ , both that wff and its negation  $\neg\alpha$ , then the path is *closed*. Otherwise it is *open*.

## 20.2 Testing arguments using truth trees, in summary

(a) Suppose, then, that we want to determine whether a given PL or PLC inference is tautologically valid. We begin a truth tree to test this inference by placing its premisses and its negated conclusion at nodes forming the initial trunk of a tree (we can imagine each of them accompanied by an invisible ‘:= T’). We then aim to systematically ‘unpack’ the implications of those initial assignments – and then the implications of these implications, etc. – exploring alternatives when necessary.

As before, we call the wff that we are unpacking at any stage the *parent* wff, and call the new wffs we add when we unpack its implications its *children*. And we do the needed ‘unpacking’ of complex wffs using the following now-familiar rules (the last two being needed only if we are working with PLC wffs):



How do we interpret these diagrammatically presented rules? To repeat, the non-branching rule ( $\neg\neg$ ) tells us that, given a parent wff of the form  $\neg\neg\alpha$  (invisibly assigned T) on a tree, we unpack its implications by adding  $\alpha$  (invisibly assigned T) as a child to new nodes at the foot of every open path containing that particular parent wff.

The non-branching rule ( $\wedge$ ) tells us that, given a parent wff of the form  $(\alpha \wedge \beta)$ , we unpack its implications by adding both  $\alpha$  and  $\beta$  to the foot of every open path containing that parent wff. Similarly for the other non-branching rules.

The branching rule ( $\vee$ ) tells us that, given a parent wff of the form  $(\alpha \vee \beta)$ , we unpack the alternatives by adding a binary fork with a branch to  $\alpha$  and a branch to  $\beta$  to the foot of every open path containing that parent wff. Similarly for the other branching rules.

(b) And how do we deploy these unpacking rules to construct an unsigned truth tree to test an inference for tautological validity? Here, in summary form, is the procedure we have been using all along (adopting the policy of definitely stopping work on a tree if and when we find an open path on which we have unpacked every complex wff).

To construct a tree for testing whether the inference  $\alpha_1, \alpha_2, \dots, \alpha_n, \therefore \gamma$  is tautologically valid, proceed as follows:

- (P1) Start the trunk of the tree by placing  $\alpha_1, \alpha_2, \dots, \alpha_n, \neg\gamma$  down the top nodes.
- (P2) Inspect every path through the tree which doesn't yet finish with an absurdity sign to see whether it involves a contradiction, i.e. to see whether it contains both  $\beta$  and also the corresponding formula  $\neg\beta$  (for some wff  $\beta$ ). If it does, then we mark that the path is closed by adding the absurdity sign '\*'.
- (P3) If every path through the tree is closed with an absurdity sign, then stop! Or if there is an open path on which every complex wff has been checked off, then also stop!
- (P4) Otherwise, we choose some unchecked complex formula  $\beta$  that occurs on an open path. We then apply the appropriate unpacking rule and add the results to new nodes at the foot of every open path containing that occurrence of the wff  $\beta$ . We then check off  $\beta$  (using a checkmark '✓').
- (P5) Loop back to step (P2).

The 'checking off' rule ensures that the appropriate unpacking rule can only be applied once to any complex wff on the tree. Further, the new wffs added when we apply a rule are always shorter than the unpacked wff they result from. It follows that our unpacking-and-tree-extending procedure must eventually terminate as we unpack and check off shorter and shorter wffs. (We are assuming, as always, that an inference has only a finite number of premisses to start with!) Call any tree which results at the end of this process a *test tree* for the given inference under test.

(c) Two more bits of terminology applying to truth trees:

- (1) A tree is *closed* if and only if every path through the tree is closed. Otherwise it stays *open*.
- (2) A path through a tree is *completed open* if and only if it is open and every complex wff on that path has been duly unpacked and checked off.

It is evident that, on any particular run of our tree building procedure, the test tree we end up with must either be closed, or else have at least one completed open path.

What do we learn from these alternative upshots?

A test tree for the inference  $\alpha_1, \alpha_2, \dots, \alpha_n, \therefore \gamma$  delivers verdicts as follows:

- (VP) If the test tree is closed, the tested inference is tautologically valid.
- (VN) If the test tree has a completed open path, the inference is tautologically invalid.

Why so? If the supposition that the  $\alpha$ s plus  $\neg\gamma$  are all true leads to contradiction whichever way we turn, then the supposition has to be false – we can't have the  $\alpha$ s true and conclusion  $\gamma$  false. Or put it this way. Suppose everything on the initial trunk of our tree is true. Then our tree-building rules ensure that, as we extend the tree, we

can choose a path where the wffs continue to be all true (see the ‘Parents-to-children’ principle in §19.1). And a path will all-true wffs cannot contain contradictory wffs, so will stay open. Hence, if there *is* a valuation which makes  $\alpha$ s plus  $\neg\gamma$  all true, the tree cannot close. Contraposing, we get (VP): if the tree does close, then there is no valuation making the  $\alpha$ s true and  $\gamma$  false.

Suppose now that, starting with the  $\alpha$ s plus  $\neg\gamma$ , we find a completed open path on the tree. Pick a valuation which makes the simple wffs on the path true – we can do this consistently as these simple wffs must be consistent with each other, or the path would be closed after all! Our tree-building rules ensure that, as we go up the path, this valuation makes all the more complex wffs we encounter true too, including the ones we started off with (see the ‘Children-to-parents’ principle in §19.1). So yes, there is a way of making the  $\alpha$ s true and  $\gamma$  false, and the tested inference is indeed invalid. As (VN) says.

Hence our tree method does indeed work as advertised.

(d) Note that, as you would expect, our procedural rules allow us a free choice of which unchecked complex wff to unpack next each time we circle round to step (P4) – at least until we get to deal with the last unchecked complex wff. So there is usually no unique test tree. However, we do have the following key result:

Starting from the same initial wffs  $\alpha_1, \alpha_2, \dots, \alpha_n, \neg\gamma$  on the trunk of the tree, either any resulting test tree is closed, or alternatively any test tree has a completed open path. We can not get some closed trees and some with completed open paths.

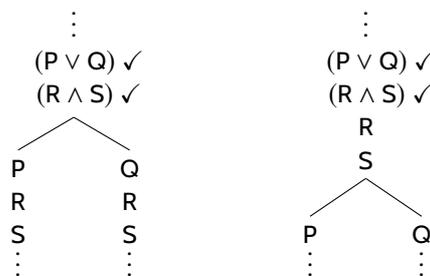
Why so? Because, given claims (VP) and (VN), if we could get both a closed tree and a (different!) tree with a completed open branch when testing a given argument, the argument would be both valid and invalid, which is impossible.

This confirms that it can’t matter – at least as far as the ultimate verdict on validity is concerned – in which order we apply our tree-building rules.

### 20.3 Trees in practice

So much for a summary review of the theory. Now for three quick comments on using trees in practice. Again there is nothing new here; again we are just gathering together points we touched on in passing in the last two chapters.

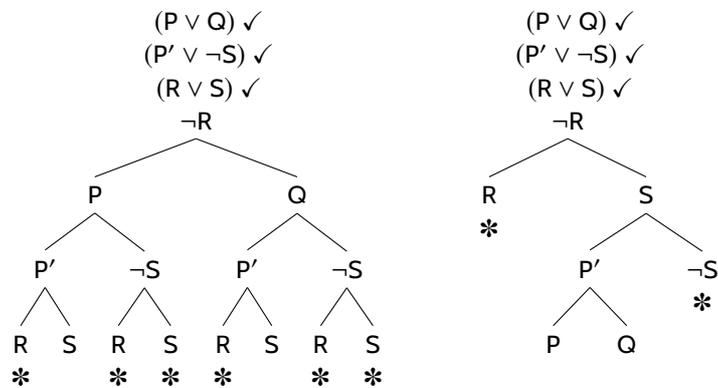
(a) Compare first the following two fragments of truth trees:



§20.3 Trees in practice

The only difference between these two (parts of) trees is the order in which we have unpacked the pair of complex wffs: we end up, either way, with two paths with just the same wffs on them. But the second tree where we apply the non-branching rule first obviously avoids some duplication across branches and so needs a bit less time and ink to write down. That’s a miniature illustration of our first maxim for getting neater trees: *in general, apply non-branching unpacking rules before branching rules, when we have the choice.*

(b) Compare next the following two completed truth trees. This time, the three complex wffs at the top of the tree all need to be unpacked using branching rules. But it still makes a difference which order we do the unpacking:



On the left, we unpack the initial wffs in the order they appear on the trunk of the tree; and on the right, we unpack them in the reverse order. We end up in each case with two open paths, with the same simple wffs on them. But obviously the right-hand tree is a lot neater, because we can close off branches sooner. For instance, by choosing to unpack  $(R \vee S)$  first, we can immediately close off one alternative at the first fork, and don’t have to explore it any further.

So that illustrates a second practical maxim: *wherever possible, choose to apply rules in an order which closes off branches as early as possible.*

(c) As we noted in §19.2(a), once we have found *one* completed open path on a tree for an argument, that is already enough to show that the argument in question is invalid. We have built this fact into our described official procedure for tree-building, but let’s highlight it again: *assuming that all we care about is a yes/no verdict on validity, stop work on a tree once we have found one completed open path.*

(d) Let’s work through a couple more annotated examples, keeping our practical maxims in mind. First, we will use a tree to test the following inference for validity:

$$\text{A} \quad (P \wedge (R \vee Q)), ((P \wedge \neg Q) \vee (\neg(P' \vee R) \vee S)), ((P' \wedge S) \vee \neg Q) \quad \therefore ((R \vee Q) \wedge S).$$

So we start in the familiar way with the premisses and negated conclusion at the top of the tree. We note that only one of these complex wffs calls for a non-branching rule to be applied, so – following our first maxim – let’s start by unpacking that one:

192

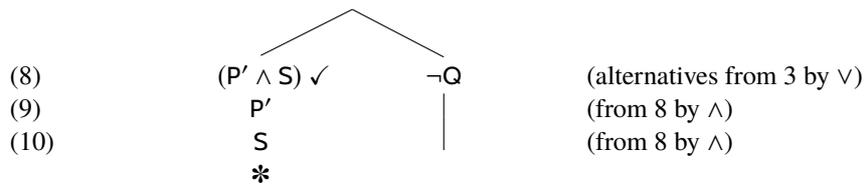
Truth trees further explored

- |     |   |                       |
|-----|---|-----------------------|
| (1) | $(P \wedge (R \vee Q)) \checkmark$                  | (premiss)             |
| (2) | $((P \wedge \neg Q) \vee (\neg(P' \vee R) \vee S))$ | (premiss)             |
| (3) | $((P' \wedge S) \vee \neg Q)$                       | (premiss)             |
| (4) | $\neg((R \vee Q) \wedge S)$                         | (negated conclusion)  |
| (5) | $P$   | (from 1 by $\wedge$ ) |
| (6) | $(R \vee Q)$  | (from 1 by $\wedge$ ) |

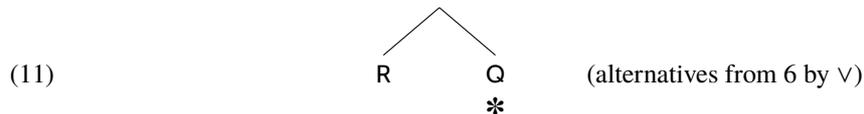
Now note that (6) and (4) each contains an occurrence of ' $(R \vee Q)$ ', naked in the case of (6) and to-be-negated-on-one-fork when we unpack (4). So – following our second practical maxim – let's continue the tree



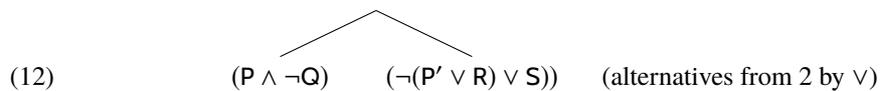
With (4) now checked off, we have a choice of three complex wffs to unpack next, those at lines (2), (3) and (6). And although the last is the simplest, our second practical maxim suggests dealing with (3) first, because it quickly leads to another closed branch:



The sensible thing to do next, obviously enough, is *now* to unpack (6). For this results in a branch which we can immediately close off:



If you are keeping track, then you will see that on our open path every complex wff has now been checked off except (2), so let's now unpack that at last:



And now the end is in sight! Let's work on the left-hand open path first. Every complex wff on that path has now been unpacked except the last one, ' $(P \wedge \neg Q)$ '. So let's deal with that. Check off (12) and add:

- |      |          |                                     |
|------|----------|-------------------------------------|
| (13) | $P$      | (alternatives from 12 by $\wedge$ ) |
| (14) | $\neg Q$ | (alternatives from 12 by $\wedge$ ) |

There are no more complex wffs on *this* path to be unpacked; and there is no contradictory pair of wffs to be found along it either. It is a completed open path.

So no more work is required: we can read off from the simple wffs on that open path the valuation  $P := T, Q := F, R := T, S := F$  which makes everything on that path true and hence shows that **A** is indeed invalid. We are done.

§20.3 Trees in practice

(e) For a final example, consider the following inference:

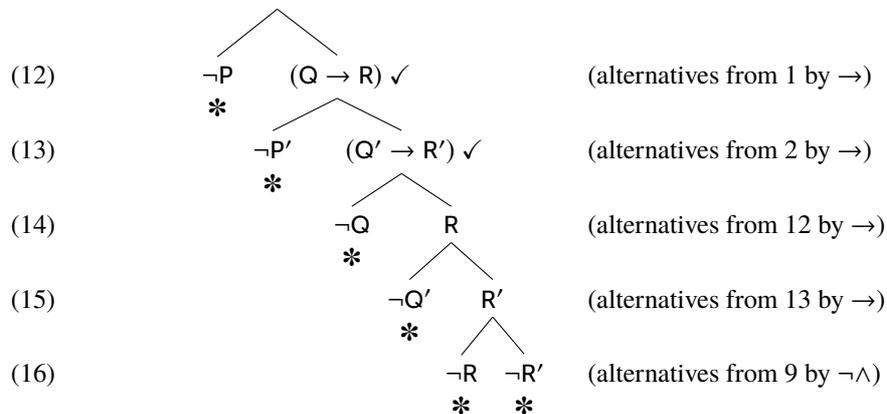
$$\mathbf{B} \quad (P \rightarrow (Q \rightarrow R)), (P' \rightarrow (Q' \rightarrow R')) \quad \therefore ((P \wedge P') \rightarrow ((Q \wedge Q') \rightarrow (R \wedge R'))).$$

There are six atoms in play, and you are very welcome to write down a sixty-four line truth-table to confirm that this inference is indeed valid. But a truth tree is a *lot* quicker if you remember our practical maxims (especially if you omit the optional commentary).

Begin the tree with the premisses and negated conclusion and, in accord with the first maxim, keep applying any possible non-branching rule:

- |      |  |                                   |
|------|--|-----------------------------------|
| (1)  | $(P \rightarrow (Q \rightarrow R))$  | (premiss)                         |
| (2)  | $(P' \rightarrow (Q' \rightarrow R'))$   | (premiss)                         |
| (3)  | $\neg((P \wedge P') \rightarrow ((Q \wedge Q') \rightarrow (R \wedge R'))) \checkmark$ | (negated conclusion) $\checkmark$ |
| (4)  | $(P \wedge P') \checkmark$   | (from 3 by $\neg\rightarrow$ )    |
| (5)  | $\neg((Q \wedge Q') \rightarrow (R \wedge R')) \checkmark$                             | (from 3 by $\neg\rightarrow$ )    |
| (6)  | $P$  | (from 4 by $\wedge$ )             |
| (7)  | $P'$   | (from 4 by $\wedge$ )             |
| (8)  | $(Q \wedge Q') \checkmark$   | (from 5 by $\neg\rightarrow$ )    |
| (9)  | $\neg(R \wedge R')$  | (from 5 by $\neg\rightarrow$ )    |
| (10) | $Q$  | (from 8 by $\wedge$ )             |
| (11) | $Q'$   | (from 8 by $\wedge$ )             |

That leaves (1), (2) and (9) as complex wffs still to check off. If we unpack (1) or (2) the tree forks with one branch immediately closing; not so if we unpack (9). So let's sensibly follow our second maxim and unpack (1) and (2) first. Then our tree neatly concludes:



So yes, **B** is valid, and can be shown to be so in a lot less than 64 lines of working.

It is worth pausing to think, though, what would have happened if we had ignored our maxims and started by fully unpacking (1) and then (2) using our branching rules first. Fully unpacking (1) splits the tree, and then splits it again. So we get three open paths. Then fully unpacking (2) adds another three-way split at the foot of those paths, so we get nine open paths, before we go on to complete the tree. Which just reinforces the point that, if we are not careful, a tree for a given argument *can* sprawl horribly, even if there's another very tidy tree which gives the same result.

## 20.4 Comparative efficiency?

We have repeatedly seen that using a tree test on an argument can be significantly faster than plodding through the corresponding truth table test. Sometimes the gain in speed is extraordinary: remember, for example, Example A in §18.2.

Yet, as just remarked, things *can* become pretty messy too. As we saw in earlier chapters, the work needed on a truth table explodes exponentially with the number of atoms in play in the argument. But the work needed on a truth tree can *also* grow exponentially as the wffs in an argument get more numerous. Imagine dealing with a lot of unnegated disjunctions using the unpacking rule ( $\vee$ ). The tree will keep forking. If things go really badly, we get two branches, then four, then eight, then sixteen, then . . . So the work needed to complete the tree can also explode. (To repeat, as we construct a tree, the added wffs get shorter and shorter; hence, even if things go really badly for us, the construction process must still eventually terminate. What we are worrying about now is *speed* of termination.)

We have seen that by following our practical maxims (stated again in §20.3), we can in good cases cut the work right down. But can we *always* tame trees and keep the needed amount of work under control? Indeed, is there *any* test for tautological validity which doesn't – in the worst cases – lead to exponential explosion as the arguments being tested get more involved?

No. Or so it is widely believed. But nobody really knows. The issue relates to a still unsolved deep problem in the foundations of the theory of computation, the so-called P vs NP problem. In fact, there is a million dollar prize on offer for the answer to *that*!

## 20.5 Summary

We have reviewed the procedures for constructing a truth tree for a PL/PLC inference, and the principles for extracting a verdict on whether the inference is tautologically valid – if the tree closes, the argument is valid; if the tree stays open, the argument is invalid.

We outlined arguments for those principles for extracting positive and negative verdicts from trees.

Test trees are not unique, but always give the same verdict.

Some practical maxims for working with trees: apply non-branching rules first; look for branchings that close quickly; stop once one completed open path is found.

Trees are typically faster than truth tables. But as we deal with trees for more and more complex inferences, in the worst cases the amount of work needed on a tree can exponentially explode.

For some elaboration of the proofs explaining why trees work see the Appendix 'Soundness and completeness for truth-trees'.

Exercises 20

195

Exercises 20

(To be added)