Metatheory

for truth-functional logic

Tim Button
University of Cambridge

In 2005, P.D. Magnus released the first version of his open-source logic textbook, forallx. In 2012, Tim Button adapted this for a first-year course at the University of Cambridge. Metatheory takes up exactly where the Cambridge implementation of forallx leaves off.

Once again, Tim would like to thank Magnus for making forallx freely available. With this in mind, Metatheory is released under the same open-source license as forallx.

Tim would also like to thank Owen Griffiths, Nicholas Heitler, Brian King and Mat Simpson for comments and corrections.

Contents

Introduction

This book is an introduction to the metatheory of truth-functional logic, also known as the propositional calculus.

This book does not itself contain an account of truth-functional logic, but instead assumes a prior understanding. More specifically, this book assumes a thorough understanding of the syntax, semantics and proof-system for TFL (truth-functional logic) that is presented in forall*x*:Cambridge. forall*x*:Cambridge can be obtained, free of charge, at www.nottub.com/forallx.shtml. There is nothing very surprising about the syntax or semantics of TFL, and the proof-system is a fairly standard Fitch-style system.

This book does not, though, presuppose any prior understanding of (meta)mathematics. Chapter 1 therefore begins by explaining induction from scratch, first with mathematical examples, and then with metatheoretical examples. Each of the remaining chapters presupposes an understanding of the material in Chapter 1.

One idiosyncrasy of this book is worth mentioning: it uses no set-theoretic notation. Where I need to talk *collectively*, I do just that, using plural-locutions. Although this is unusual, I think Metatheory benefits from being free of unnecessary epsilons and curly brackets. So, where some books say:

$\{\mathcal{A}_1, \mathcal{A}_2, ..., \mathcal{A}_n\}$ is inconsistent

to indicate that no valuation makes all of $\mathcal{A}_1, \mathcal{A}_2, ..., \mathcal{A}_n$ true, I say:

$\mathcal{A}_1, \mathcal{A}_2, ..., \mathcal{A}_n$ are jointly inconsistent

I use upper-case Greek letters for plural terms, so I might equally say:

Γ are jointly inconsistent

Inconsistency is therefore a (collective) property of sentences, rather than a property of sets (of sentences). Semantic entailment also comes out as a (collective) property of sentences. To indicate that every valuation that makes all of Γ true also makes $\mathcal{C}$ true, we write:

$$\Gamma \vDash \mathcal{C}$$

where, of course, Γ *are* some sentences. If there is a valuation that makes all of Γ true but $\mathcal{C}$ false, we write:

$$\Gamma \nvDash \mathcal{C}$$

The corresponding notions from proof-theory are also properties of sentences. Thus we write:

$$\Gamma \vdash \mathcal{C}$$

to indicate that there is a TFL-proof which starts with assumptions among Γ and ends with C on no further assumptions. If there is no such proof, we write:

$$\Gamma \nvdash C$$

A special case of this notation is used when there is a TFL-proof which starts with assumptions among Γ and ends with '$\bot$' (i.e. with an absurdity). In this situation, we say that $\Gamma \vdash \bot$, or that Γ are *jointly contrary*. (I should note that '$\bot$' is not a primitive 0-place connective of TFL. Rather, '$\bot$' shows up only in TFL's proof-theory. For more details, see forall$x$:Cambridge§26.7)

Induction 1

This book proves a bunch of results *about* formal logical systems. Specifically, it proves results about truth-functional logic(s): that is, logics where all of the connectives are truth functional. More specifically still, it proves results about TFL, a particular truth-functional logic whose syntax, semantics and proof-system are detailed in forall*x*:Cambridge.

When we prove results *about* formal logical systems, we inevitably have to rely upon some background logical reasoning. This reasoning need not, itself, be completely formalised. Nonetheless, the logical principles that we shall invoke informally – like *modus ponens*, or *universal generalisation* – should all feel familiar from their formalised counterparts.

However, when proving results about formal logical systems, we also rely upon *mathematical* reasoning. And there is an important mathematical principle which we shall use frequently: *induction*.[1] Indeed, almost every result in this book will invoke induction, in some form or other. So, the task of this first chapter is to familiarise ourselves with induction.

1.1 Stating and justifying the Principle of Induction

I mentioned that induction is a mathematical principle. In its plain vanilla form, it governs the behaviour of the *natural numbers*. These are all the finite, whole, non-negative numbers: 0, 1, 2, 3, and so on. (In what follows, whenever I say 'number', I always mean 'natural number'.) Here is the principle:

Principle of Induction. Let φ be any property such that both:

- 0 has φ; and
- for every natural number n: if n has φ, then $n + 1$ has φ.

Then every natural number has φ.

Here is an intuitive way to think about the Principle of Induction. Imagine that we had infinitely many dominoes – one for every number – balanced on their ends. Whenever we prove 'n has φ', we get to knock down the corresponding domino. Our aim is to knock down all of them. But there are too many of them for us to hope to knock them down one at a time. Fortunately, dominoes have a nice property: *if* we can arrange them in a line , so

---

[1]NB: *mathematical* induction must be sharply distinguished from the kind of *empirical* induction that David Hume discussed.

that when a domino falls it knocks over the next in line, *then* we can knock all the dominoes down just by knocking the first one down. To knock down the first domino corresponds with proving that 0 has φ. And to arrange the dominoes in this way corresponds with proving the conditional: if n has φ then $n + 1$ has φ.

Of course, this is just an analogy, and a spatial one at that. Why, then, should we *believe* the Principle of Induction? Here is an argument in its favour. Suppose you have proved both that 0 has φ and also that, for every n, if n has φ then $n + 1$ has φ. Now, pick any number, m. Given what you have proved, you would be justified in writing down all of the following:

> 0 has φ
> if 0 has φ, then 1 has φ
> if 1 has φ, then 2 has φ
> $\vdots$
> if $m - 1$ has φ, then m has φ

It would now follow, by m-many applications of *modus ponens*, that m has φ. And since m was arbitrary, this point generalises: every number has φ.

I doubt this line of argument will convince the determined sceptic.[2] And, in the end, the question 'Why believe in the Principle of Induction?' is a deep question in the philosophy of mathematics. That question is interesting, but it is a question for another day. For the purposes of this book, we shall take the Principle of Induction for granted.

1.2   Arithmetical proofs involving induction

To invoke induction in a proof, we must first identify some property, φ, that concerns us. We call this *the induction property*. We then do two things:

- Prove the *base case*: we need to show that 0 has φ. Often, this is a fairly straightforward task.
- Prove the *induction case*: we need to show that, for any n, if n has φ, so does $n + 1$. This is normally the harder step. We typically proceed by 'choosing some arbitrary n', and assuming 'for induction' that n has φ. If, on just this basis, we can show that $n + 1$ has φ, then we have proved the conditional: if n has φ, then $n + 1$ has φ. And since n was arbitrary, this holds for *all n*.

To illustrate the strategy, I shall offer a few simple examples of the use of induction, in establishing some arithmetical results. (Just to be clear: I am not going to make any use of the arithmetical results *themselves*; I just want to show how induction is used in proofs.) So, for example, suppose we want to prove a claim that should be very obvious:

Proposition 1.1. Every number is either even or odd. That is: for every number n, *either* there is a number k such that $n = 2k$ (i.e. n is even), *or* there is a number k such that $n = 2k + 1$ (i.e. n is odd).

---

[2]It is worth spending some time asking yourself: Why won't this convince the sceptic?

I shall prove this Proposition by induction. The induction property will be *x is either odd or even*. It suffices to prove the base case and the induction case; then all we need to complete the proof is to appeal to induction. Here goes:

Proof. Evidently, $0 = 2 \times 0$. So 0 has the induction property.

Now pick an arbitrary number, n, and suppose (for induction) that n has the induction property. There are two cases to consider

- Suppose n is even, i.e. $n = 2k$, for some number k. Then $n + 1 = 2k + 1$, i.e. $n + 1$ is odd.
- Suppose n is odd, i.e. $n = 2k+1$, for some number k. Then $n+1 = 2k+2 = 2(k+1)$, i.e. $n + 1$ is even.

Either way, $n + 1$ has the induction property. So if n has the induction property, then so does $n + 1$.

This completes the proof of Proposition 1.1 by induction. ∎

In addition to the use of induction, a couple of things should strike you about this proof. In particular, you should note that it is not a *formal* proof in any particular deductive system. It is a bit of informal, *mathematical*, reasoning, written in mathematised English. That said, it is perfectly rigorous. (The symbol '∎' at the end of the proof just indicates that the proof is complete; other texts use other symbols, e.g. 'QED' or '□'.)

To cement ideas, let me offer two more examples of proofs by induction:

Proposition 1.2. $2^0 + 2^1 + 2^2 + \dots + 2^n = 2^{n+1} - 1$, for any number n.

Proof. Evidently $2^0 = 1 = 2^{0+1} - 1$. So 0 has the induction property.

Now, pick any arbitrary number, k, and suppose (for induction) that

$$2^0 + 2^1 + 2^2 + \dots + 2^k = 2^{k+1} - 1$$

Then we have:

$$
\begin{aligned}
2^0 + 2^1 + 2^2 + \dots + 2^k + 2^{k+1} &= 2^{k+1} - 1 + 2^{k+1} \\
&= (2^{k+1} \times 2) - 1 \\
&= 2^{k+2} - 1
\end{aligned}
$$

So the induction property holds for $k + 1$ as well. ∎

Proposition 1.3. $0 + 1 + 2 + \dots + n = \frac{n^2+n}{2}$, for any number n.

Proof. Evidently $0 = \frac{0^2+0}{2}$. So 0 has the induction property.

Now, pick any arbitrary number, k, and suppose (for induction) that

$$0 + 1 + 2 + \dots + k = \frac{k^2 + k}{2}$$

Then we have:

$$0 + 1 + 2 + \ldots + k + (k+1) = \frac{k^2 + k}{2} + (k+1)$$
$$= \frac{k^2 + k + 2(k+1)}{2}$$
$$= \frac{(k^2 + 2k + 1) + (k+1)}{2}$$
$$= \frac{(k+1)^2 + (k+1)}{2}$$

So the induction property holds for $k + 1$ as well.   ■

We could keep going with examples from arithmetic, but these illustrate the use of induction, and give some sense of how powerful it is.

1.3   Strong Induction

Now let me introduce a second inductive principle:

Principle of Strong Induction. Let φ be any property. Suppose that, for every number n, if all the numbers less than n have φ, then n itself has φ. Then *every* number has φ.

Why should we believe *this*? In fact, it follows directly from the original Principle of Induction, as I shall now show:[3]

Proof. Suppose that, for every number n, if all the numbers less than n have φ, then n itself has φ. We aim to show that this entails that every number has φ (thereby vindicating the Principle of Strong Induction).
    To do this, we shall define a new property, ψ, as follows:

    every number less than x has φ

We shall now prove by (ordinary) induction that every number has ψ.

    Base case: 0 has ψ, vacuously, since 0 is the smallest number.
    Induction case: Suppose, for (ordinary) induction, that n has ψ. That is to say that every number less than n has φ. So, by the supposition at the start of this proof, n has φ. It follows that every number less than $n + 1$ has φ, i.e. that $n + 1$ has ψ.

It follows, by (ordinary) induction, that *every* number has ψ. And since every number has ψ, and for every number there is a larger number, every number has φ. Hence the use of Strong Induction is vindicated.   ■

It is worth illustrating the power of the Principle of Strong Induction, with an arithmetical example:[4]

---

[3]And it entails the original Principle too, though I shall not show this.
[4]This is one half of *the* Fundamental Theorem of Arithmetic; the other half is that prime factorisations are *unique*.

Proposition 1.4. Every number greater than 1 has a prime factorisation (i.e. it can be expressed as the product of one or more prime numbers).

Proof. Fix *n*, and suppose for (strong) induction that every number less than *n* has the induction property (i.e. the property that *if x > 1 then x has a prime factorisation*). There are now three (exhaustive) cases to consider:

> *Case 1: $n \leq 1$.* Then *n* obviously has the induction property.
> *Case 2: n is prime.* Then *n* obviously has the induction property.
> *Case 3: $n > 1$ and n is not prime.* Then $n = ab$ with both $n > a > 1$ and $n > b > 1$. By assumption, *a* and *b* have prime factorisations, i.e. $a = p_1 \times \ldots \times p_j$ and $b = q_1 \times \ldots \times q_k$, for some prime numbers $p_1, \ldots, p_j, q_1, \ldots, q_k$. So now $n = p_1 \times \ldots \times p_j \times q_1 \times \ldots \times q_k$, i.e. *n* has a prime factorisation.

So, in every case, *n* has the induction property on the assumption that every smaller number has the induction property. Hence, by strong induction, every number has the induction property. ■

Note that we needed to use *strong* induction here, because in the case where $n = ab$, we know nothing about *a* and *b except* that $n > a > 1$ and $n > b > 1$.

1.4 Induction on length

I now want to move on from arithmetical propositions to claims about TFL. By the end of this book, we will have proved some rather sophisticated results. But we need to start at the shallow end. In particular, we shall start by proving a syntactic result, which should already be familiar:

Proposition 1.5. In any TFL sentence, there are exactly as many instances of '(' as instances of ')'.

When I mention 'TFL sentences' here I have in mind the the objects that were defined recursively in forall*x*:Cambridge §6. In particular, then, I shall not allow us to rely upon the convenient bracketing conventions that we introduced later in forall*x*:Cambridge; so, we cannot, for example, allow for arbitrarily long conjunctions, and we cannot drop the outermost brackets of a sentence. To save space, in what follows I shall speak of 'sentences' rather than 'TFL sentences as given by the recursive definition of forall*x*:Cambridge §6'.

Now, we want to prove Proposition 1.5 by induction. But there is an immediate obstacle in our path. The property mentioned in Proposition 1.5 is *x has as many instances of '(' as of ')'*, and that is a property of *sentences*, rather than a property of *numbers*. So, if we are to prove this Proposition using induction, we need a way to associate sentences with numbers.

The easiest way to associate sentences with numbers is just by considering their length. More precisely, by the LENGTH of a sentence, we shall mean the number of instances truth-functional connectives (i.e. '¬', '∧', '∨', '→' and '↔') that the sentence contains.[5] Thus, for example:

---

[5]Many books refer to this as the *complexity* of a sentence.

- 'B' has length 0
- 'A_{234}' has length 0
- '$\neg(B \lor C_4)$' has length 2
- '$\neg\neg B$' has length 2
- '$\neg\neg\neg B$' has length 3

Having defined the notion of the length of a sentence, we could now prove Proposition 1.5 by (strong) induction on numbers, considered as lengths of sentences. But things will go easier for us if we introduce *yet another* inductive principle that relates more directly to TFL:

Principle of Strong Induction on Length. Let φ be any property. Suppose that, for any sentence, $\mathcal{A}$, if every shorter sentence has φ, then $\mathcal{A}$ has φ. Then every sentence has φ.

This principle follows directly from (ordinary) strong induction:

Proof. Suppose that, for any sentence, $\mathcal{A}$, if every shorter sentence has φ, then $\mathcal{A}$ has φ. We aim to show that this entails that every sentence has φ (thereby vindicating the Principle of Strong Induction on Length).

To do this, we shall define a new property of the numbers, ψ, thus:

every sentence of length x has φ

Pick any arbitrary number, n, and suppose for (strong) induction that every number less than n has ψ. That is, every sentence of length $< n$ has φ. So, by supposition, every sentence of length n has φ. So n has ψ. It follows, by strong induction, that every number has ψ.

Since every sentence has some number as its length, it now follows that every sentence has φ, as required. ∎

So, armed with the notion of the length of a sentence, and the Principle of Strong Induction on Length, we can now prove Proposition 1.5:

Proof of Proposition 1.5. Let $\mathcal{A}$ be any sentence, and suppose for (strong) induction (on length) that every shorter sentence has exactly as many instances of '(' as of ')'. We shall show that $\mathcal{A}$ has exactly as many instances of '(' as of ')', by considering all the different forms that $\mathcal{A}$ might have. (These six cases are given to us by the *recursive definition* of the notion of a sentence; for a reminder of the definition, see forall*x*:Cambridge §6.)

Case 1: $\mathcal{A}$ *is atomic.* Then $\mathcal{A}$ contains no brackets, so it has exactly as many instances of '(' as of ')'.

Case 2: $\mathcal{A}$ *is* $\neg\mathcal{B}$, *for some sentence* $\mathcal{B}$. Clearly, $\mathcal{B}$ is shorter than $\mathcal{A}$. By supposition, then, $\mathcal{B}$ contains exactly as many instances of '(' as of ')'. Hence $\mathcal{A}$ does too.

Case 3: $\mathcal{A}$ *is* $(\mathcal{B} \land \mathcal{C})$, *for some sentences* $\mathcal{B}$ *and* $\mathcal{C}$. Clearly, both $\mathcal{B}$ and $\mathcal{C}$ are shorter than $\mathcal{A}$. By supposition, then, $\mathcal{B}$ contains exactly as many instances of '(' as of ')', say i; and $\mathcal{C}$ contains exactly as many instances of '(' as of ')', say j. Hence $\mathcal{A}$ contains $i + j + 1$ instances of '(', and $i + j + 1$ instances of ')'.

Case 4: $\mathcal{A}$ *is* $(\mathcal{B} \lor \mathcal{C})$, *for some sentences* $\mathcal{B}$ *and* $\mathcal{C}$. Exactly similar to case 3.

Case 5: 𝒜 is (ℬ → 𝒞), for some sentences ℬ and 𝒞. Exactly similar to case 3.
Case 6: 𝒜 is (ℬ ↔ 𝒞), for some sentences ℬ and 𝒞. Exactly similar to case 3.

There are no other forms that 𝒜 could take. Hence, for any sentence, 𝒜, if every shorter sentence has exactly as many instances of '(' as of ')', then so does 𝒜. The result now follows by strong induction on length.   ∎

Now, I hope that it was obvious that the Proposition was true, before we set about proving it. The purpose of this proof was not, then, to *convince* you of its truth. Rather, I gave you this proof primarily in order to illustrate the Principle of Strong Induction on Length. Moreover, this proof exhibits three features that are worth mentioning explicitly.

First, even though the result was *obvious*, it took quite a lot of effort to prove it. This is fairly typical of metalogical results: showing them with sufficient rigour can be hard!

Second, we teamed an inductive principle with a recursive definition. This is fairly typical of metalogical results. Indeed, it is hardly surprising that induction and recursion should be best of friends. Recursive definitions tell you how to build objects up, stage-by-stage; proofs by (strong) induction involve showing that something holds at later stages if it has held at every stage so far.

Third, when we broke the proof down into cases, cases 3–6 were a bit repetitious. Indeed, all four of our two-place connectives have the same *syntactic* behaviour, and that is all we cared about during the course of the proof. So we could simply have rolled cases 3–6 together in something like the following:

Case 3: 𝒜 is (ℬ ⊲ 𝒞), for some two-place connective ⊲ and sentences ℬ and 𝒞. Clearly, both ℬ and 𝒞 are shorter than 𝒜. By supposition, then, ℬ contains exactly as many instances of '(' as of ')', say i; and 𝒞 contains exactly as many instances of '(' as of ')', say j. Hence 𝒜 contains $i + j + 1$ instances of '(', and $i + j + 1$ instances of ')'.

saving ourselves some time and effort.[6] Generally, if we need to do a proof by cases – and we often shall – we should pause before ploughing in, to think about exactly what the relevant cases are.

1.5   Unique readability

I shall close this chapter by proving two more results about the syntax of TFL.

When 𝒜 and ℬ are expressions, say that ℬ is an INITIAL SEGMENT of 𝒜 *iff* 𝒜 can be obtained by concatenating some (possibly no) symbols to the (right-hand) end of ℬ. For example:

- '¬' is an initial segment of '¬¬A' and an initial segment of '¬(P ∧ Q)', but not an initial segment of '(¬P → Q)'.
- '(P ∧ ¬Q)' is an initial segment of '(P ∧ ¬Q)', but not an initial segment of '((P ∧ ¬Q) ∨ R)'

Here is a cute result about initial segments.

---

[6]Note that '⊲', like '𝒜', is a symbol of the *meta*language, not the *object* language. For a reminder of the distinction, forall*x*:Cambridge §7.

Lemma 1.6. For any sentence $\mathcal{A}$, the only sentence which is an initial segment of $\mathcal{A}$ is $\mathcal{A}$ itself.

Proof. Let $\mathcal{A}$ be any sentence, and suppose (for induction) that every shorter sentence, $\mathcal{B}$, has the property that the only sentence which is an initial segment of $\mathcal{B}$ is $\mathcal{B}$ itself (this is the induction property). There are three cases:

> *Case 1: $\mathcal{A}$ is atomic..* Then the only initial segment of $\mathcal{A}$ is $\mathcal{A}$ itself.
>
> *Case 2: $\mathcal{A}$ is $\neg\mathcal{B}$, for some sentence $\mathcal{B}$.* Let $\mathcal{A}^*$ be any sentence that is an initial segment of $\mathcal{A}$. So $\mathcal{A}^*$ is $\neg\mathcal{B}^*$, for some sentence $\mathcal{B}^*$. Evidently, $\mathcal{B}^*$ is an initial segment of $\mathcal{B}$. But $\mathcal{B}$ and $\mathcal{B}^*$ are both sentences that are shorter than $\mathcal{A}$; so by the induction hypothesis, $\mathcal{B}^*$ must just be $\mathcal{B}$ itself. Hence $\mathcal{A}^*$ is $\mathcal{A}$ itself, and so $\mathcal{A}$ has the induction property.
>
> *Case 3: $\mathcal{A}$ is $(\mathcal{B} \triangleleft \mathcal{C})$, for some sentences $\mathcal{B}$ and $\mathcal{C}$ and some two-place connective $\triangleleft$.* Let $\mathcal{A}*$ be any sentence that is an initial segment of $\mathcal{A}$. Evidently, $\mathcal{A}^*$ starts with a '('; so, by the recursion rules, $\mathcal{A}^*$ must be $(\mathcal{B}^* \blacktriangleleft \mathcal{C}^*)$, for some sentences $\mathcal{B}^*$ and $\mathcal{C}^*$ and some two-place connective $\blacktriangleleft$. Now, exactly as in Case 2, we can see that $\mathcal{B}^*$ must just be $\mathcal{B}^*$ itself (since $\mathcal{B}^*$ is an initial segment of $\mathcal{B}$). Consequently, $\triangleleft$ and $\blacktriangleleft$ must be the same connective. And so, by a similar argument, $\mathcal{C}^*$ must just be $\mathcal{C}$ itself (since $\mathcal{C}^*$ is an initial segment of $\mathcal{C}$). Hence $\mathcal{A}^*$ is $\mathcal{A}$ itself, and so $\mathcal{A}$ has the induction property.

This completes the strong induction on length.   ∎

Considered on its own, this result is not very thrilling. But I called it a 'lemma' because it will be used to prove a result that is sufficiently important to be called a 'theorem', and to have its own name.

Before I state and prove the Theorem, I shall try to motivate it. When I introduced the idea of the MAIN LOGICAL OPERATOR of a sentence (in forall$x$:Cambridge §6), I described it as the symbol that 'was introduced *last,* when constructing the sentence' by the recursion rules. However, I gave you no guarantee that there was only *one* way to construct a sentence using those rules. Consequently, I gave you no guarantee that the main logical operator is *unique.* And, at that point in your lives, we were not in a position to guarantee this result. Now, though, we are; and here is that overdue guarantee.

Theorem 1.7. Unique Readability Theorem. For every sentence $\mathcal{A}$, exactly one of the following obtains:

> 1. $\mathcal{A}$ is an atomic sentence; or
> 2. there is a unique sentence, $\mathcal{B}$, such that $\mathcal{A}$ is $\neg\mathcal{B}$; or
> 3. there are unique sentences, $\mathcal{B}$ and $\mathcal{C}$, and a unique two-place connective, $\triangleleft$, such that $\mathcal{A}$ is $(\mathcal{B} \triangleleft \mathcal{C})$.

Proof of Theorem 1.7. Let $\mathcal{A}$ be any sentence. It is clear that *at most* one of (1), (2) or (3) obtains, since each requires that $\mathcal{A}$ start with a different symbol. It remains to show that *at least* one of them obtains. There are three cases to consider:

> *Case 1: $\mathcal{A}$ is atomic.* i.e. case (1) obtains.

Case 2: $\mathcal{A}$ is ¬$\mathcal{B}$, for some sentence $\mathcal{B}$. Then it is clear that there is exactly *one* sentence that is formed by removing the negation, i.e. $\mathcal{B}$ is unique. So case (2) obtains.

Case 3: $\mathcal{A}$ is ($\mathcal{B} \lhd \mathcal{C}$), for some sentences $\mathcal{B}$ and $\mathcal{C}$ and some two-place connective $\lhd$. To establish that case (3) obtains, we need to show that this decomposition is *unique*. So suppose that there are also sentences, $\mathcal{B}^*$ and $\mathcal{C}^*$, and a two-place connective, $\blacktriangleleft$, such that $\mathcal{A}$ is ($\mathcal{B}^* \blacktriangleleft \mathcal{C}^*$). We can use exactly the same technique as in the proof of Case 3 of Lemma 1.6 to show that $\mathcal{B}$ is $\mathcal{B}^*$, that $\lhd$ is $\blacktriangleleft$ and that $\mathcal{C}$ is $\mathcal{C}^*$; the only difference is that we now invoke Lemma 1.6 itself (rather than an induction hypothesis). Hence, case (3) obtains.   ■

This uniqueness really matters. It shows that there can be no syntactic ambiguity in TFL, so that each sentence admits of one (and only one) reading. Contrast one of the most famous newspaper headlines of all time: '8th Army push bottles up German rear'.

Practice exercises

A. Fill in the details in Case 3 of the proof of Lemma 1.6, and in Case 3 of the proof of Theorem 1.7.

B. Prove the following, by induction on length: for any sentence $\mathcal{A}$, every valuation which assigns truth values to all of the atomic sentences in $\mathcal{A}$ assigns a *unique* truth value to $\mathcal{A}$.

C. For any sentence $\mathcal{A}$, let:

- $a_{\mathcal{A}}$ be the number of instances of atomic sentences in $\mathcal{A}$;
- $b_{\mathcal{A}}$ be the number of instances of '(' in $\mathcal{A}$;
- $c_{\mathcal{A}}$ be the number of instances of two-place connectives in $\mathcal{A}$.

As an example: if $\mathcal{A}$ is '(¬$B \rightarrow (B \wedge C)$)', then $a_{\mathcal{A}} = 3$, $b_{\mathcal{A}} = 2$ and $c_{\mathcal{A}} = 2$. Prove the following, by induction on length: for any sentence $\mathcal{A}$: $a_{\mathcal{A}} - 1 = b_{\mathcal{A}} = c_{\mathcal{A}}$.

D. Suppose a sentence is built up by conjoining together all of the sentences $\mathcal{A}_1, \ldots, \mathcal{A}_n$, in any order. So the sentence might be, e.g.:

$$(\mathcal{A}_5 \wedge ((\mathcal{A}_3 \wedge \mathcal{A}_2) \wedge (\mathcal{A}_4 \wedge \mathcal{A}_1)))$$

Prove, by induction on the number of conjuncts, that the truth value of the resulting sentence (on any valuation which assigns a truth value to all of $\mathcal{A}_1, \ldots, \mathcal{A}_n$) will be the same, no matter how, exactly, they are conjoined. Then prove that the same is true if we consider disjunction rather than conjunction. Which notational conventions that were introduced in forall*x*:Cambridge are vindicated by these proofs?

Substitution 2

In this chapter, we shall put our new familiarity with induction to good work, and prove a cluster of results that connect with the idea of substituting one sentence in for another.

Before going further, we should state the idea of 'substitution' a bit more precisely. Where $\mathcal{A}$, C and S are sentences, $\mathcal{A}[S/C]$ is the result of searching through $\mathcal{A}$ and finding and replacing every occurrence of C with S. As an example, when $\mathcal{A}$ is '$(Q \wedge (\neg P \vee (Q \leftrightarrow R)))$', when C is 'Q' and when S is '$\neg (D \rightarrow B)$', then $\mathcal{A}[S/C]$ is '$(\neg (D \rightarrow B) \wedge (\neg P \vee (\neg (D \rightarrow B) \leftrightarrow R)))$'.

All of the results in this chapter employ this surprisingly powerful idea.

2.1 Two crucial lemmas about substitution

Given your familiarity with TFL, here is a line of thought which should strike you as plausible. Suppose you have a valid argument in TFL. Now, run through that argument, uniformly substituting every instance of some atomic sentence that occurs within that argument, with some *other* sentence. The resulting argument is still valid.

It is worth realising that this line of thought is *not* generally plausible in natural languages. For example, here is a valid argument in English:

- Isaac is an eft; so Isaac is a newt.

Its validity, plausibly, connects with the *meanings* of its component words. But recall that the semantics of TFL is purely extensional, and knows nothing about meanings. Its only dimension of freedom is over how we assign truth and falsity to components of sentences. And uniform substitution – of the sort mentioned in the previous paragraph – cannot, in a relevant sense, *increase* our options for assigning truth values to such components.

But to convince ourselves that this line of thought is correct, we of course require a rigorous proof. Here it is.

Lemma 2.1. For any atomic sentence C and any sentences $\mathcal{A}_1, \ldots, \mathcal{A}_{n+1}, S$:
if $\mathcal{A}_1, \ldots, \mathcal{A}_n \vDash \mathcal{A}_{n+1}$, then $\mathcal{A}_1[S/C], \ldots, \mathcal{A}_n[S/C] \vDash \mathcal{A}_{n+1}[S/C]$.

Proof. Suppose $\mathcal{A}_1, \ldots, \mathcal{A}_n \vDash \mathcal{A}_{n+1}$. Suppose also that v is a valuation which makes all of $\mathcal{A}_1[S/C], \ldots, \mathcal{A}_n[S/C]$ true. Let w be a valuation which differs from v (if at all) only by assigning a value to C as follows: w makes C true iff v makes S true. I now make a claim:

Claim: for any sentence $\mathcal{A}$: w makes $\mathcal{A}$ true iff v makes $\mathcal{A}[S/C]$ true.

I shall prove the Claim by induction on the length of $\mathcal{A}$. Suppose, for induction, that the Claim is true of every shorter sentence. Now there are three cases:

> *Case 1: $\mathcal{A}$ is atomic.* Then either $\mathcal{A}$ is $\mathcal{C}$, in which case $\mathcal{A}[S/c]$ is just S; or $\mathcal{A}[S/c]$ is just $\mathcal{A}$. Either way, the Claim holds of $\mathcal{A}$.
>
> *Case 2: $\mathcal{A}$ is $\neg\mathcal{B}$, for some sentence $\mathcal{B}$.* Then, by assumption, w makes $\mathcal{B}$ true iff v makes $\mathcal{B}[S/c]$ true. By considering the truth-table for '$\neg$', it follows that w makes $\neg\mathcal{B}$ true iff v makes $\neg\mathcal{B}[S/c]$ true.
>
> *Case 3: $\mathcal{A}$ is $(\mathcal{B} \lhd \mathcal{C})$, for some sentences $\mathcal{B}$ and $\mathcal{C}$ and some two-place connective $\lhd$.* This is much like Case 3; I leave the details as an exercise.

This proves the Claim, by induction. It follows from the claim that w makes all of $\mathcal{A}_1, \ldots, \mathcal{A}_n$ true. And, since $\mathcal{A}_1, \ldots, \mathcal{A}_n \vDash \mathcal{A}_{n+1}$, we see that w makes $\mathcal{A}_{n+1}$ true. So, appealing to the Claim again, v makes $\mathcal{A}_{n+1}[S/c]$ true.   ∎

Let's see this Lemma in action. Consider these three claims:

$$(\neg A \to (A \lor B)) \vDash (A \lor B)$$
$$(\neg\neg A \to (\neg A \lor \neg B)) \vDash (\neg A \lor \neg B)$$
$$(\neg(P \leftrightarrow \neg Q) \to ((P \leftrightarrow \neg Q) \lor \neg(B \land (C \leftrightarrow D)))) \vDash ((P \leftrightarrow \neg Q) \lor \neg(B \land (C \leftrightarrow D)))$$

The first of these is easy to check using a truth table; the second is marginally harder; and the third is a pain! But if we recognise that the second and third are the result of uniform substitution into the first, then we can use Lemma 2.1 to save ourselves a lot of work.

There is a deeper point here. Lemma 2.1 vindicates the practice of describing an *entire argument structure* as valid. For example, we can say that the following is a valid *pattern* of entailment:

$$(\neg\mathcal{A} \to (\mathcal{A} \lor \mathcal{B})) \vDash (\mathcal{A} \lor \mathcal{B})$$

We know already that at least one instance of this pattern of entailment is valid. Lemma 2.1 then tells us that the entailment will hold, *regardless* of what sentences we use for $\mathcal{A}$ and $\mathcal{B}$. Lemma 2.1 therefore underscores the purely structural nature of entailments, in our semantics for TFL.

Lemma 2.1 concerns tautological *entailment*. We shall also want to have a result to hand which concerns tautological *equivalence*. It will help us if we introduce the symbol '$\dashv\vDash$' to symbolise tautological equivalence. So we write:

$$\mathcal{A} \dashv\vDash \mathcal{B}$$

to symbolise that $\mathcal{A}$ and $\mathcal{B}$ are tautologically equivalent; or, in other words, that $\mathcal{A} \vDash \mathcal{B}$ and $\mathcal{B} \vDash \mathcal{A}$. Using this notation, we can prove a nifty little result:

Lemma 2.2. For any sentences $\mathcal{A}$, $\mathcal{C}$ and $\mathcal{S}$: if $\mathcal{C} \dashv\vDash \mathcal{S}$, then $\mathcal{A} \dashv\vDash \mathcal{A}[S/c]$.

Proof. By an induction on the length of $\mathcal{A}$; I leave this as an exercise.   ∎

In a slogan: substitution of tautological equivalents preserves tautological equivalence. We shall rely upon this frequently, in what follows.

2.2   Interpolation

The next result uses the idea of substitution to prove something rather deep. When working with an argument, one can get the idea that there is an 'active core' to the argument, which should be distinguished from some of the irrelevant surrounding matter. The next result gives some formal content to this claim

I shall start with a little result which might be put as telling us the following: if I am to draw a conclusion *from* some claim, there had better be something in the original claim that my conclusion builds upon. More precisely:

Lemma 2.3. Suppose $\mathcal{A} \vDash \mathcal{B}$, that $\mathcal{A}$ is not a contradiction, and that $\mathcal{B}$ is not a tautology. Then at least one atomic sentence occurs in both $\mathcal{A}$ and $\mathcal{B}$.

Proof. Supposing $\mathcal{A}$ is not a contradiction, there is a valuation, v, that makes $\mathcal{A}$ true. Supposing $\mathcal{B}$ is not a tautology, there is a valuation, w, that makes $\mathcal{B}$ false. Supposing also $\mathcal{A}$ and $\mathcal{B}$ have no atomic sentences in common, then there is a valuation, x, which assigns the same values as v to all atomic sentence letters in $\mathcal{A}$, and the same values as w to all atomic sentence letters in $\mathcal{B}$. By hypothesis, x makes $\mathcal{A}$ true and $\mathcal{B}$ false. So $\mathcal{A} \nvDash \mathcal{B}$. ∎

Our main result is going to employ a (fairly lengthy) induction to extend Lemma 2.3 rather radically. First, we need a definition. Suppose that $\mathcal{A} \vDash \mathcal{B}$. Then an INTERPOLANT linking $\mathcal{A}$ to $\mathcal{B}$ is any sentence $\mathcal{I}$ meeting all three of these conditions:

(INT1) $\mathcal{A} \vDash \mathcal{I}$
(INT2) $\mathcal{I} \vDash \mathcal{B}$
(INT3) every atomic sentence in $\mathcal{I}$ occurs both in $\mathcal{A}$ and in $\mathcal{B}$

Our result guarantees the existence of interpolants.

Theorem 2.4. Interpolation Theorem. Suppose $\mathcal{A} \vDash \mathcal{B}$, that $\mathcal{A}$ is not a contradiction, and that $\mathcal{B}$ is not a tautology. Then there is an interpolant linking $\mathcal{A}$ to $\mathcal{B}$.

Proof. Our proof consists in an (ordinary) induction on the number of atomic sentences that occur in $\mathcal{A}$ but not $\mathcal{B}$.

In the base case, there are no atomic sentences in $\mathcal{A}$ that are not also in $\mathcal{B}$. Hence we can let $\mathcal{A}$ itself be our interpolant, since $\mathcal{A} \vDash \mathcal{A}$ vacuously, and $\mathcal{A} \vDash \mathcal{B}$ by hypothesis.

Now fix n, and suppose (for induction) that we can find an interpolant whenever there are n atomic sentences that occur in the 'premise' but not the 'conclusion' of the entailment. Additionally, suppose that $\mathcal{A} \vDash \mathcal{B}$, that $\mathcal{A}$ is not a contradiction, that $\mathcal{B}$ is not a tautology, and that there are $n + 1$ atomic sentences that occur in $\mathcal{A}$ but not in $\mathcal{B}$. Lemma 2.3 tells us that there is at least one atomic sentence, $\mathcal{C}$, common to both $\mathcal{A}$ and $\mathcal{B}$. Now, where:

- $\mathcal{D}$ is any sentence that occurs in $\mathcal{A}$ but not in $\mathcal{B}$
- $\top$ is a tautology which contains no atomic sentence apart from $\mathcal{C}$, e.g. $(\mathcal{C} \to \mathcal{C})$
- $\bot$ is a contradiction which contains no atomic sentence apart from $\mathcal{C}$, e.g. $(\mathcal{C} \land \neg \mathcal{C})$

I make the following claim:

    Claim 1: $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}]) \vDash \mathcal{B}$

To prove Claim 1, let v be any valuation that makes $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}])$ true. If v makes $\mathcal{A}[\top/\mathcal{D}]$ true, let w differ from v just by making $\mathcal{D}$ true; otherwise, let w differ from v just by making $\mathcal{D}$ false. Either way, w makes $\mathcal{A}$ true, so w makes $\mathcal{B}$ true too (since $\mathcal{A} \vDash \mathcal{B}$). Finally, since $\mathcal{D}$ does not occur in $\mathcal{B}$, v makes $\mathcal{B}$ true too. This establishes the Claim.

Next, because there were $n + 1$ atomic sentences in $\mathcal{A}$ but not $\mathcal{B}$, and $\mathcal{D}$ was one of these atomic sentences, there must be exactly n atomic sentences in $\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}]$ but not $\mathcal{B}$. Hence, by Claim 1 and the induction hypothesis, we can find an interpolant, $\mathcal{I}$, linking $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}])$ to $\mathcal{B}$, i.e.:

1. $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}]) \vDash \mathcal{I}$
2. $\mathcal{I} \vDash \mathcal{B}$
3. every atomic sentence in $\mathcal{I}$ occurs in both $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}])$ and in $\mathcal{B}$

I shall now show that $\mathcal{I}$ is an interpolant linking $\mathcal{A}$ itself to $\mathcal{B}$. Point (2), above, tells us that $\mathcal{I} \vDash \mathcal{B}$. Moreover, given point (3), it is obvious that every atomic sentence in $\mathcal{I}$ occurs in both $\mathcal{A}$ and in $\mathcal{B}$. So it remains to show that $\mathcal{A} \vDash \mathcal{I}$. Given that $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}]) \vDash \mathcal{I}$, it suffices, by the transitivity of entailment, to prove the following claim:

Claim 2. $\mathcal{A} \vDash (\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}])$

To prove Claim 2, let v be any valuation that makes $\mathcal{A}$ true. If v makes $\mathcal{D}$ true, then since $\top$ is a tautology, v makes $\mathcal{A}[\top/\mathcal{D}]$ true. If v makes $\mathcal{D}$ false, then v makes $\mathcal{A}[\bot/\mathcal{D}]$ true. Either way, any valuation that makes $\mathcal{A}$ true makes $(\mathcal{A}[\top/\mathcal{D}] \lor \mathcal{A}[\bot/\mathcal{D}])$ true. This establishes Claim 2.

Hence $\mathcal{I}$ is an interpolant linking $\mathcal{A}$ to $\mathcal{B}$. This completes the proof of the induction case. The entire Theorem now follows by induction. ∎

This is easily the longest proof that we have seen so far. So it is worth pausing, to make sure both that you understand every step in the proof, and also how the proof is put together.

In fact, this proof merits particular scrutiny since, implicitly, the proof does more than *just* prove the Theorem. The Interpolation Theorem states that, when $\mathcal{A} \vDash \mathcal{B}$, an interpolant *exists*. If this were all we knew, then we might have no idea about how to *find* an interpolant. But, viewed correctly, the proof yields an algorithm for finding an interpolant. An ALGORITHM is a step-by-step process, followed 'thoughtlessly' (i.e. without any need for insight, just by implementing very simple little rules), that is guaranteed to terminate in finitely many steps. I shall describe the algorithm for finding an interpolant; and having done so, it should be clear how it relates to the proof.

Given that $\mathcal{A} \vDash \mathcal{B}$, linking $\mathcal{A}$ to $\mathcal{B}$ as follows:

Step 1: Designate $\mathcal{A}$ as the *input* sentence

Step 2: Let $\mathcal{D}$ be the first atomic sentence (alphabetically) that appears in the input sentence, but that does not appear in $\mathcal{B}$. Rewrite the input sentence twice, once uniformly substituting a tautology in place of $\mathcal{D}$ and once uniformly substituting a contradiction in place of $\mathcal{D}$. Let the disjunction of these two rewritten sentences be your new input sentence.

Step 3: Repeat step 2, until there are no atomic sentences that appear in the input sentence but not in $\mathcal{B}$. At that point, stop. The last sentence you wrote down is an interpolant linking $\mathcal{A}$ to $\mathcal{B}$.

It is worth seeing this in action. Suppose we start with the entailment:

$$((A \land (\neg B \lor C)) \land (\neg C \land D)) \vDash ((B \lor D) \land (\neg B \lor E))$$

Running the algorithm, our initial input sentence is:

$$((A \land (\neg B \lor C)) \land (\neg C \land D))$$

We would first eliminate A, obtaining (using '$\top$' to abbreviate some tautology containing 'B', and '$\bot$' to abbreviate some contradiction containing 'B'):

$$(((\top \land (\neg B \lor C)) \land (\neg C \land D)) \lor ((\bot \land (\neg B \lor C)) \land (\neg C \land D)))$$

This would become our new input sentence; from which we would eliminate C:

$$((((\top \land (\neg B \lor \top)) \land (\neg \top \land D)) \lor ((\bot \land (\neg B \lor \top)) \land (\neg \top \land D))) \lor$$
$$(((\top \land (\neg B \lor \bot)) \land (\neg \bot \land D)) \lor ((\bot \land (\neg B \lor \bot)) \land (\neg \bot \land D))))$$

And this is our interpolant! But, as the example illustrates, the output of this algorithm can be quite a mouthful. To simplify it, we can apply some 'clean-up' rules. Where '$\top$' abbreviates any tautology, and '$\bot$' abbreviates any contradiction, note, for example, that:

$$\top \lor \mathcal{A} \dashv\vDash \top$$
$$\bot \land \mathcal{A} \dashv\vDash \bot$$
$$\neg \top \dashv\vDash \bot$$
$$\neg \bot \dashv\vDash \top$$

Using these equivalences and Lemma 2.2, we can simplify the subsentences of our interpolant.[1] First, applying the rules for negation:

$$((((\top \land (\neg B \lor \top)) \land (\bot \land D)) \lor ((\bot \land (\neg B \lor \top)) \land (\bot \land D))) \lor$$
$$(((\top \land (\neg B \lor \bot)) \land (\top \land D)) \lor ((\bot \land (\neg B \lor \bot)) \land (\top \land D))))$$

Next, applying the rules for conjunction and disjunction:

$$((((\top \land \top) \land \bot) \lor (\bot \land \bot)) \lor (((\top \land \neg B) \land D) \lor (\bot \land D)))$$

And, then, applying them again and again, we get, successively:

$$((\bot \lor \bot) \lor ((\neg B \land D) \lor \bot))$$
$$(\bot \lor (\neg B \land D))$$
$$(\neg B \land D)$$

Which is much simpler!

---

[1]Recall that a sentence $\mathcal{A}$ is a SUBSENTENCE of $\mathcal{B}$ iff $\mathcal{A}$ is a part of $\mathcal{B}$. We here read 'part' in a generous sense, so that every sentence is a part of itself.

2.3   Duality

The final result of this chapter will use substitution to explore a very powerful idea. To explain the idea, recall the following equivalences, known as the De Morgan Laws:

$$(\mathcal{A} \wedge \mathcal{B}) \Dashv\vDash \neg(\neg\mathcal{A} \vee \neg\mathcal{B})$$
$$(\mathcal{A} \vee \mathcal{B}) \Dashv\vDash \neg(\neg\mathcal{A} \wedge \neg\mathcal{B})$$

These indicate that there is a tight 'symmetry' between conjunction and disjunction. The notion of *duality* is a generalisation of this tight 'symmetry'. And in fact, it is worth introducing it with complete generality at the outset.

In TFL, our only primitive connectives are one-place (i.e. '¬') and two-place (i.e. '∧', '∨', '→' and '↔'). But nothing stops us from introducing three-, four-, or five-place connectives; or, more generally, n-place connectives, for any number n we like. (We shall, however, always assume that our connectives are *truth-functional*.)

So, let ◁ be an n-place connective. The DUAL of ◁, which I shall generically denote by underlining it, i.e. ◁, is defined to have a characteristic truth table as follows:

$$\underline{\triangleleft}(\mathcal{A}_1, \dots, \mathcal{A}_n) \Dashv\vDash \neg\triangleleft(\neg\mathcal{A}_1, \dots, \neg\mathcal{A}_n)$$

This is the obvious generalisation of the De Morgan Laws to many-place connectives.[2] And it is immediate, on this definition, that '∧' is the dual of '∨', and that '∨' is the dual of '∧'. I encourage you also to check that '↑' is the dual of '↓', and vice versa; and that '¬' is its *own* dual.

Inspired by these examples, we might correctly conjecture the following:

Lemma 2.5. The dual of a dual is always the original connective; that is, $\underline{\underline{\triangleleft}}$ is the same connective as ◁.

Proof. I start by noting:

$$
\begin{array}{ll}
\underline{\triangleleft}(\mathcal{A}_1, \dots, \mathcal{A}_n) \Dashv\vDash \neg\triangleleft(\neg\mathcal{A}_1, \dots, \neg\mathcal{A}_n) & \text{by definition} \\[4pt]
\underline{\triangleleft}(\neg\mathcal{A}_1, \dots, \neg\mathcal{A}_n) \Dashv\vDash \neg\triangleleft(\neg\neg\mathcal{A}_1, \dots, \neg\neg\mathcal{A}_n) & \text{by Lemma 2.1} \\[4pt]
\neg\underline{\triangleleft}(\neg\mathcal{A}_1, \dots, \neg\mathcal{A}_n) \Dashv\vDash \neg\neg\triangleleft(\neg\neg\mathcal{A}_1, \dots, \neg\neg\mathcal{A}_n) & \text{by the definition of '¬'} \\[4pt]
\neg\underline{\triangleleft}(\neg\mathcal{A}_1, \dots, \neg\mathcal{A}_n) \Dashv\vDash \triangleleft(\mathcal{A}_1, \dots, \mathcal{A}_n) & \text{by Lemma 2.2}
\end{array}
$$

Next, observe that, by definition:

$$\underline{\underline{\triangleleft}}(\mathcal{A}_1, \dots, \mathcal{A}_n) \Dashv\vDash \neg\underline{\triangleleft}(\neg\mathcal{A}_1, \dots, \neg\mathcal{A}_n)$$

Hence, by the transitivity of entailment:

$$\underline{\underline{\triangleleft}}(\mathcal{A}_1, \dots, \mathcal{A}_n) \Dashv\vDash \triangleleft(\mathcal{A}_1, \dots, \mathcal{A}_n)$$

as required.   ∎

---

[2]Hence some books talk of the 'De Morgan dual', rather than simply of the 'dual'.

We shall use the same kind of proof-technique in what follows, to great effect.

First, we need a little more notation, connecting duality with negation. Given any sentence, $\mathcal{A}$, we say $\underline{\mathcal{A}}$ is the result of replacing *every* connective in $\mathcal{A}$ with its dual, and $\overline{\mathcal{A}}$ is the result of replacing every atomic sentence in $\mathcal{A}$ with the negation of that atomic sentence. So, when $\mathcal{A}$ is the sentence '$(A \vee (\neg B \wedge C))$', $\underline{\mathcal{A}}$ is '$(A \wedge (\neg B \vee C))$' and $\overline{\mathcal{A}}$ is '$(\neg A \vee (\neg\neg B \wedge \neg C))$'.

Theorem 2.6. $\underline{\mathcal{A}} \vDash\!\!\dashv \neg\overline{\mathcal{A}}$, for any sentence $\mathcal{A}$.

Proof. Let $\mathcal{A}$ be any sentence; we suppose, for induction on length, that every shorter sentence $\mathcal{B}$ is such that $\underline{\mathcal{B}} \vDash\!\!\dashv \neg\overline{\mathcal{B}}$. There are two cases to consider:

Case 1: $\mathcal{A}$ is atomic. Then $\underline{\mathcal{A}}$ is just $\mathcal{A}$ itself; and $\neg\overline{\mathcal{A}}$ is just $\neg\neg\mathcal{A}$; and we know that $\mathcal{A} \vDash\!\!\dashv \neg\neg\mathcal{A}$.

Case 2: $\mathcal{A}$ is $\triangleleft(\mathcal{B}_1, \dots, \mathcal{B}_n)$, for some n-place connective $\triangleleft$. So by hypothesis, $\underline{\mathcal{B}_i} \vDash\!\!\dashv \neg\overline{\mathcal{B}_i}$ for each $\mathcal{B}_i$. Now observe that:

$$\underline{\triangleleft(\mathcal{B}_1, \dots, \mathcal{B}_n)} \vDash\!\!\dashv \underline{\triangleleft}(\underline{\mathcal{B}_1}, \dots, \underline{\mathcal{B}_n}) \qquad \text{by definition}$$
$$\vDash\!\!\dashv \underline{\triangleleft}(\neg\overline{\mathcal{B}_1}, \dots, \neg\overline{\mathcal{B}_n}) \qquad \text{by Lemma 2.2}$$
$$\vDash\!\!\dashv \neg\triangleleft(\neg\neg\overline{\mathcal{B}_1}, \dots, \neg\neg\overline{\mathcal{B}_n}) \qquad \text{by definition of } \underline{\triangleleft}$$
$$\vDash\!\!\dashv \neg\triangleleft(\overline{\mathcal{B}_1}, \dots, \overline{\mathcal{B}_n}) \qquad \text{by Lemma 2.2}$$
$$\vDash\!\!\dashv \neg\overline{\triangleleft(\mathcal{B}_1, \dots, \mathcal{B}_n)} \qquad \text{by definition}$$

and hence, given the transitivity of entailment, $\underline{\mathcal{A}} \vDash\!\!\dashv \neg\overline{\mathcal{A}}$.

This completes the induction on length.   ∎

And we can now obtain a very elegant result:

Theorem 2.7. $\mathcal{A} \vDash \mathcal{B}$ iff $\underline{\mathcal{B}} \vDash \underline{\mathcal{A}}$, for any sentences $\mathcal{A}$ and $\mathcal{B}$. Hence, in particular, $\mathcal{A} \vDash\!\!\dashv \mathcal{B}$ iff $\underline{\mathcal{A}} \vDash\!\!\dashv \underline{\mathcal{B}}$.

Proof. Left to right. Suppose $\mathcal{A} \vDash \mathcal{B}$. Then $\overline{\mathcal{A}} \vDash \overline{\mathcal{B}}$, by Lemma 2.1. So $\neg\overline{\mathcal{B}} \vDash \neg\overline{\mathcal{A}}$. So $\underline{\mathcal{B}} \vDash \underline{\mathcal{A}}$, by Theorem 2.6 and Lemma 2.2.

Right to left. Suppose $\underline{\mathcal{B}} \vDash \underline{\mathcal{A}}$. Then, by the preceding, $\underline{\underline{\mathcal{A}}} \vDash \underline{\underline{\mathcal{B}}}$. Hence $\mathcal{A} \vDash \mathcal{B}$, by Lemmas 2.5 and 2.2.   ∎

This ends my theorem-proving. But to close the chapter, I want to offer some sketch some reasons for thinking that duality is pretty great.

Illustration 1. You may have already noticed the following equivalences:

$$(\mathcal{A} \wedge (\mathcal{B} \vee \mathcal{C})) \vDash\!\!\dashv ((\mathcal{A} \wedge \mathcal{B}) \vee (\mathcal{A} \wedge \mathcal{C}))$$
$$((\mathcal{A} \vee \mathcal{B}) \wedge \mathcal{C}) \vDash\!\!\dashv ((\mathcal{A} \wedge \mathcal{C}) \vee (\mathcal{B} \wedge \mathcal{C}))$$

If not, then you can test these using a truth table. But, having done so, Theorem 2.7 immediately yields two further equivalences for free:

$$(\mathcal{A} \vee (\mathcal{B} \wedge \mathcal{C})) \vDash\!\!\dashv ((\mathcal{A} \vee \mathcal{B}) \wedge (\mathcal{A} \vee \mathcal{C}))$$
$$((\mathcal{A} \wedge \mathcal{B}) \vee \mathcal{C}) \vDash\!\!\dashv ((\mathcal{A} \vee \mathcal{C}) \wedge (\mathcal{B} \vee \mathcal{C}))$$

These four equivalences are the Distribution Laws. We shall use them repeatedly in the next chapter.

Illustration 2. Here is a playful way to think about negation, due to Frank Plumpton Ramsey.[3] Suppose that we allow ourselves to represent a sentence's negation by flipping the entire sentence upside-down, rather than by prefixing the sentence with '¬'. With this notation, and supposing '♡' to be an n-place connective, we would have:

$$\underline{\heartsuit}(\mathcal{A}_1, ..., \mathcal{A}_n) \dashv\vDash \neg\heartsuit(\neg\mathcal{A}_1, ..., \neg\mathcal{A}_n) \qquad \text{the original definition}$$
$$\dashv\vDash \heartsuit(\text{⅁}^\text{I}, ..., \text{⅁}^\text{u}) \qquad \text{replacing the interior negations}$$
$$\dashv\vDash \triangle(\mathcal{A}_1\text{`}...\text{`}\mathcal{A}_n) \qquad \text{replacing the outer negation}$$

and with a little imagination concerning punctuation (commas and brackets), we see that the same effect could be achieved more simply by merely inverting the connective itself, i.e. we would represent the dual of '♡' by '⌂'. A 'fully perspicuous' notation would, then, represent connectives and their duals by symbols which are mirror-images of each other. But that is exactly how we *do* represent conjunction and disjunction![4] Note, too, that if we adopted this notation, then Lemma 2.5 would be *visually* immediate. And this suggests (correctly) that the success of our Ramsey-inspired notation system depends only upon Lemma 2.2 and the equivalence: $\neg\neg\mathcal{A} \dashv\vDash \mathcal{A}$. Once we realise this, we can see that *all* of our results about duality are extremely general: they hold for any extensional language which respects the equivalence $\neg\neg\mathcal{A} \dashv\vDash \mathcal{A}$ and in which Lemmas 2.1 and 2.2 hold.

Illustration 3. An immediate consequence of our results so far is

$$\mathcal{A} \dashv\vDash \neg\overline{\mathcal{A}}$$

I leave the proof of this consequence as an exercise. But the consequence prompts a subtle thought. Typically, when I hear you say 'London is the capital of the UK', or whatever, I take you to mean *that* London is the capital of the UK, and to be *asserting* this. What if, instead, I took you to mean that London is *not* the capital of the UK, and to be *rejecting* this? More generally, where before I took you to be *asserting* that $\mathcal{A}$, I now take you to be *rejecting* that $\overline{\mathcal{A}}$. Would anything be wrong with my reinterpretation? If not: why should I assume that you mean *and*, rather than *or*, by the word 'and'?[5]

Practice exercises

A. Complete the inductions on length mentioned in the proofs of Lemmas 2.1 and 2.2.

B. Prove that $\mathcal{A} \dashv\vDash \neg\overline{\mathcal{A}}$, for any sentence $\mathcal{A}$.

---

[3]'Facts and Propositions' (1927, *Proceedings of the Aristotelian Society Supplementary Volume* 7.1).

[4]There is something, then, to be said for representing negation with a horizontally symmetric '−'. But I have no great ideas for how to represent (bi)conditionals.

[5]For discussion of all this, see Gerald Massey, 'The Indeterminacy of Translation', pp.329ff. (1992, *Philosophical Topics* 20.1) and Hilary Putnam, 'Replies', pp.399–402 (1992, *Philosophical Topics* 20.1).

Normal forms 3

In this chapter, I prove two *normal form* theorems for TFL. These guarantee that, for any sentence, there is a tautologically equivalent sentence in some canonical normal form. Moreover, I shall give methods for finding these normal-form equivalents.

3.1 Disjunctive Normal Form

Say that a sentence is in DISJUNCTIVE NORMAL FORM *iff* it meets all of the following conditions:

(DNF1) No connectives occur in the sentence other than negations, conjunctions and disjunctions;

(DNF2) Every occurrence of negation has minimal scope (i.e. any '¬' is immediately followed by an atomic sentence);

(DNF3) No disjunction occurs within the scope of any conjunction.

(For a reminder of the definition of the SCOPE of a connective, see forall*x*:Cambridge §6.) So, here are are some sentences in disjunctive normal form:

$$A$$
$$(A \land B) \lor (A \land \neg B)$$
$$(A \land B) \lor (A \land B \land C \land \neg D \land \neg E)$$
$$A \lor (C \land \neg P_{234} \land P_{233} \land Q) \lor \neg B$$

Note that I have here broken one of the maxims of this book (see §1.4) and *temporarily* allowed myself to employ the relaxed bracketing-conventions that allow conjunctions and disjunctions to be of arbitrary length (see forall*x*:Cambridge §10.3). These conventions make it easier to see when a sentence is in disjunctive normal form. I shall continue to help myself to these relaxed conventions, without further comment, until §3.3.

To further illustrate the idea of disjunctive normal form, I shall introduce some more notation. I write '$\pm\mathcal{A}$' to indicate that $\mathcal{A}$ is an atomic sentence which may or may not be prefaced with an occurrence of negation. Then a sentence in disjunctive normal form has the following shape:

$$\left(\pm\mathcal{A}_1 \land ... \land \pm\mathcal{A}_i \right) \lor \left(\pm\mathcal{A}_{i+1} \land ... \land \pm\mathcal{A}_j \right) \lor ... \lor \left(\pm\mathcal{A}_{m+1} \land ... \land \pm\mathcal{A}_n \right)$$

We now know what it is for a sentence to be in disjunctive normal form. The result that we are aiming at is:

Theorem 3.1. Disjunctive Normal Form Theorem. For any sentence, there is a tautologically equivalent sentence in disjunctive normal form.

Henceforth, I shall abbreviate 'Disjunctive Normal Form' by 'DNF'.

3.2 Proof of DNF Theorem via truth tables

My first proof of the DNF Theorem employs truth tables. I shall first illustrate the technique for finding an equivalent sentence in DNF, and then turn this illustration into a rigorous proof.

Let's suppose we have some sentence, S, which contains three atomic sentences, 'A', 'B' and 'C'. The very first thing to do is fill out a complete truth table for S. Maybe we end up with this:

A	B	C	S
T	T	T	T
T	T	F	F
T	F	T	T
T	F	F	F
F	T	T	F
F	T	F	F
F	F	T	T
F	F	F	T

As it happens, S is true on four lines of its truth table, namely lines 1, 3, 7 and 8. Corresponding to each of those lines, I shall write down four sentences, whose only connectives are negations and conjunctions, where every negation has minimal scope:

- '$A \land B \land C$'   which is true on line 1 (and only then)
- '$A \land \neg B \land C$'   which is true on line 3 (and only then)
- '$\neg A \land \neg B \land C$'   which is true on line 7 (and only then)
- '$\neg A \land \neg B \land \neg C$'   which is true on line 8 (and only then)

But if I now disjoin all of these conjunctions, like so:

$$(A \land B \land C) \lor (A \land \neg B \land C) \lor (\neg A \land \neg B \land C) \lor (\neg A \land \neg B \land \neg C)$$

I have a sentence in DNF which is true on exactly those lines where one of the disjuncts is true, i.e. it is true on (and only on) lines 1, 3, 7, and 8. So this sentence has exactly the same truth table as S. So I have a sentence in DNF that is tautologically equivalent to S. Which is exactly what I required!

Now, the strategy that I just adopted did not depend on the specifics of S; it is perfectly general. Consequently, we can use it to obtain a simple proof of the DNF Theorem.

Proof of DNF Theorem via truth tables. Pick any arbitrary sentence, S, and let $\mathcal{A}_1, \ldots, \mathcal{A}_n$ be the atomic sentences that occur in S. To obtain a sentence in DNF that is tautologically equivalent S, we consider S's truth table. There are two cases to consider:

Case 1: S is false on every line of its truth table. Then, S is a contradiction. In that
    case, the contradiction $(\mathcal{A}_1 \wedge \neg \mathcal{A}_1) \dashv\vDash S$, and $(\mathcal{A}_1 \wedge \neg\mathcal{A}_1)$ is in DNF.

Case 2: S is true on at least one line of its truth table. For each line i of the truth table,
    let $\mathcal{B}_i$ be a conjunction of the form

$$(\pm\mathcal{A}_1 \wedge \ldots \wedge \pm\mathcal{A}_n)$$

where the following rules determine whether or not to include a negation in front
of each atomic sentence:

$$\mathcal{A}_m \text{ is a conjunct of } \mathcal{B}_i \text{ iff } \mathcal{A}_m \text{ is true on line } i$$
$$\neg\mathcal{A}_m \text{ is a conjunct of } \mathcal{B}_i \text{ iff } \mathcal{A}_m \text{ is false on line } i$$

Given these rules, a trivial proof by induction shows that $\mathcal{B}_i$ is true on (and only
on) line i of the truth table which considers all possible valuations of $\mathcal{A}_1, \ldots, \mathcal{A}_n$
(i.e. S's truth table).

Next, let $i_1, i_2, \ldots, i_m$ be the numbers of the lines of the truth table where S is *true*.
Now let $\mathcal{D}$ be the sentence:

$$\mathcal{B}_{i_1} \vee \mathcal{B}_{i_2} \vee \ldots \vee \mathcal{B}_{i_m}$$

Since S is true on at least one line of its truth table, $\mathcal{D}$ is indeed well-defined; and
in the limiting case where S is true on exactly one line of its truth table, $\mathcal{D}$ is just
$\mathcal{B}_{i_1}$, for some i_1.

By construction, $\mathcal{D}$ is in DNF. Moreover, by construction, for each line i of the truth
table: S is true on line i of the truth table *iff* one of $\mathcal{D}$'s disjuncts (namely, $\mathcal{B}_i$) is
true on, and only on, line i. (Again, this is shown by a trivial proof by induction.)
Hence S and $\mathcal{D}$ have the same truth table, and so are tautologically equivalent.

These two cases are exhaustive and, either way, we have a sentence in DNF that is tauto-
logically equivalent to S. ∎

So we have proved the DNF Theorem. Before I say any more, though, I should imme-
diately flag that I am hereby returning to the austere definition of a (TFL) sentence, ac-
cording to which we can assume that any conjunction has exactly two conjuncts, and any
disjunction has exactly two disjuncts.[1]

3.3   Proof of DNF Theorem via substitution

I now want to offer a second proof of the DNF Theorem.

At this point, you might reasonably ask: Why bother with *two* proofs? Certainly one
proof is sufficient to establish a result. But there are several possible reasons for offering
new proofs of known results. For example:

- a new proof might offer more insight (or different insights) than the original proof;

---

[1]NB: if you did the practice exercises for chapter 1, you will know how to justify these conventions rig-
orously; and you will also, essentially, have performed the two 'trivial proofs by induction' mentioned in the
preceding proof.

- a new proof-technique might be useful in other ways: it might help you to prove other results (that you have not yet proved);
- sheer curiosity;

and more besides. All these reasons apply here.

Our first proof made use of truth tables. Our second proof will employ Lemma 2.2. The aim is to define an *algorithm* which will slowly replace subsentences with tautologically equivalent subsentences, pushing us in the direction of a DNF sentence, whilst preserving tautological equivalence, until eventually we end up with a tautologically equivalent sentence in DNF. (For a reminder about what an algorithm is, see §2.2.)

That is the big idea. But the details are conceptually trickier than those in the first proof. So we shall proceed slowly. I shall start with a relatively informal overview of the algorithm, which we shall call IntoDNF, and shall show how it applies to a particular input sentence. With the ideas fixed by the example, I shall offer a rigorous proof that the algorithm does what we require.

Our algorithm, IntoDNF, takes an input sentence, S, and then brings us closer and closer to DNF. The initial step of the algorithm, then, is obvious:

Step 0. Write down your initial sentence, S.

To make matters tractable, we shall work a particular application of the algorithm as we describe it. So, let us suppose that this is our input sentence:

$$((\neg A \to \neg(B \lor \neg C)) \to \neg(D \leftrightarrow \neg E))$$

The next step of the algorithm is designed to yield a tautologically equivalent sentence which satisfies (DNF1). Otherwise put, our aim is to eliminate all occurrences of conditionals and biconditionals from a sentence, whilst preserving tautological equivalence. So, here is what we should do:

Step 1a. Replace all subsentences of the form $(\mathcal{A} \to \mathcal{B})$ with $(\neg\mathcal{A} \lor \mathcal{B})$.
Step 1b. Replace all subsentences of the form $(\mathcal{A} \leftrightarrow \mathcal{B})$ with $((\mathcal{A} \land \mathcal{B}) \lor (\neg\mathcal{A} \land \neg\mathcal{B}))$.

Here is the result of running through Step 1 on our example. (The markers on the left-side indicate the Step we have just finished.)

$$0: \qquad ((\neg A \to \neg(B \lor \neg C)) \to \neg(D \leftrightarrow \neg E))$$
$$1a: \qquad (\neg(\neg\neg A \lor \neg(B \lor \neg C)) \lor \neg(D \leftrightarrow \neg E))$$
$$1b: \qquad (\neg(\neg\neg A \lor \neg(B \lor \neg C)) \lor \neg((D \land \neg E) \lor (\neg D \land \neg\neg E)))$$

We are a little closer to DNF, since this clearly satisfies (DNF1).

The next step of IntoDNF is designed to bring us to a tautologically equivalent sentence which satisfies (DNF2). Roughly, we need to 'push' the negations as deeply into the sentence as we possibly can. This is what the second step of IntoDNF does:

Step 2a: Replace all subsentences of the form $\neg\neg\mathcal{A}$ with $\mathcal{A}$.
Step 2b: Replace the first subsentence of the form $\neg(\mathcal{A} \land \mathcal{B})$ with $(\neg\mathcal{A} \lor \neg\mathcal{B})$.[2]

---

[2] Here 'first' means 'first, working from left to right'. In fact, inspecting the proof of Lemma 3.3, it is clear that the *order* in which we perform these manipulations does not matter at all; we just need to do them one at a time.

Step 2c: Replace the first subsentence of the form $\neg(\mathcal{A} \vee \mathcal{B})$ with $(\neg\mathcal{A} \wedge \neg\mathcal{B})$.
Step 2d: Repeat Steps 2a–2c, until no further change can occur.[3]

Applying Step 2 to to our example, we get the following sequence:

$$1: \quad (\neg(\neg\neg A \vee \neg(B \vee \neg C)) \vee \neg((D \wedge \neg E) \vee (\neg D \wedge \neg\neg E)))$$
$$2a: \quad (\neg(A \vee \neg(B \vee \neg C)) \vee \neg((D \wedge \neg E) \vee (\neg D \wedge E)))$$
$$2b: \quad (\neg(A \vee \neg(B \vee \neg C)) \vee \neg((D \wedge \neg E) \vee (\neg D \wedge E)))$$
$$2c: \quad ((\neg A \wedge \neg\neg(B \vee \neg C)) \vee \neg((D \wedge \neg E) \vee (\neg D \wedge E)))$$
$$2a: \quad ((\neg A \wedge (B \vee \neg C)) \vee \neg((D \wedge \neg E) \vee (\neg D \wedge E)))$$
$$2b: \quad ((\neg A \wedge (B \vee \neg C)) \vee \neg((D \wedge \neg E) \vee (\neg D \wedge E)))$$
$$2c: \quad ((\neg A \wedge (B \vee \neg C)) \vee (\neg(D \wedge \neg E) \wedge \neg(\neg D \wedge E)))$$
$$2a: \quad ((\neg A \wedge (B \vee \neg C)) \vee (\neg(D \wedge \neg E) \wedge \neg(\neg D \wedge E)))$$
$$2b: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \wedge \neg\neg E) \wedge \neg(\neg D \wedge E)))$$
$$2c: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \vee \neg\neg E) \wedge \neg(\neg D \wedge E)))$$
$$2a: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \vee E) \wedge \neg(\neg D \wedge E)))$$
$$2b: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \vee E) \wedge (\neg\neg D \vee \neg E)))$$
$$2c: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \vee E) \wedge (\neg\neg D \vee \neg E)))$$
$$2a: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \vee E) \wedge (D \vee \neg E)))$$

Note that the sentence does not necessarily change at every sub-step. But after the last written step, no further changes *can* occur; so we are done with Step 2. And we are even closer to DNF, since this satisfies (DNF2).

All that remains is to meet condition (DNF3), i.e. to ensure that no disjunction occurs within the scope of any conjunction. Roughly, we need to 'push' the conjunctions into the scope of disjunctions. And that is the idea behind the final step of IntoDNF:

Step 3a: Replace the first subsentence of the form $(\mathcal{A} \wedge (\mathcal{B} \vee \mathcal{C}))$ with $((\mathcal{A} \wedge \mathcal{B}) \vee (\mathcal{A} \wedge \mathcal{C}))$.
Step 3b: Replace the first subsentence of the form $((\mathcal{A} \vee \mathcal{B}) \wedge \mathcal{C})$ with $((\mathcal{A} \wedge \mathcal{C}) \vee (\mathcal{B} \wedge \mathcal{C}))$.
Step 3c: Repeat Steps 3a–3b, until no further changes can occur.

Applying Step 3 to our example, we get:

$$2: \quad ((\neg A \wedge (B \vee \neg C)) \vee ((\neg D \vee E) \wedge (D \vee \neg E)))$$
$$3a: \quad (((\neg A \wedge B) \vee (\neg A \wedge \neg C)) \vee ((\neg D \vee E) \wedge (D \vee \neg E)))$$
$$3b: \quad (((\neg A \wedge B) \vee (\neg A \wedge \neg C)) \vee ((\neg D \wedge (D \vee \neg E)) \vee (E \wedge (D \vee \neg E))))$$
$$3a: \quad (((\neg A \wedge B) \vee (\neg A \wedge \neg C)) \vee (((\neg D \wedge D) \vee (\neg D \wedge \neg E)) \vee (E \wedge (D \vee \neg E))))$$
$$3b: \quad (((\neg A \wedge B) \vee (\neg A \wedge \neg C)) \vee (((\neg D \wedge D) \vee (\neg D \wedge \neg E)) \vee (E \wedge (D \vee \neg E))))$$
$$3a: \quad (((\neg A \wedge B) \vee (\neg A \wedge \neg C)) \vee (((\neg D \wedge D) \vee (\neg D \wedge \neg E)) \vee ((E \wedge D) \vee (E \wedge \neg E))))$$

And this is in DNF, as may become clearer if we *temporarily* allow ourselves to remove brackets, following the notational conventions of forall*x*:Cambridge §10.3:

$$(\neg A \wedge B) \vee (\neg A \wedge \neg C) \vee (\neg D \wedge D) \vee (\neg D \wedge \neg E) \vee (E \wedge D) \vee (E \wedge \neg E)$$

---

[3]And we can check *that* no further change can occur, just by repeatedly running through Steps 2a–2c, until changes stop occurring.

So we are done! Finally, it is worth noting that this algorithm – although lengthy – was *much* more efficient at generating a DNF sentence than the truth-table method would have been. In fact, the truth table for our example would have 32 lines, 22 of which are true.

So far, I have defined the algorithm, IntoDNF, via Steps 0 through to 3c. I have also provided a worked-application of IntoDNF. But it remains to *prove* that IntoDNF really is fit for purpose.

First, we need to prove that any outputs from IntoDNF are tautologically equivalent to the inputs. This is actually quite simple:

Lemma 3.2. For any input sentence: any outputs of IntoDNF are tautologically equivalent to the input sentence.

Proof. At Step 0, we simply write our input. Each instance of Steps 1a–3b then involves replacing a subsentence with a tautologically equivalent subsentence. In particular, we make use of the following equivalences:

$$\begin{aligned}
\text{in Step 1a:} \qquad & (\mathcal{A} \rightarrow \mathcal{B}) \vDash\!\!\dashv (\neg\mathcal{A} \vee \mathcal{B}) \\
\text{in Step 1b:} \qquad & (\mathcal{A} \leftrightarrow \mathcal{B}) \vDash\!\!\dashv ((\mathcal{A} \wedge \mathcal{B}) \vee (\neg\mathcal{A} \wedge \neg\mathcal{B})) \\
\text{in Step 2a:} \qquad & \neg\neg\mathcal{A} \vDash\!\!\dashv \mathcal{A} \\
\text{in Step 2b:} \qquad & \neg(\mathcal{A} \wedge \mathcal{B}) \vDash\!\!\dashv (\neg\mathcal{A} \vee \neg\mathcal{B}) \\
\text{in Step 2c:} \qquad & \neg(\mathcal{A} \vee \mathcal{B}) \vDash\!\!\dashv (\neg\mathcal{A} \wedge \neg\mathcal{B}) \\
\text{in Step 3a:} \qquad & (\mathcal{A} \wedge (\mathcal{B} \vee \mathcal{C})) \vDash\!\!\dashv ((\mathcal{A} \wedge \mathcal{B}) \vee (\mathcal{A} \wedge \mathcal{C})) \\
\text{in Step 3b:} \qquad & ((\mathcal{A} \vee \mathcal{B}) \wedge \mathcal{C}) \vDash\!\!\dashv ((\mathcal{A} \wedge \mathcal{C}) \vee (\mathcal{B} \wedge \mathcal{C}))
\end{aligned}$$

Now, by Lemma 2.2, if we replace some subsentence of a sentence with a tautologically equivalent subsentences, we obtain a sentence that is tautologically equivalent to what we started with. So any instance of any Step of IntoDNF preserves tautological equivalence. So, provided the algorithm yields an output within finitely many steps, that output is indeed (by a trivial induction) tautologically equivalent to the input. ∎

The next task is to show that IntoDNF *always* yields an output. We need to ensure, for example, that the algorithm never gets trapped on an infinite loop. At the same time as we prove this, we shall also show that the output is in DNF, as we would hope.

Unfortunately, the proof is surprisingly difficult. So I shall start with an intuitive overview of how the proof. We will aim, intuitively, to show each of the following:

- *The output of Step 1 is a sentence satisfying* (DNF1).
  This will be fairly trivial.
- *The output of Step 2 is a sentence also satisfying* (DNF2).
  Intuitively, this is because the negation-signs get pushed progressively deeper into the formula, until they precede only atomic sentences.
- *The output of Step 3 is a sentence also satisfying* (DNF3).
  Intuitively, this is because the conjunction-signs get pushed progressively deeper into the formula, until they have no disjunctions in their scope.

That is the main idea. As a student, your aim should be to grasp this main idea fully. Morally speaking, it just *must* be right. And, for just this reason, many books at this level leave matters here, and claim that they have proved the DNF Theorem. But, *just because we can*,[4] here is a full proof, which brings these intuitive points into sharp relief:

Lemma 3.3. On any input, IntoDNF is guaranteed to terminate, yielding an output in DNF.

Proof. Concerning Step 0: trivially, we will complete Step 0 (eventually), since this Step just involves writing down a single (finitely long) input sentence.

Concerning Step 1: The input sentence only contained a finite number of instances of '→' and '↔'. These will all be eliminated (eventually) during Step 1. As such, the output at the end of Step 1 satisfies (DNF1).

Concerning Step 2: This is a bit harder. For each sentence $\mathcal{A}$, we define its not-height, $nh(\mathcal{A})$, recursively:[5]

$$nh(\mathcal{A}) = 1, \text{if } \mathcal{A} \text{ is atomic}$$
$$nh(\neg\mathcal{A}) = 3 \times nh(\mathcal{A}) - 1$$
$$nh(\mathcal{A} \wedge \mathcal{B}) = nh(\mathcal{A} \vee \mathcal{B}) = nh(\mathcal{A}) + nh(\mathcal{B})$$

I claim that, *if an instance of Step 2a, 2b or 2c has any effect on the sentence, then it reduces the sentence's not-height.* In the case of Step 2a, observe that:

$$\begin{aligned} nh(\neg\neg\mathcal{A}) &= 3 \times nh(\neg\mathcal{A}) - 1 \\ &= 3 \times (3 \times nh(\mathcal{A}) - 1) - 1 \\ &= 9 \times nh(\mathcal{A}) - 4 \end{aligned}$$

and this is greater than $nh(\mathcal{A})$ for any $\mathcal{A}$, since a sentence's not-height must always be at least 1. In the case of Step 2b, observe that:

$$\begin{aligned} nh(\neg(\mathcal{A} \wedge \mathcal{B})) &= 3 \times (nh(\mathcal{A} \wedge \mathcal{B})) - 1 \\ &= 3 \times (nh(\mathcal{A}) + nh(\mathcal{B})) - 1 \end{aligned}$$

and this is always greater than:

$$\begin{aligned} nh(\neg\mathcal{A} \vee \neg\mathcal{B}) &= nh(\neg\mathcal{A}) + nh(\neg\mathcal{B}) \\ &= 3 \times nh(\mathcal{A}) - 1 + 3 \times nh(\mathcal{B}) - 1 \\ &= 3 \times (nh(\mathcal{A}) + nh(\mathcal{B})) - 2 \end{aligned}$$

The case of Step 2c is exactly similar, and this establishes the claim.

---

[4]Can, and *therefore should*? There is a deep philosophical-cum-technical question: how much rigour does a proof require? But this is a book in formal logic, not philosophical logic; so I leave this as a cliffhanger.

[5]Why do we define not-height this way? The glib answer is: because *it works* in the proof. In fact, this glib answer really says it all. I have deliberately manufactured an arithmetical function, *nh*, with no other end in mind, than that successive applications of Steps 2a–2c of the algorithm are guaranteed to yield a decreasing sequence of *nh* values.

Consequently, repeatedly running through Steps 2a–2c yields a sequence of sentences with decreasing not-heights. Since the input sentence to Step 2 (i.e. the output from Step 1) has some finite not-height, after a finite number of applications of Steps 2a–2c, it will be impossible to reduce the not-height any further. This brings Step 2 to an end. Trivially, the resulting sentence still satisfies (DNF1). And it also satisfies (DNF2), since the not-height of any sentence satisfying (DNF1) but not satisfying (DNF2) clearly *can* be decreased by some application of Step 2a, 2b or 2c.

Concerning Step 3: This is like Step 2. For each sentence $\mathcal{A}$, we define its conjunction-height, $ch(\mathcal{A})$, recursively:

$$ch(\mathcal{A}) = 2, \text{ if } \mathcal{A} \text{ is atomic}$$
$$ch(\neg\mathcal{A}) = ch(\mathcal{A})$$
$$ch(\mathcal{A} \wedge \mathcal{B}) = ch(\mathcal{A}) \times ch(\mathcal{B})$$
$$ch(\mathcal{A} \vee \mathcal{B}) = ch(\mathcal{A}) + ch(\mathcal{B}) + 1$$

I claim that, *if an instance of Step 3a or 3b has any effect on the sentence, then it reduces the sentence's conjunction-height.* In the case of Step 3a, observe that

$$ch(\mathcal{A} \wedge (\mathcal{B} \vee \mathcal{C})) = ch(\mathcal{A}) \times (ch(\mathcal{B}) + ch(\mathcal{C}) + 1)$$
$$= ch(\mathcal{A}) \times ch(\mathcal{B}) + ch(\mathcal{A}) \times ch(\mathcal{C}) + ch(\mathcal{A})$$

and, since $ch(\mathcal{A}) > 1$ for any $\mathcal{A}$, this must be larger than:

$$ch((\mathcal{A} \wedge \mathcal{B}) \vee (\mathcal{A} \wedge \mathcal{C})) = ch(\mathcal{A}) \times ch(\mathcal{B}) + ch(\mathcal{A}) \times ch(\mathcal{C}) + 1$$

The case of Step 3b is exactly similar. This establishes the claim.

Now, as above: running through Steps 3a–3b repeatedly yields a sequence of sentences with decreasing conjunction-heights. So, after a finite number of applications of Steps 3a–3b, it will be impossible to reduce the conjunction-height any further. Trivially, the resulting sentence still satisfies (DNF1) and (DNF2). And it also satisfies (DNF3), since the conjunction-height of any sentence satisfying both (DNF1) and (DNF2) but not (DNF3) *can* be decreased by some application of Step 3a or 3b.   ∎

Now if we just combine combine Lemma 3.2 with Lemma 3.3, we have a new proof of the DNF Theorem, via substitution.

3.4   Cleaning up DNF sentences

I have offered two different proofs of the DNF Theorem. They also described two different algorithms for putting sentences into DNF. But these algorithms do not always leave you with the most elegant output. Indeed, just as we can 'clean up' the output of the algorithm for calculating interpolants (see §2.2), we can 'clean up' the output of our algorithms for calculating DNF-equivalents. Here are the clean-up rules; I leave it as an exercise to *justify* them. First, we rewrite our DNF sentence using the relaxed notational conventions of forall*x*:Cambridge §10.3. Now we apply the following:

Rule 1: remove repetitious conjuncts. If any disjunct contains the same conjunct more than

once, remove all but one instance of that conjunct (respecting bracketing conventions as appropriate).

Example: if the output is:

$$(A \wedge B) \vee (A \wedge \neg C) \vee D \vee (A \wedge B)$$

then applying Rule 1 we get:

$$(A \wedge B) \vee (A \wedge \neg C) \vee D$$

Rule 2: remove overly-specific disjuncts. Delete any disjunct which entails any other disjunct. (If two or more disjuncts entail each other, delete all but the first of those disjuncts.)

Example: the output of our example in §3.3 was:

$$(A \wedge B) \vee (A \wedge \neg C) \vee (\neg D \wedge D) \vee (\neg E \wedge D) \vee (\neg D \wedge E) \vee (\neg E \wedge E)$$

Noting that a contradiction entails *everything*, we should remove both contradictory disjuncts, obtaining (still with relaxed conventions):

$$(A \wedge B) \vee (A \wedge \neg C) \vee (\neg E \wedge D) \vee (\neg D \wedge E)$$

Rule 3: invoke excluded middle. Suppose there is a disjunct whose conjuncts are just $\mathcal{A}_1, \ldots, \mathcal{A}_m$ and $\mathcal{B}$ (in any order), and another disjunct whose conjuncts are just $\mathcal{A}_1, \ldots, \mathcal{A}_m$ and $\neg\mathcal{B}$ (in any order), so that the two disjuncts disagree only on whether whether $\mathcal{B}$ should be prefixed with a negation. In that case, delete both disjuncts and replace them with the simpler disjunct $(\mathcal{A}_1 \wedge \ldots \wedge \mathcal{A}_m)$. (However, if we have *just* $\mathcal{B}$ and $\neg\mathcal{B}$ as disjuncts, then delete the entire sentence and leave the tautology $(\mathcal{B} \vee \neg\mathcal{B})$.)

Example: the output of the example in §3.2 was:

$$(A \wedge B \wedge C) \vee (A \wedge \neg B \wedge C) \vee (\neg A \wedge \neg B \wedge C) \vee (\neg A \wedge \neg B \wedge \neg C)$$

which can be simplified to:

$$(A \wedge C) \vee (\neg A \wedge \neg B)$$

3.5   Conjunctive Normal Form

So far in this chapter, I have discussed *disjunctive* normal form. Given the duality of disjunction and conjunction (see §2.3), it may not come as a surprise to hear that there is also such a thing as *conjunctive normal form* (CNF).

The definition of CNF is exactly analogous to the definition of DNF. So, a sentence is in CNF *iff* it meets all of the following conditions:

(CNF1) No connectives occur in the sentence other than negations, conjunctions and disjunctions;
(CNF2) Every occurrence of negation has minimal scope;
(CNF3) No conjunction occurs within the scope of any disjunction.

Generally, then, a sentence in CNF looks like this

$$\left(\pm\mathcal{A}_1 \vee \ldots \vee \pm\mathcal{A}_i\right) \wedge \left(\pm\mathcal{A}_{i+1} \vee \ldots \vee \pm\mathcal{A}_j\right) \wedge \ldots \wedge \left(\pm\mathcal{A}_{m+1} \vee \ldots \vee \pm\mathcal{A}_n\right)$$

where each $\mathcal{A}_k$ is an atomic sentence.

Since '¬' is its own dual, and '∨' and '∧' are the duals of each other, it is immediate clear that if a sentence is in DNF, then its dual is in CNF; and *vice versa*. Armed with this insight, we can immediately prove another normal form theorem:

Theorem 3.4. Conjunctive Normal Form Theorem. For any sentence, there is a tautologically equivalent sentence in conjunctive normal form.

Proof. Let S be any sentence. Applying the DNF Theorem to $\underline{S}$, there is a DNF sentence, $\mathcal{D}$, such that $\underline{S} \vDash\dashv \mathcal{D}$. So, by Theorem 2.7, $\underline{\underline{S}} \vDash\dashv \underline{\mathcal{D}}$. Whence, by Lemma 2.5, $S \vDash\dashv \underline{\mathcal{D}}$. But, since $\mathcal{D}$ is in DNF, $\underline{\mathcal{D}}$ is in CNF. ∎

This slick proof is a further illustration of the power of duality. However, it might suggest that the DNF Theorem enjoys some kind of 'precedence' over the CNF Theorem. That would be misleading. We can easily prove the CNF Theorem directly, using either of the two proof techniques that we used to prove the DNF Theorem (whereupon the DNF Theorem could be proved as a consequence of the CNF Theorem and duality). I shall sketch the main ideas of the direct proofs of the CNF Theorem; I leave it as an exercise to make these ideas more precise.

Proof sketch of CNF Theorem via truth tables. Given a TFL sentence, S, we begin by writing down the complete truth table for S.

If S is *true* on every line of the truth table, then $S \vDash\dashv \left(\mathcal{A}_1 \vee \neg\mathcal{A}_1\right)$.

If S is *false* on at least one line of the truth table then, for every line on the truth table where S is false, write down a disjunction $\left(\pm\mathcal{A}_1 \vee \ldots \vee \pm\mathcal{A}_n\right)$ which is *false* on (and only on) that line. Let C be the conjunction of all of these disjuncts; by construction, C is in CNF and $S \vDash\dashv C$. ∎

Proof sketch of CNF Theorem via substitution. Using Steps 0–2 of IntoDNF, we obtain a sentence satisfying (CNF1) and (CNF2). To turn this into a sentence in CNF, we simply modify Step 3, using substitutions based on the duals (what else!) of the Distribution Laws used in Step 3 of IntoDNF; i.e.:

$$\left((\mathcal{A} \wedge \mathcal{B}) \vee \mathcal{C}\right) \vDash\dashv \left((\mathcal{A} \vee \mathcal{C}) \wedge (\mathcal{B} \vee \mathcal{C})\right)$$
$$\left(\mathcal{A} \vee (\mathcal{B} \wedge \mathcal{C})\right) \vDash\dashv \left((\mathcal{A} \vee \mathcal{B}) \wedge (\mathcal{A} \vee \mathcal{C})\right)$$

This allows us to pull conjunctions outside the scope of any disjunction. ∎

Of course, after obtaining a sentence in CNF – however we chose to do it – we can clean-up the resulting sentence using rules similar to those of §3.4. I leave it as an exercise, to determine what the appropriate rules should be.

Practice exercises

A. Consider the following sentences:

1. $(A \rightarrow \neg B)$
2. $\neg(A \leftrightarrow B)$
3. $(\neg A \vee \neg(A \wedge B))$
4. $(\neg(A \rightarrow B) \wedge (A \rightarrow C))$
5. $(\neg(A \vee B) \leftrightarrow ((\neg C \wedge \neg A) \rightarrow \neg B))$
6. $((\neg(A \wedge \neg B) \rightarrow C) \wedge \neg(A \wedge D))$

For each sentence:

- use both algorithms (by truth table, and by substitution) to write down sentences in DNF that are tautologically equivalent to these sentences.
- use both algorithms to write down sentences in CNF that are tautologically equivalent to these sentences.

B. Offer proofs of the two 'trivial inductions' mentioned in the proof of Theorem 3.1 via truth tables.

C. Explain why the proof of Theorem 3.2 required 'a trivial induction' (in the last sentence). Prove that induction.

D. Justify each of the rules mentioned in §3.4, i.e. prove that applying each of them preserves tautological equivalence.

E. Fill out the details of the two proof sketches of the CNF Theorem. Determine the appropriate 'clean up' rules for CNF sentences.

Expressive adequacy   4

In chapter 3, we established two normal form theorems. In this chapter, we shall use them to demonstrate the expressive power of TFL.

4.1   The expressive adequacy of TFL

In discussing duality (§2.3), I introduced the general idea of an *n*-place connective. We might, for example, define a three-place connective, '♡', into existence, by stipulating that it is to have the following characteristic truth table:

$\mathcal{A}$	$\mathcal{B}$	$\mathcal{C}$	$♡(\mathcal{A}, \mathcal{B}, \mathcal{C})$
T	T	T	F
T	T	F	T
T	F	T	T
T	F	F	F
F	T	T	F
F	T	F	T
F	F	T	F
F	F	F	F

Probably this new connective would not correspond with any natural English expression (in the way that '∧' corresponds with 'and'). But a question arises: if we wanted to employ a connective with this characteristic truth table, must we add a *new* connective to TFL? Or can we get by with the connectives we *already have*?

Let us make this question more precise. Say that some connectives are JOINTLY EXPRESSIVELY ADEQUATE *iff*, for any possible truth function, there is a scheme containing only those connectives which expresses that truth function. Since we can represent truth functions using characteristic truth tables, we could equivalently say the following: some connectives are jointly expressively adequate *iff*, for any possible truth table, there is a scheme containing only those connectives with that truth table.

I say 'scheme' rather than 'sentence', because we are not concerned with something as specific as a sentence (it might help to reread *forallx:Cambridge* §7, and this book's §2.1). To see why, consider the characteristic truth table for conjunction; this schematically encodes the information that a conjunction $(\mathcal{A} \wedge \mathcal{B})$ is true iff both $\mathcal{A}$ and $\mathcal{B}$ are true (whatever $\mathcal{A}$ and $\mathcal{B}$ might be). When we discuss expressive adequacy, we are considering something at the same level of generality.

The general point is, when we are armed with some jointly expressively adequate connectives, no truth function lies beyond our grasp. And in fact, we are in luck.[1]

Theorem 4.1. Expressive Adequacy Theorem. The connectives of TFL are jointly expressively adequate. Indeed, the following pairs of connectives are jointly expressively adequate:

1. '¬' and '∨'
2. '¬' and '∧'
3. '¬' and '→'

Proof. Given any truth table, we can use the method of proving the DNF Theorem (or the CNF Theorem) via truth tables, to write down a scheme which has the same truth table. For example, employing the truth table method for proving the DNF Theorem, I can tell you that the following scheme has the same characteristic truth table as $\heartsuit(\mathcal{A}, \mathcal{B}, \mathcal{C})$, above:

$$(\mathcal{A} \wedge \mathcal{B} \wedge \neg\mathcal{C}) \vee (\mathcal{A} \wedge \neg\mathcal{B} \wedge \mathcal{C}) \vee (\neg\mathcal{A} \wedge \mathcal{B} \wedge \neg\mathcal{C})$$

It follows that the connectives of TFL are jointly expressively adequate. I now prove each of the subsidiary results.

Subsidiary Result 1: expressive adequacy of '¬' and '∨'. Observe that the scheme that we generate, using the truth table method of proving the DNF Theorem, will only contain the connectives '¬', '∧' and '∨'. So it suffices to show that there is an equivalent scheme which contains only '¬' and '∨'. To show do this, we simply consider the following equivalence:

$$(\mathcal{A} \wedge \mathcal{B}) \dashv\vDash \neg(\neg\mathcal{A} \vee \neg\mathcal{B})$$

and reason as in §3.2 (I leave the details as an exercise).

Subsidiary Result 2: expressive adequacy of '¬' and '∧'. Exactly as in Subsidiary Result 1, making use of this equivalence instead:

$$(\mathcal{A} \vee \mathcal{B}) \dashv\vDash \neg(\neg\mathcal{A} \wedge \neg\mathcal{B})$$

Subsidiary Result 3: expressive adequacy of '¬' and '→'. Exactly as in Subsidiary Result 1, making use of these equivalences instead:

$$(\mathcal{A} \vee \mathcal{B}) \dashv\vDash (\neg\mathcal{A} \rightarrow \mathcal{B})$$
$$(\mathcal{A} \wedge \mathcal{B}) \dashv\vDash \neg(\mathcal{A} \rightarrow \neg\mathcal{B})$$

Alternatively, we could simply rely upon one of the other two subsidiary results, and (repeatedly) invoke only one of these two equivalences. ∎

In short, there is never any *need* to add new connectives to TFL. Indeed, there is already some redundancy among the connectives we have: we could have made do with just two connectives, if we had been feeling really austere.

---

[1]This, and Theorem 4.5, follow from a more general result, due to Emil Post in 1941. For a modern presentation of Post's result, see Pelletier and Martin 'Post's Functional Completeness Theorem' (1990, *Notre Dame Journal of Formal Logic* 31.2).

4.2 Individually expressively adequate connectives

In fact, some two-place connectives are *individually* expressively adequate. These connectives are not standardly included in TFL, since they are rather cumbersome to use. But their existence shows that, if we had wanted to, we could have defined a truth-functional language that was expressively adequate, which contained only a single primitive connective.

The first such connective we shall consider is '↑', which has the following characteristic truth table.

$\mathcal{A}$	$\mathcal{B}$	$\mathcal{A} \uparrow \mathcal{B}$
T	T	F
T	F	T
F	T	T
F	F	T

This is often called 'the Sheffer stroke', after Harry Sheffer, who used it to show how to reduce the number of logical connectives in Russell and Whitehead's *Principia Mathematica*.[2] (In fact, Charles Sanders Peirce had anticipated Sheffer by about 30 years, but never published his results.)[3] It is quite common, as well, to call it 'nand', since its characteristic truth table is the negation of the truth table for '∧'.

Proposition 4.2. '↑' is expressively adequate all by itself.

Proof. Theorem 4.1 tells us that '¬' and '∨' are jointly expressively adequate. So it suffices to show that, given any scheme which contains only those two connectives, we can rewrite it as a tautologically equivalent scheme which contains only '↑'. As in the proof of the subsidiary cases of Theorem 4.1, then, we simply apply the following equivalences:

$$\neg \mathcal{A} \vDash\dashv (\mathcal{A} \uparrow \mathcal{A})$$
$$(\mathcal{A} \vee \mathcal{B}) \vDash\dashv ((\mathcal{A} \uparrow \mathcal{A}) \uparrow (\mathcal{B} \uparrow \mathcal{B}))$$

to Subsidiary Result 1 of Theorem 4.1. ■

Similarly, we can consider the connective '↓':

$\mathcal{A}$	$\mathcal{B}$	$\mathcal{A} \downarrow \mathcal{B}$
T	T	F
T	F	F
F	T	F
F	F	T

This is sometimes called the 'Peirce arrow' (Peirce himself called it 'ampheck'). More often, though, it is called 'nor', since its characteristic truth table is the negation of '∨'.

---

[2] Sheffer, 'A Set of Five Independent Postulates for Boolean Algebras, with application to logical constants,' (1913, *Transactions of the American Mathematical Society* 14.4)

[3] See Peirce, 'A Boolian Algebra with One Constant', which dates to c.1880; and Peirce's *Collected Papers*, 4.264–5.

Proposition 4.3. '↓' is expressively adequate all by itself.

Proof. As in Proposition 4.2, although invoking the dual equivalences:

$$\neg \mathcal{A} \Eqvv (\mathcal{A} \downarrow \mathcal{A})$$
$$(\mathcal{A} \wedge \mathcal{B}) \Eqvv ((\mathcal{A} \downarrow \mathcal{A}) \downarrow (\mathcal{B} \downarrow \mathcal{B}))$$

and Subsidiary Result 2 of Theorem 4.1. ∎

4.3 Failures of expressive adequacy

In fact, the *only* two-place connectives which are individually expressively adequate are '↑' and '↓'. But how would we show this? More generally, how can we show that some connectives are *not* jointly expressively adequate?

The obvious thing to do is to try to find some truth table which we *cannot* express, using just the given connectives. But there is a bit of an art to this. Moreover, in the end, we shall have to rely upon induction; for we shall need to show that *no* scheme – no matter how *long* – is capable of expressing the target truth table.

To make this concrete, let's consider the question of whether '∨' is expressively adequate all by itself. After a little reflection, it should be clear that it is not. In particular, it should be clear that any scheme which only contains disjunctions cannot have the same truth table as negation, i.e.:

$\mathcal{A}$	$\neg \mathcal{A}$
T	F
F	T

The intuitive reason, why this should be so, is simple: the top line of the desired truth table needs to have the value False; but the top line of any truth table for a scheme which *only* contains disjunctions will always be True. But so far, this is just hand-waving. To make it rigorous, we need to reach for induction. Here, then, is our rigorous proof.

Proposition 4.4. '∨' is not expressively adequate by itself.

Proof. Let $\mathcal{A}$ by any scheme containing no connective other than disjunctions. Suppose, for induction on length, that every shorter scheme containing only disjunctions is true whenever all its atomic constituents are true. There are two cases to consider:

> *Case 1: $\mathcal{A}$ is atomic.* Then there is nothing to prove.
> *Case 2: $\mathcal{A}$ is $(\mathcal{B} \vee \mathcal{C})$, for some schemes $\mathcal{B}$ and $\mathcal{C}$ containing only disjunctions.* Then, since $\mathcal{B}$ and $\mathcal{C}$ are both shorter than $\mathcal{A}$, by the induction hypothesis they are both true when all their atomic constituents are true. Now the atomic constituents of $\mathcal{A}$ are just the constituents of both $\mathcal{B}$ and $\mathcal{C}$, and $\mathcal{A}$ is true whenever $\mathcal{B}$ and $\mathcal{C}$. So $\mathcal{A}$ is true when all of its atomic constituents are true.

It now follows, by induction on length, that any scheme containing no connective other than disjunctions is true whenever all of its atomic constituents are true. Consequently, no scheme containing only disjunctions has the same truth table as that of negation. Hence '∨' is not expressively adequate by itself. ∎

In fact, we can generalise Proposition 4.4:

Theorem 4.5. The *only* two-place connectives that are expressively adequate by themselves are '↑' and '↓'.

Proof. There are sixteen distinct two-place connectives. We shall run through them all, considering whether or not they are individually expressively adequate, in four groups.

Group 1: the top line of the truth table is True. Consider those connectives where the top line of the truth table is True. There are eight of these, including '∧', '∨', '→' and '↔', but also the following:

$\mathcal{A}$	$\mathcal{B}$	$\mathcal{A} \circ_1 \mathcal{B}$	$\mathcal{A} \circ_2 \mathcal{B}$	$\mathcal{A} \circ_3 \mathcal{B}$	$\mathcal{A} \circ_4 \mathcal{B}$
T	T	T	T	T	T
T	F	T	T	T	F
F	T	T	F	F	T
F	F	T	T	F	F

(obviously the names for these connectives were chosen arbitrarily). But, exactly as in Proposition 4.4, none of these connectives can express the truth table for negation. So there is a connective whose truth table they cannot express. So none of them is individually expressively adequate.

Group 2: the bottom line of the truth table is False. Having eliminated eight connectives, eight remain. Of these, four are false on the bottom line of their truth table, namely:

$\mathcal{A}$	$\mathcal{B}$	$\mathcal{A} \circ_5 \mathcal{B}$	$\mathcal{A} \circ_6 \mathcal{B}$	$\mathcal{A} \circ_7 \mathcal{B}$	$\mathcal{A} \circ_8 \mathcal{B}$
T	T	F	F	F	F
T	F	T	T	F	F
F	T	T	F	T	F
F	F	F	F	F	F

As above, though, none of these connectives can express the truth table for negation. To show this we prove that any scheme whose only connective is one of these (perhaps several times) is false whenever all of its atomic constituents are false. We can show this by induction, exactly as in Proposition 4.4 (I leave the details as an exercise).

Group 3: connectives with redundant positions. Consider two of the remaining four connectives:

$\mathcal{A}$	$\mathcal{B}$	$\mathcal{A} \circ_9 \mathcal{B}$	$\mathcal{A} \circ_{10} \mathcal{B}$
T	T	F	F
T	F	F	T
F	T	T	F
F	F	T	T

These connectives have redundant positions, in the sense that the truth value of the overarching scheme only depends upon the truth value of one of the atomic constituents. More precisely:

$$\mathcal{A} \circ_9 \mathcal{B} \dashv\vDash \neg\mathcal{A}$$
$$\mathcal{A} \circ_{10} \mathcal{B} \dashv\vDash \neg\mathcal{B}$$

Consequently, there are many truth functions that they cannot express. In particular, they cannot express either the tautologous truth function (given by '$\circ_1$'), or the contradictory truth function (given by '$\circ_8$'). To show this, it suffices to prove that any scheme whose only connective is either '$\circ_9$' or '$\circ_{10}$' (perhaps several times) is contingent, i.e. it is true on at least one line and false on at least one other line. I leave the details of this proof as an exercise.

Group 4. Only two connectives now remain, namely '↑' and '↓', and Propositions 4.2 and 4.3 show that both are individually expressively adequate.   ■

The above inductions on complexity were fairly routine. To close this chapter, I shall offer some similar results about expressive limitation that are only slightly more involved.

Proposition 4.6. Exactly four 2-place truth functions can be expressed using schemes whose only connectives are biconditionals.

Proof. I claim that the four 2-place truth functions that can be so expressed are expressed by:

$$\mathcal{A} \leftrightarrow \mathcal{A}$$
$$\mathcal{A}$$
$$\mathcal{B}$$
$$\mathcal{A} \leftrightarrow \mathcal{B}$$

It is clear that we can express all four, and that they are distinct.

To see that these are the *only* functions which can be expressed, I shall perform an induction on the length of our scheme. Let $\mathcal{S}$ be any scheme containing only (arbitrarily many) biconditionals and the atomic constituents $\mathcal{A}$ and $\mathcal{B}$ (arbitrarily many times). Suppose that any such shorter scheme is equivalent to one of the four schemes listed above. There are two cases to consider:

> *Case 1: $\mathcal{S}$ is atomic.* Then certainly it is equivalent to either $\mathcal{A}$ or to $\mathcal{B}$.
> *Case 2: $\mathcal{S}$ is $(\mathcal{C} \leftrightarrow \mathcal{D})$, for some schemes $\mathcal{C}$ and $\mathcal{D}$.* Since $\mathcal{C}$ and $\mathcal{D}$ are both shorter than $\mathcal{S}$, by supposition they are equivalent to one of the four mentioned schemes. And it is now easy to check, case-by-case, that, whichever of the four schemes each is equivalent to, $(\mathcal{C} \leftrightarrow \mathcal{D})$ is also equivalent to one of those four schemes. (I leave the details as an exercise.)

This completes the proof of my claim, by induction on length of scheme.   ■

Proposition 4.7. Exactly eight 2-place truth functions can be expressed using schemes whose only connectives are negations and biconditionals.

Proof. Observe the following equivalences:

$$\neg\neg\mathcal{A} \dashv\vDash \mathcal{A}$$
$$(\neg\mathcal{A} \leftrightarrow \mathcal{B}) \dashv\vDash \neg(\mathcal{A} \leftrightarrow \mathcal{B})$$
$$(\mathcal{A} \leftrightarrow \neg\mathcal{B}) \dashv\vDash \neg(\mathcal{A} \leftrightarrow \mathcal{B})$$

We can use these tautological equivalences along the lines of Lemma 3.2 to show the following: for any scheme containing only biconditionals and negations, there is a tautologically equivalent scheme where no negation is inside the scope of any biconditional. (The details of this are left as an exercise.)

So, given any scheme containing only the atomic constituents $\mathcal{A}$ and $\mathcal{B}$ (arbitrarily many times), and whose only connectives are biconditionals and negations, we can produce a tautologically equivalent scheme which contains at most *one* negation, and that at the start of the scheme. By Proposition 4.6, the subscheme following the negation (if there is a negation) must be equivalent to one of only four schemes. Hence the original scheme must be equivalent to one of only eight schemes.   ∎

Practice exercises

A. Where '$\circ_7$' has the characteristic truth table defined in the proof of Theorem 4.5, show that the following are jointly expressively adequate:

1. '$\circ_7$' and '$\neg$'.
2. '$\circ_7$' and '$\rightarrow$'.
3. '$\circ_7$' and '$\leftrightarrow$'.

B. Show that the connectives '$\circ_7$', '$\wedge$' and '$\vee$' are not jointly expressively adequate.

C. Complete the exercises left in each of the proofs of this chapter.

Soundness   5

So far, I have discussed only the syntax and semantics of TFL. In this chapter, and the next, I shall relate TFL's semantics to its *proof system* (as defined in forall$x$:Cambridge). I shall prove that the semantics and the formal proof system mirror are perfectly suited to one another, in a sense to be made more precise.

5.1   Soundness defined, and setting up the proof

Intuitively, a formal proof system is sound iff it does not allow you to prove any invalid arguments. This is obviously a highly desirable property. It tells us that our proof system will never lead us astray. Indeed, if our proof system were not sound, then we would not be able to trust our proofs. The aim of this chapter is to prove that our proof system is sound.

Let me make the idea more precise. A formal proof system is SOUND (relative to a given semantics) *iff*, if there is a formal proof of C from assumptions among Γ, then Γ genuinely entails C (given that semantics). Otherwise put, to prove that TFL's proof system is sound, I need to prove the following

Theorem 5.1. Soundness. For any sentences Γ and C: if $\Gamma \vdash C$, then $\Gamma \vDash C$

To prove this, we shall check each of the rules of TFL's proof system individually. We want to show that no application of those rules ever leads us astray. Since a proof just involves repeated application of those rules, this will show that no proof ever leads us astray. Or at least, that is the general idea.

To begin with, we must make the idea of 'leading us astray' more precise. Say that a line of a proof is SHINY iff the assumptions on which that line depends tautologically entail the sentence on that line.[1] To illustrate the idea, consider the following TFL-proof:

1	$F \rightarrow (G \wedge H)$	
2	F	
3	$G \wedge H$	$\rightarrow$E 1, 2
4	G	$\wedge$E 3
5	$F \rightarrow G$	$\rightarrow$I 2–4

---

[1] The word 'shiny' is not standard among logicians.

38

Line 1 is shiny iff $F \rightarrow (G \wedge H) \vDash F \rightarrow (G \wedge H)$. You should be easily convinced that line 1 is, indeed, shiny! Similarly, line 4 is shiny iff $F \rightarrow (G \wedge H), F \vDash G$. Again, it is easy to check that line 4 is shiny. As is every line in this TFL-proof. I want to show that this is no coincidence. That is, I want to prove:

Lemma 5.2. Every line of every TFL-proof is shiny.

Then we will know that we have never gone astray, on any line of a proof. Indeed, given Lemma 5.2, it will be easy to prove the Soundness Theorem:

Proof of Soundness Theorem, given Lemma 5.2. Suppose $\Gamma \vdash C$. Then there is a TFL-proof, with C appearing on its last line, whose only undischarged assumptions are among Γ. Lemma 5.2 tells us that every line on every TFL-proof is shiny. So this last line is shiny, i.e. $\Gamma \vDash C$. ∎

It remains to prove Lemma 5.2.

To do this, we observe that every line of any TFL-proof is obtained by applying some rule. So what I want to show is that no application of a rule of TFL's proof system will lead us astray. More precisely, say that a rule of inference is RULE-SOUND *iff* for all TFL-proofs, if we obtain a line on a TFL-proof by applying that rule, and every earlier line in the TFL-proof is shiny, then our new line is also shiny. What I need to show is that *every* rule in TFL's proof system is rule-sound.

I shall do this in the next section. But having demonstrated the rule-soundness of every rule, Lemma 5.2 will follow immediately:

Proof of Lemma 5.2, given that every rule is rule-sound. Fix any line, line n, on any TFL-proof. The sentence written on line n must be obtained using a formal inference rule which is rule-sound. This is to say that, if every earlier line is shiny, then line n itself is shiny. Hence, by strong induction on the length of TFL-proofs, every line of every TFL-proof is shiny. ∎

Note that this proof appeals to a principle of strong induction on the length of TFL-proofs. This is the first time we have seen that principle, and you should pause to confirm that it is, indeed, justified. (Hint: compare the justification for the Principle of Strong Induction on Length *of sentences* from §1.4.)

5.2   Checking each rule

It remains to show that every rule is rule-sound. This is actually much easier than many of the proofs in earlier chapters. But it is also time-consuming, since we need to check each rule individually, and TFL's proof system has plenty of rules! To speed up the process marginally, I shall introduce a convenient abbreviation: 'Δ_i' will abbreviate the assumptions (if any) on which line i depends in our TFL-proof (context will indicate which TFL-proof I have in mind).

Lemma 5.3. Introducing an assumption is rule-sound.

Proof. If $\mathcal{A}$ is introduced as an assumption on line n, then $\mathcal{A}$ is among Δ_n, and so $\Delta_n \vDash$ $\mathcal{A}$.   ∎

Lemma 5.4. ∧I is rule-sound.

Proof. Consider any application of ∧I in any TFL-proof, i.e. something like:

$$
\begin{array}{c|l}
i & \mathcal{A} \\
j & \mathcal{B} \\
n & \mathcal{A} \wedge \mathcal{B} \qquad \text{∧I } i, j
\end{array}
$$

To show that ∧I is rule-sound, we assume that every line before line n is shiny; and we aim to show that line n is shiny, i.e. that $\Delta_n \vDash \mathcal{A} \wedge \mathcal{B}$.

So, let v be any valuation that makes all of Δ_n true.

I first show that v makes $\mathcal{A}$ true. To prove this, note that all of Δ_i are among Δ_n. By hypothesis, line i is shiny. So any valuation that makes all of Δ_i true makes $\mathcal{A}$ true. Since v makes all of Δ_i true, it makes $\mathcal{A}$ true too.

We can similarly see that v makes $\mathcal{B}$ true.

So v makes $\mathcal{A}$ true and v makes $\mathcal{B}$ true. Consequently, v makes $\mathcal{A} \wedge \mathcal{B}$ true. So any valuation that makes all of the sentences among Δ_n true also makes $\mathcal{A} \wedge \mathcal{B}$ true. That is: line n is shiny.   ∎

All of the remaining lemmas establishing rule-soundness will have, essentially, the same structure as this one did.

Lemma 5.5. ∧E is rule-sound.

Proof. Assume that every line before line n on some TFL-proof is shiny, and that ∧E is used on line n. So the situation is:

$$
\begin{array}{c|l}
i & \mathcal{A} \wedge \mathcal{B} \\
n & \mathcal{A} \qquad \text{∧E } i
\end{array}
$$

(or perhaps with $\mathcal{B}$ on line n instead; but similar reasoning will apply in that case). Let v be any valuation that makes all of Δ_n true. Note that all of Δ_i are among Δ_n. By hypothesis, line i is shiny. So any valuation that makes all of Δ_i true makes $\mathcal{A} \wedge \mathcal{B}$ true. So v makes $\mathcal{A} \wedge \mathcal{B}$ true, and hence makes $\mathcal{A}$ true. So $\Delta_n \vDash \mathcal{A}$.   ∎

Lemma 5.6. ∨I is rule-sound.

Proof. I leave this as an exercise.   ∎

Lemma 5.7. ∨E is rule-sound.

Proof. Assume that every line before line n on some TFL-proof is shiny, and that ∧E is used on line n. So the situation is:

$$
\begin{array}{r|l}
m & \mathcal{A} \lor \mathcal{B} \\
i & \quad \mathcal{A} \\
j & \quad \mathcal{C} \\
k & \quad \mathcal{B} \\
l & \quad \mathcal{C} \\
n & \mathcal{C} \qquad \lor\text{E } m,\, i\text{–}j,\, k\text{–}l
\end{array}
$$

Let v be any valuation that makes all of Δ_n true. Note that all of Δ_m are among Δ_n. By hypothesis, line m is shiny. So any valuation that makes Δ_n true makes $\mathcal{A} \lor \mathcal{B}$ true. So in particular, v makes $\mathcal{A} \lor \mathcal{B}$ true, and hence either v makes $\mathcal{A}$ true, or v makes $\mathcal{B}$ true. We now reason through these two cases:

> *Case 1: v makes $\mathcal{A}$ true.* All of Δ_i are among Δ_n, with the possible exception of $\mathcal{A}$. Since v makes all of Δ_n true, and also makes $\mathcal{A}$ true, v makes all of Δ_i true. Now, by assumption, line j is shiny; so $\Delta_j \vDash \mathcal{C}$. But the sentences Δ_i are just the sentences Δ_j, so $\Delta_i \vDash \mathcal{C}$. So, any valuation that makes all of Δ_i true makes $\mathcal{C}$ true. But v is just such a valuation. So v makes $\mathcal{C}$ true.
>
> *Case 2: v makes $\mathcal{B}$ true.* Reasoning in exactly the same way, considering lines k and l, v makes $\mathcal{C}$ true.

Either way, v makes $\mathcal{C}$ true. So $\Delta_n \vDash \mathcal{C}$. ∎

Lemma 5.8. $\perp$I is rule-sound.

Proof. Assume that every line before line n on some TFL-proof is shiny, and that $\perp$I is used on line n. So the situation is:

$$
\begin{array}{r|l}
i & \mathcal{A} \\
j & \neg\mathcal{A} \\
n & \perp \qquad \perp\text{I } i,\, j
\end{array}
$$

Note that all of Δ_i and all of Δ_j are among Δ_n. By hypothesis, lines i and j are shiny. So any valuation which makes all of Δ_n true would have to make both $\mathcal{A}$ and $\neg\mathcal{A}$ true. But no valuation can do that. So no valuation makes all of Δ_n true. So $\Delta_n \vDash \perp$, vacuously. ∎

Lemma 5.9. $\perp$E is rule-sound.

Proof. I leave this as an exercise. ∎

Lemma 5.10. $\neg$I is rule-sound.

Proof. Assume that every line before line n on some TFL-proof is shiny, and that $\neg$I is used on line n. So the situation is:

$$
\begin{array}{c|c}
i & \quad \mathcal{A} \\
& \quad \rule{1.2cm}{0.4pt} \\
j & \quad \bot \\
\hline
n & \neg\mathcal{A} \qquad\qquad \neg\text{I } i\text{–}j
\end{array}
$$

Let v be any valuation that makes all of Δ_n true. Note that all of Δ_n are among Δ_i, with the possible exception of $\mathcal{A}$ itself. By hypothesis, line j is shiny. But no valuation can make '$\bot$' true, so no valuation can make all of Δ_j true. Since the sentences Δ_i are just the sentences Δ_j, no valuation can make all of Δ_i true. Since v makes all of Δ_n true, it must therefore make $\mathcal{A}$ false, and so make $\neg\mathcal{A}$ true. So $\Delta_n \vDash \neg\mathcal{A}$. ∎

Lemma 5.11. TND is rule-sound.

Proof. Assume that every line before line n on some TFL-proof is shiny, and that TND is used on line n. So the situation is:

$$
\begin{array}{c|c}
i & \quad \mathcal{A} \\
& \quad \rule{1.2cm}{0.4pt} \\
j & \quad \mathcal{B} \\
k & \quad \neg\mathcal{A} \\
& \quad \rule{1.2cm}{0.4pt} \\
l & \quad \mathcal{B} \\
\hline
n & \mathcal{B} \qquad\qquad \text{TND } i\text{–}j,\ k\text{–}l
\end{array}
$$

Let v be any valuation that makes all of Δ_n true. Either v makes $\mathcal{A}$ true, or it makes $\mathcal{A}$ false. We shall reason through these (sub)cases separately.

> *Case 1: v makes $\mathcal{A}$ true.* The sentences Δ_i are just the sentences Δ_j. By hypothesis, line j is shiny. So $\Delta_i \vDash \mathcal{B}$. Furthermore, all of Δ_i are among Δ_n, with the possible exception of $\mathcal{A}$. So v makes all of Δ_i true, and hence makes $\mathcal{B}$ true.
>
> *Case 2: v makes $\mathcal{A}$ false.* Then v makes $\neg\mathcal{A}$ true. Reasoning in exactly the same way as above, considering lines k and l, v makes $\mathcal{B}$ true.

Either way, v makes $\mathcal{B}$ true. So $\Delta_n \vDash \mathcal{B}$. ∎

Lemma 5.12. $\rightarrow$I, $\rightarrow$E, $\leftrightarrow$I, and $\leftrightarrow$E are all rule-sound.

Proof. I leave these as exercises. ∎

This establishes that all the basic rules of our proof system are rule-sound. Finally, we show:

Lemma 5.13. All of the derived rules of our proof system are rule-sound.

Proof. Suppose that we used a derived rule to obtain some sentence, $\mathcal{A}$, on line n of some TFL-proof, and that every earlier line is shiny. Every use of a derived rule can be replaced (at the cost of long-windedness) with multiple uses of basic rules. (This was proved in forall$x$:Cambridge §30.) That is to say, we could have used basic rules to write $\mathcal{A}$ on some line $n + k$, without introducing any further assumptions. So, applying Lemmas 5.3–5.12

several times ($k + 1$ times, in fact), we can see that line $n + k$ is shiny. Hence the derived rule is rule-sound. ∎

And that's that! We have shown that every rule – basic or otherwise – is rule-sound, which is all that we required to establish Lemma 5.2, and hence the Soundness Theorem (see §5.1).

But it might help to round off this chapter if I repeat my informal explanation of what we have done. A formal proof is just a sequence – of arbitrary length – of applications of rules. We have spent this section showing that any application of any rule will not lead you astray. It follows (by induction) that no formal proof will lead you astray. That is: our proof system is sound.

Practice exercises

A. Complete the Lemmas left as exercises in this chapter. That is, show that the following are rule-sound:

- ∨I. (*Hint*: this is similar to the case of ∧E.)
- ⊥E. (*Hint*: this is similar to the case of ⊥I.)
- →I. (*Hint*: this is similar to ∨E.)
- →E.
- ↔I.
- ↔E.

Completeness 6

Our formal proof system is *sound*: it will never lead us astray. In this chapter, I shall show it is *complete*: every entailment is captured by some formal proof.

6.1 Completeness defined, and setting up the proof

We say that a formal proof system is COMPLETE (relative to a given semantics) *iff* whenever Γ entails C, there is some formal proof of C whose assumptions are among Γ. To prove that TFL's proof system is complete, I need to show:

Theorem 6.1. Completeness. For any sentences Γ and C: if $\Gamma \vDash C$, then $\Gamma \vdash C$.

This is the converse of the Soundness Theorem. Together, these two theorems vindicate our proof system entirely: it doesn't prove too much (it is sound); it doesn't prove too little (it is complete); it is just right.

As with the Soundness Theorem, the proof of the Completeness Theorem turns on a single important lemma. Here is that Lemma:

Lemma 6.2. For any sentences Δ: if $\Delta \nvdash \bot$, then there is a valuation that makes all of Δ true.

I shall explain how we might prove Lemma 6.2 in due course. But, once we have proved Lemma 6.2, the Completeness Theorem will follow almost immediately:

Proof of the Completeness Theorem, given Lemma 6.2. We shall prove the contrapositive, i.e. if $\Gamma \nvdash C$, then $\Gamma \nvDash C$.

So, suppose that $\Gamma \nvdash C$. Now, if there were a proof of '$\bot$' from Γ with $\neg C$, we could manipulate this proof (using ¬I and DNE) to obtain a proof of C from Γ, contradicting our supposition. So it must be that:

$$\Gamma, \neg C \nvdash \bot$$

Now, by Lemma 6.2, there is some valuation that makes all of Γ and $\neg C$ true. Hence $\Gamma \nvDash C$. ∎

All the hard work, then, will go into proving Lemma 6.2. And – fair warning – that *will* be hard work. In fact, I shall prove Lemma 6.2 in two different ways. One of these ways is more concrete, and gives us quite a lot of information; the other is more compressed and abstract, but also a bit more general.

6.2   An algorithmic approach

I start with the more concrete approach to Lemma 6.2. This works by specifying an algorithm, SimpleSearch, which determines, for any arbitrary sentences, whether or not those sentences are jointly contrary. (Iam assuming that we are given only *finitely many* sentences. I shall revisit this point in §6.4, but shall not mention it much until then.) Moreover, SimpleSearch two very desirable properties. For any sentences, Δ, that we feed in:

(A1) if Δ are jointly contrary, then the algorithm outputs a TFL-proof which starts with Δ as assumptions and terminates in '$\bot$' on just those assumptions; and

(A2) if Δ are not jointly contrary, then the algorithm yields a valuation which makes all of Δ true together.

Clearly, any algorithm with both of these properties is independently interesting. But the important thing, for our purposes, is that if we can prove that there is an algorithm with property (A2), then we will have proved Lemma 6.2 (pause for a moment to check this!) and hence the Completeness Theorem. So, my aim is to specify an algorithm, and then prove that it has these properties.

First steps towards an algorithm

To say that $\Delta \vdash \bot$ is to say that there is a TFL-proof, whose assumptions are all among Δ, and which terminates in '$\bot$' without any additional assumptions.

It obvious that, if '$\bot$' appears anywhere in a TFL-proof, then the TFL-proof involves an application of the rule $\bot$I. So our algorithm will have to look for situations in which we might be able to apply the rule $\bot$I.

Of course, to show that some sentences are jointly contrary, it is not enough simply to find a situation in which we can apply $\bot$I. For if we only obtain '$\bot$' within the scope of *additional* assumptions, we will not have shown that our *original* sentences are jointly contrary. Fortunately, there are rules that allow us to take instances of '$\bot$' that appear in the scope of additional assumptions, and derive an instance of '$\bot$' that appears within the scope of fewer assumptions. As an example, consider the last line of this TFL-proof:

1	$A \lor B$	
2	$(\neg A \land \neg B)$	
3	$\neg A$	$\land$E 2
4	$\neg B$	$\land$E 2
5	$\quad A$	
6	$\quad \bot$	$\bot$I 5, 3
7	$\quad B$	
8	$\quad \bot$	$\bot$I 7, 4
9	$\bot$	$\lor$E 1, 5–6, 7–8

In fact, this TFL-proof is exactly the sort of proof that our algorithm will look for. In particular, the core of our algorithm involves only the techniques used in this proof: ∧E, introducing assumptions, using ⊥I within the scope of those assumptions, and then using ∨E to manipulate the instances of '⊥'.

Illustrations of **SimpleSearch**, assuming CNF

In order to keep things simple, at this stage *I shall assume that the sentences that we want to test for joint contrariness are all in CNF.* (For a reminder of what this means, see §3.5.) I shall explain why this assumption is harmless towards the end of this section.

Given that we are testing some CNF sentences for joint contrariness, our first task is of course to start a TFL-proof with all of these sentences as assumptions.

Our second task is equally plain. Since the sentences are in CNF, we should eliminate all of the conjunctions that occur in them. More precisely: we should apply ∧E repeatedly, until no further application of ∧E would yield a sentence that we have not already obtained.

Given that our original sentences were in CNF, the output of all of these applications of ∧E will be a series of sentences containing only negations and disjunctions, where the scope of every occurrence of negation is minimal. Otherwise put: we will now be dealing with a bunch of sentences that are either atomic sentences, negations of atomic sentences, or disjunctions. Here is where things get interesting, and it will help us to have an example. So, suppose we start with the following sentences:

$$
\begin{array}{l|l}
1 & (A \lor B) \\
2 & (\neg B \lor \neg C) \\
3 & \neg A \\
4 & C
\end{array}
$$

There are no conjunctions to deal with, so our next task is to position ourselves so that we might be able to use ∨E further down the line. In particular, we must try to position ourselves to use ∨E with a citation that mentions line 1, and then two subproofs, one starting with the assumption '*A*', the other starting with the assumption that '*B*'. So we now continue the TFL-proof as follows (I have added some annotations to explain our motivations):

$$
\begin{array}{l|l}
5 & \quad A \qquad \text{want to use ∨E with line 1} \\
  & \\
  & \quad B \qquad \text{want to use ∨E with line 1} \\
\end{array}
$$

When we set ourselves up for potential uses of ∨E in this way, we shall henceforth say that the disjunction in question (here, the disjunction on line 1) has been *expanded*. Note that, when we expand a disjunction by introducing new assumptions, we must leave gaps

underneath these new assumptions. We will fill in these gaps as we continue to implement our algorithm. But, since we do not yet know how these gaps will be filled in, we cannot yet assign definite numbers to all the lines of our proof. So what we have written down cannot any longer be described as a TFL-proof; rather, it is a PROOF-SKELETON. This is how I shall refer to it, in what follows.

We have expanded our first disjunction, thereby positioning ourselves for a (potential) later use of ∨E with line 1. But there is another disjunction on line 2. We need to expand this disjunction too, positioning ourselves for a (potential) later use of ∨E with line 2. At this point, however, some care is needed. We have just made two new assumptions – '*A*', and '*B*', separately – and we might be able to use ∨E with line 2 inside the scope of *either* assumption. So we must continue our proof-skeleton as follows:

5		*A*	want to use ∨E with line 1
6		¬*B*	want to use ∨E with line 2
		¬*C*	want to use ∨E with line 2
		B	want to use ∨E with line 1
		¬*B*	want to use ∨E with line 2
		¬*C*	want to use ∨E with line 2
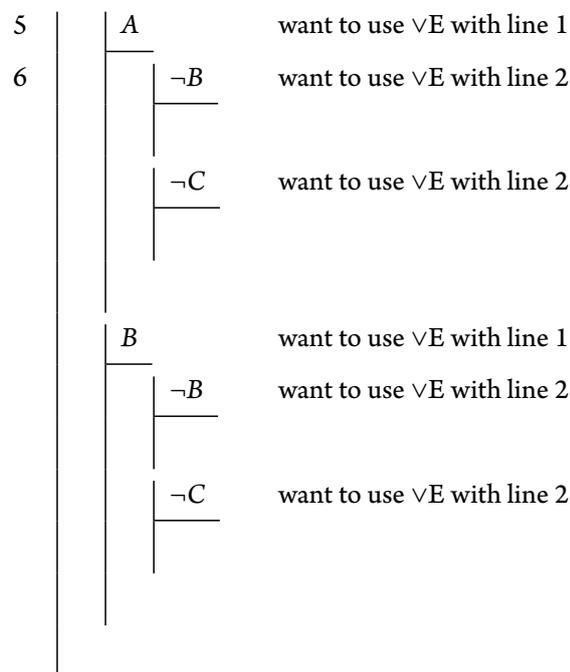
At this point, all the disjunctions in our proof-skeleton have been expanded. So we can move on to the next task.

Recall that we started with sentences in CNF. We first deal with all the conjunctions (by ∧E). We then unpacked all the disjunctions. So our proof-skeleton will now have lots of lines on it that contain simply an atomic sentence, or its negation. So the next task is to apply the rule ⊥I wherever possible. In our example, we can apply it four times:

5	A		want to use ∨E with line 1
6		$\neg B$	want to use ∨E with line 2
7		$\bot$	⊥I 5, 3
8		$\neg C$	want to use ∨E with line 2
9		$\bot$	⊥I 5, 3
i	B		want to use ∨E with line 1
$i+1$		$\neg B$	want to use ∨E with line 2
$i+2$		$\bot$	⊥I i, $i+1$
$i+3$		$\neg C$	want to use ∨E with line 2
$i+4$		$\bot$	⊥I 4, $i+3$

Recall that our aim is to end up with a proof of '⊥' on no assumptions. So the most we can hope for, from each assumption, is that it generates '⊥'. When an assumption does generate '⊥', we will have no further use for it, except in a potential application of ∨E. So, at this point, we can start to close some of the gaps in our proof-skeleton, and put some flesh on its bones by adding in line numbers. I have done this in the above example. I have also added in some *dummy line numbers*, lines i through to $i+4$, for ease of reference.)

Given what I have just said, we must next apply ∨E as often as we can:

5	A		want to use ∨E with line 1
6		$\neg B$	want to use ∨E with line 2
7		$\bot$	⊥I 5, 3
8		$\neg C$	want to use ∨E with line 2
9		$\bot$	⊥I 5, 3
10	$\bot$		∨E 2, 6–7, 8–9
11	B		want to use ∨E with line 1
12		$\neg B$	want to use ∨E with line 2
13		$\bot$	⊥I 10, 11
14		$\neg C$	want to use ∨E with line 2
15		$\bot$	⊥I 4, 13
16	$\bot$		∨E 2, 12–13, 14–15

We have managed to apply it twice here, citing line 2 on both occasions. Again, we can close some gaps and fill in some line numbers (removing the dummy line numbers where appropriate). And now we can see that we are able to apply ∨E again:

> 17 | ⊥   ∨E 1, 5–10, 11–16

And now we have generated a complete TFL-proof of '⊥' from our initial assumptions. That is, *our initial assumptions were jointly contrary*.

Now, the preceding TFL-proof is not as slick as it could have been (in particular, steps 5–10 took longer than they needed to). But the important point is that we did, indeed, end up with a proof of '⊥' from our initial assumptions, without really having to *think*.

This is our first run-through of our algorithm, SimpleSearch. To get a better – but still informal – idea of how SimpleSearch works, I shall now consider an application of Simple-Search to some sentences that are *not* jointly contrary. Suppose we start with the following sentences:

> 1 | $(A \lor (\neg B \lor C))$
>
> 2 | $(\neg A \lor \neg C)$
>
> 3 | B

There are no conjunctions to deal with, so we immediately move to the task of expanding the disjunctions on line 1 and line 2. This yields:

> 4 | | A want to use ∨E with line 1
>
> 5 | | | $\neg A$ want to use ∨E with line 2
>
> | | | $\neg C$ want to use ∨E with line 2
>
> *i* | | $(\neg B \lor C)$ want to use ∨E with line 1
>
> | | | $\neg A$ want to use ∨E with line 2
>
> | | | $\neg C$ want to use ∨E with line 2

Again, I have inserted a dummy line-number for ease of reference. And this line merits comment since, in expanding the disjunction on line 1, we have generated a *new* disjunction on line *i*. So, even though we have expanded all the disjunctions we *began* with, we still have an unexpanded disjunction. This, too, must be expanded. But at this point, we

must take particular care. Our aim, in expanding a disjunction, is to set ourselves up for all possible future uses of ∨E. However, we plainly would not be able to use ∨E with i at any line *earlier* than line i. Consequently, we should expand that disjunction by continuing our proof-skeleton as follows:

4	A	want to use ∨E with line 1
5	$\neg A$	want to use ∨E with line 2
	$\neg C$	want to use ∨E with line 2
i	$(\neg B \lor C)$	want to use ∨E with line 1
	$\neg A$	want to use ∨E with line 2
	$\neg B$	want to use ∨E with line i
	C	want to use ∨E with line i
	$\neg C$	want to use ∨E with line 2
	$\neg B$	want to use ∨E with line i
	C	want to use ∨E with line i

And, at this point, we have expanded *all* the disjunctions that occur on our proof. So we move on to the next stage of our algorithm: we apply ⊥I wherever we can; and then we apply ∨E wherever we can. The final result (adding in a few more dummy line-numbers) is:

1	$(A \lor (\neg B \lor C))$	
2	$(\neg A \lor \neg C)$	
3	B	
4	A	want to use $\lor$E with line 1
5	$\neg A$	want to use $\lor$E with line 2
6	$\bot$	$\bot$I 4, 5
7	$\neg C$	want to use $\lor$E with line 2
i	$(\neg B \lor C)$	want to use $\lor$E with line 1
$i + 1$	$\neg A$	want to use $\lor$E with line 2
$i + 2$	$\neg B$	want to use $\lor$E with line i
$i + 3$	$\bot$	$\bot$I 3, $i + 2$
$i + 4$	C	want to use $\lor$E with line i
k	$\neg C$	want to use $\lor$E with line 2
$k + 1$	$\neg B$	want to use $\lor$E with line i
$k + 2$	$\bot$	$\bot$I 3, $k + 1$
$k + 3$	C	want to use $\lor$E with line i
$k + 4$	$\bot$	$\bot$I $k + 3$, k
$k + 5$	$\bot$	$\lor$E 2, $k + 1$–$k + 2$, $k + 3$–$k + 4$

At this point, SimpleSearch has run its course. But the output is not a TFL-proof of '$\bot$' from the initial assumptions: it is a proof-skeleton with several gaps in it. I claim that I can infer from this that the initial assumptions are *not* jointly contrary. I further claim that there is a valuation that makes all of the initial assumptions true. Moreover, I claim that I can read off this valuation directly from the proof-skeleton!

I shall *justify* all of these claims in a moment. First, though, let me explain how to read off the valuation from the proof-skeleton. Let us say that a line in a proof-skeleton is OPEN iff there is gap underneath it.[1] In my example, lines 7 and $i + 4$ are both open. I choose

---

[1]More precisely a line is open *iff* it meets all of the following four conditions. (i) The line is not '$\bot$'. (ii) No instance of $\land$E occurs after the line. (iii) The next line in the proof-skeleton (if there is one) is not '$\bot$'. (iv)

one of these open lines to work with; it might as well be the first, i.e. line 7. I then make a list of all the atomic sentences and negations of atomic sentences on lines that either occur on the open line or that occur on earlier lines (but not in a subproof that has been closed before the open line). So, working with line 7 of my example, my list is:[2]

$$\neg C, A, B$$

There is, indeed, a unique valuation which makes all of these sentences true. And this valuation will make all of my original assumptions true too. In the case of my example, this is easily verified.

In a moment, I shall prove that this was not just good fortune, but that SimpleSearch is guaranteed to behave in this way.

Relaxing the assumption of CNF

First, I need to explain why it was harmless to restrict our attention to sentences in CNF. The simple reason for this is that, given any sentences, Δ, we can always find some equivalent sentences, Δ^*, in CNF.

However, we need to be a little careful here. Given any sentences Δ, the CNF Theorem of §3.5 tells us that there are *tautologically* equivalent sentences, Δ^*, in CNF. But, given our aims, we cannot afford to rely upon this fact. After all, if TFL's proof-system cannot itself *prove* the equivalences between Δ and Δ^*, then manipulating Δ^* with TFL's proof-system will teach us *nothing* about the proof-theoretic consequences of Δ. And what we want is an algorithm which decides whether or not Δ are jointly contrary.

The upshot of this is that we need to prove a *proof-theoretic* analogue of the CNF Theorem.

Lemma 6.3. For any sentence, there is a *provably* equivalent sentence in CNF. Moreover, there is an algorithm for determining these equivalents.

Proof sketch. First, prove a proof-theoretic analogue of Lemma 2.2. That is, show the following: If $C \dashv\vdash S$, then $\mathcal{A} \dashv\vdash \mathcal{A}[S/C]$.

Then offer proof-theoretic analogues of Lemmas Lemmas 3.2–3.3, noting that each of the tautological equivalences appealed to (namely, Double Negation Elimination, the De Morgan Laws, and the Distribution Laws) are also *provable* equivalences. (Indeed, all but the Distribution Laws are *derived* rules of TFL's proof-system.) ∎

So, given any sentences, we can *first* put them into CNF, and then apply the steps that were sketched just earlier: apply $\wedge$E, expand disjunctions, apply $\bot$I, and apply $\vee$E. Let's make this precise.

Statement of **SimpleSearch**

Given any sentences Δ:

---

the next line in the proof-skeleton (if there is one) does not depend upon all of the assumptions of the former line.

[2]If I had worked instead with line $i + 4$, my list would have been: $C, \neg A, B$.

Step 1. Start a TFL-proof, with Δ as your assumptions.

Step 2. For each sentence $\mathcal{D}$ among Δ, apply the algorithm which finds a CNF-sentence $\mathcal{D}^*$ such that $\mathcal{D}^* \dashv\vdash \mathcal{D}$. This will leave you with a single TFL-proof where each $\mathcal{D}^*$ appears on some line.

Step 3: Now apply ∧E as often as you can to lines of your TFL-proof that are in CNF, until any further application of ∧E would yield a sentence already on your proof.

Step 4: Expand the first disjunction in CNF which you have not yet expanded. More precisely: let $(\mathcal{A} \vee \mathcal{B})$ be the first disjunction in CNF that you have not yet expanded and let *l* be any line in your proof-skeleton that is both beneath the line on which the disjunction occurs; and which depends upon (at least) the assumptions on which the disjunction occurs. Then, beneath line *l*, introduce two new assumptions – $\mathcal{A}$ and $\mathcal{B}$ – separately, and with an empty line beneath both. At this point, we say that the disjunction '$(\mathcal{A} \vee \mathcal{B})$' has been *expanded*.

Repeat this step until every disjunction in CNF has been expanded.

Step 5: Go through every *open* line and check to see whether ⊥I can be applied. If so, apply it (filling in line numbers).

Step 6: Go through every *open* line and check to see whether ∨E can be applied (allowing you to generate a line with '⊥' on it). If so, apply it (filling in line numbers). Repeat this step until ∨E cannot be applied any further.

Step 7: Output the proof or proof-skeleton. This terminates SimpleSearch.

This is the algorithm SimpleSearch.

Proving that **SimpleSearch** behaves as required

I now want to prove that SimpleSearch has the properties (A1) and (A2), which I advertised for it.

The very first thing I need to do is to make sure that, when we run SimpleSearch, there is no chance of our getting trapped in an infinite loop. That is, we need a guarantee that SimpleSearch will output something *in the end*, regardless of the inputs we feed it.

Lemma 6.4. SimpleSearch is guaranteed to terminate, on any inputs.

Proof sketch. Step 1 is guaranteed to terminate eventually, since it just involves writing down finitely many sentences.

Step 2 is similarly guaranteed to terminate, since it only ever takes finitely many steps to generate a CNF equivalent from each of the finitely many assumptions.

Step 3 is guaranteed to terminate, since are only finitely many conjunction signs among the CNF sentences you have written down, and so you only need to apply ∧E finitely many times.

Step 4 is the tricky step. The worry is that, in expanding a disjunction (in CNF), you may have to introduce *many* new assumptions; and if one of the disjuncts is itself a disjunction, this can increase the total number of unexpanded disjunctions.[3] However, when we expand a disjunction, each of the disjuncts that we introduce as an assumption is *shorter*

---

[3]To see this in action, see what happens when your initial assumptions are '$(A \vee B)$' and '$((C \vee D) \vee (D \vee E))$'.

than the newly-expanded disjunction. So, even though expanding a disjunction may increase the number of unexpanded disjunctions, the unexpanded disjunctions keep getting shorter. Eventually, then, these will dwindle into atomic sentences or negations of atomic sentences, so that every disjunction has been expanded.[4]

Step 5 is guaranteed to terminate, since as Step 4 terminated, there will only be finitely many open lines, and inspecting whether or not $\perp$I can be applied on a given line is a bounded procedure (you only need to inspect all the earlier lines). Step 6 is similarly guaranteed to terminate. ∎

So we now know that, given any inputs, SimpleSearch *will* yield an output. We next need to confirm that we can *trust* that output.

Lemma 6.5. If SimpleSearch outputs a proof or proof-skeleton whose last line is '$\perp$' on no assumptions other than Δ, then $\Delta \vdash \perp$

Proof sketch. By inspection, it is clear that SimpleSearch simply applies the rules of the TFL-proof system (in a very specific order) to a proof that starts with Δ. So if the proof-skeleton's last line is '$\perp$' on no assumptions other than Δ, then SimpleSearch has simply outputted a TFL-proof that $\Delta \vdash \perp$. ∎

Lemma 6.6. If SimpleSearch outputs a proof or proof-skeleton whose last line is *not* '$\perp$' on no assumptions other than Δ, then:

1. there is at least one open line on the proof-skeleton; and
2. every open line yields (in the manner described) a valuation which makes all of Δ true; and
3. $\Delta \nvdash \perp$.

Proof. I shall prove each point separately, on the assumption that SimpleSearch terminates by outputting a proof-skeleton whose last line is *not* '$\perp$' on no assumptions other than Δ.
   Proof of (1). There are two cases to consider:

   Case 1: *None of the CNF-sentences that we generated in Step 2 of* SimpleSearch *contained any occurrences of '$\lor$'.* Then, since the last line of our proof is not '$\perp$', the last line of our proof-skeleton is open.
   Case 2: *At least one of the CNF-sentences that we generated in Step 2 of* SimpleSearch *contained an occurrence of '$\lor$'.* Then, if every assumption that we made when expanding disjunctions led to '$\perp$', we would have applied $\lor$E repeatedly until we obtained '$\perp$' on no assumptions. So at least one assumption did not lead to '$\perp$', and so there is an open line.

   Proof of (2). Let Ω be the sentences that we read off (in the manner described earlier) from an open line on our proof (whose existence is guaranteed by point (1)). They are all atomic sentences and their negations. Now, since the line in question is open, this means that we cannot find both a sentence and its negation among Ω, since otherwise we could have applied $\perp$I. Hence there is a valuation, v, which makes all of them true.[5]

---

[4]One way to make this rigorous would be to prove, by induction, that if Δ contains exactly n instances of '$\lor$', Step 3 will require you to introduce at most $2^{n+1} - 2$ assumptions.
[5]For details of this last step, see Lemma 6.8, below.

Now let $\mathcal{D}$ be any sentence among Δ; I claim that v makes $\mathcal{D}$ true. To prove this claim, I first consider the CNF sentence $\mathcal{D}^*$ generated in Step 1 of SimpleSearch, such that $\mathcal{D} \dashv\vdash \mathcal{D}^*$. It is easy to show that, for every (conjoined) disjunction in $\mathcal{D}^*$, at least one disjunct of that disjunction appears in Ω; it follows that each of the disjuncts is made true by v, so that $\mathcal{D}^*$ is made true by v (I leave the details of this as an exercise). Now, since $\mathcal{D}^* \vdash \mathcal{D}$, the Soundness Theorem tells us that $\mathcal{D}^* \vDash \mathcal{D}$; so v must make $\mathcal{D}$ true too.

Proof of (3). Suppose, for reductio, that $\Delta \vdash \bot$. Then, by the Soundness Theorem, there is no valuation that makes all of the Δ true. This contradicts what we have just shown in point (2). ∎

And now – at long last – since we can trust the outputs of SimpleSearch, we can show that SimpleSearch has properties (A1) and (A2):

Theorem 6.7. Given any sentences Δ:

1. If $\Delta \vdash \bot$, then SimpleSearch yields a TFL-proof which starts with Δ as assumptions and terminates in '$\bot$'.
2. If $\Delta \nvdash \bot$, then SimpleSearch shows us as much and yields a valuation which makes all of Δ true.

Proof. Suppose $\Delta \vdash \bot$. Then by Lemma 6.6, SimpleSearch does *not* output a proof or proof-skeleton whose last line is not '$\bot$' on no assumptions other than Δ. But by Lemma 6.4, SimpleSearch does terminate. So the last line of the output must be '$\bot$', on no assumptions other than Δ, which, by Lemma 6.5, is an explicit TFL-proof that $\Delta \vdash \bot$.

Suppose $\Delta \nvdash \bot$. Then by Lemma 6.5, SimpleSearch does not output a proof-skeleton whose last line is '$\bot$' on no assumptions other than Δ. Now apply Lemma 6.6. ∎

And with this result, we have achieved what we set out to achieve! But it might be worth reminding ourselves *why* we set out on this course. Our target for this chapter was to prove the Completeness Theorem. To prove that, we needed to prove Lemma 6.2. But Theorem 6.7 immediately entails Lemma 6.2 (just look at the second condition). So we have proved the Completeness Theorem, and acquired a useful little algorithm along the way!

6.3   A more abstract approach: polarising sentences

We could stop now. But I want to offer a second proof of the Completeness Theorem. It is more abstract, more general, and in a sense, more direct.

From the perspective of proving the Completeness Theorem, it is massive overkill to deploy SimpleSearch. For our purposes, all that matters is the following: when $\Delta \nvdash \bot$, there are some atomic sentences and negations of atomic sentences, Ω, from which we can read off a valuation that makes all of Δ true. From an abstract point of view, we needn't care how SimpleSearch *generates* Ω; only that Ω exists. The second proof of the Completeness Theorem builds upon this observation.

It will make our lives easier if we introduce a bit more terminology. Where Ω are some sentences, say that Ω are JOINTLY POLARISING *iff* both:[6]

---

[6]NB: this terminology is not standard among logicians.

(P1) every sentence among Ω is either an atomic sentence or a negation of an atomic sentence; and

(P2) if a sentence is among Ω, then its negation is not among Ω

Where Ω are jointly polarising, say that $\mathcal{A}$ is IN THE FIELD OF Ω iff $\mathcal{A}$ is an atomic sentence such that either $\mathcal{A}$ or $\neg\mathcal{A}$ is among Ω.

It is worth pausing for a moment to think about how this terminology, and this result, connects with SimpleSearch. When $\Delta \nvdash \bot$, SimpleSearch generates a proof-skeleton with at least one open line. From each open line, we can read off some jointly polarising sentences, Ω. Moreover, when $\mathcal{D}$ is any sentence among Δ, every atomic subsentence of $\mathcal{D}$ is in the field of Ω. At this point, we claimed (in the proof of Lemma 6.6) that there is bound to be some valuation that makes all of Ω true. That was right, and the more general result is simply this very easy observation.

Lemma 6.8. Suppose Ω are jointly polarising. Then there is exactly one valuation of the sentences in the field of Ω that makes every sentence among Ω true.

Proof. Define a valuation, v, as follows: for each sentence $\mathcal{A}$ in the field of Ω,

- v makes $\mathcal{A}$ true if $\mathcal{A}$ is among Ω
- v makes $\mathcal{A}$ false if $\neg\mathcal{A}$ is among Ω

Given (P2), v is well-defined. Now, if $\mathcal{B}$ is a sentence among Ω, then $\mathcal{B}$ is either $\mathcal{A}$ or $\neg\mathcal{A}$, for some sentence $\mathcal{A}$ in the field of Ω. Either way, v makes $\mathcal{B}$ true. Moreover, any valuation that assigned a different value to $\mathcal{A}$ would make $\mathcal{B}$ false. Hence v is unique. ∎

Continuing, though, to think through how SimpleSearch relates to this terminology: the next stage in the proof of Lemma 6.6 was for us to show that the (unique) valuation that makes all of Ω true also makes all of Δ true too. This required a bit of work (some of which I left as an exercise). But let me now part company with the proof of Lemma 6.6, and instead make a much more simple observation.

We are considering a situation where all of the atomic subsentences of some sentence (in this case, $\mathcal{D}$) are in the field of some jointly polarising sentences (in this case, Ω). We can obtain a very easy and general result about such situations.

Lemma 6.9. Suppose Ω are jointly polarising and that every atomic subsentence of $\mathcal{A}$ is in the field of Ω. Then either $\Omega \vDash \mathcal{A}$ or $\Omega \vDash \neg\mathcal{A}$.

Proof. Using Lemma 6.8, let v be the unique valuation of the sentences in the field of Ω that makes all of Ω true . Either v makes $\mathcal{A}$ true, or v makes $\mathcal{A}$ false. If v makes $\mathcal{A}$ true, then since v is the only valuation that makes all of Ω true, $\Omega \vDash \mathcal{A}$. If v makes $\mathcal{A}$ false, then it makes $\neg\mathcal{A}$ true, and similar reasoning shows that $\Omega \vDash \neg\mathcal{A}$. ∎

Now, when we are in the business of proving the Completeness Theorem, our aim is to tie '$\vDash$' to '$\vdash$'. And so, in this regard, it is very striking that there is a perfect *proof-theoretic* analogue of Lemma 6.9:

Lemma 6.10. Suppose Ω are jointly polarising and that every atomic subsentence of $\mathcal{A}$ is in the field of Ω. Then either $\Omega \vdash \mathcal{A}$ or $\Omega \vdash \neg\mathcal{A}$.

Proof. By (strong) induction on the length of $\mathcal{A}$.

Fix any sentence $\mathcal{A}$, and assume that all of $\mathcal{A}$'s atomic subsentences are in the field of Ω. For induction, suppose that, where $\mathcal{B}$ is any sentence shorter than $\mathcal{A}$ and where all of $\mathcal{B}$'s atomic subsentences are among the sentences that jointly span Ω, either $\Omega \vdash \mathcal{B}$ or $\Omega \vdash \neg\mathcal{B}$. There are now six cases to consider:

> *Case 1: $\mathcal{A}$ is atomic.* Then either $\mathcal{A}$ is among Ω (in which case $\Omega \vdash \mathcal{A}$, by the rule R), or $\neg\mathcal{A}$ is among Ω (in which case $\Omega \vdash \neg\mathcal{A}$, by the rule R). So the atomic case is straightforward.
>
> *Case 2: $\mathcal{A}$ is $\neg C$.* Suppose $\Omega \nvdash \mathcal{A}$, i.e. $\Omega \nvdash \neg C$. Then $\Omega \vdash C$, by the induction hypothesis. It follows that $\Omega \vdash \neg\neg C$ (just expand the proof of C from Ω by supposing $\neg C$ and then using $\bot$I and $\neg$I). So $\Omega \vdash \neg\mathcal{A}$. So either $\Omega \vdash \mathcal{A}$ or $\Omega \vdash \neg\mathcal{A}$.
>
> *Case 3: $\mathcal{A}$ is $(C \land D)$.* Suppose $\Omega \nvdash \mathcal{A}$, i.e. $\Omega \nvdash C \land D$. Then either $\Omega \nvdash C$, or $\Omega \nvdash D$. For if $\Omega \vdash C$ and $\Omega \vdash D$, then $\Omega \vdash C \land D$ (just combine the two proofs and then use $\land$I). So, by the induction hypothesis, either $\Omega \vdash \neg C$, or $\Omega \vdash \neg D$. Either way, $\Omega \vdash \neg(C \land D)$. (Showing this is left as an exercise.) That is, $\Omega \vdash \neg\mathcal{A}$. So either $\Omega \vdash \mathcal{A}$, or $\Omega \vdash \neg\mathcal{A}$.
>
> *Case 4: $\mathcal{A}$ is $(C \lor D)$.* Suppose $\Omega \nvdash \mathcal{A}$, i.e. $\Omega \nvdash C \lor D$. Then $\Omega \nvdash C$ and $\Omega \nvdash D$. So, by hypothesis, $\Omega \vdash \neg C$ and $\Omega \vdash \neg D$. So $\Omega \vdash \neg(C \lor D)$. (Showing this is left as an exercise). That is, $\Omega \vdash \neg\mathcal{A}$. So either $\Omega \vdash \mathcal{A}$, or $\Omega \vdash \neg\mathcal{A}$.
>
> *Case 5: $\mathcal{A}$ is $(C \to D)$.* I leave this as an exercise.
> *Case 6: $\mathcal{A}$ is $(C \leftrightarrow D)$.* I leave this as an exercise.

This completes the (strong) induction on length.   ∎

We here have a very tight connection between the behaviour of '⊨' and '⊢' in the case of jointly polarising sentences. In fact, if we like, we can prove a kind of mini-version of the Soundness and Completeness Theorems:

Lemma 6.11. Suppose Ω are jointly polarising and that every atomic subsentence of $\mathcal{A}$ is in the field of Ω. Then $\Omega \vdash \mathcal{A}$ iff $\Omega \vDash \mathcal{A}$.

Proof. Left-to-right. If $\Omega \vdash \mathcal{A}$, then $\Omega \vDash \mathcal{A}$ by the Soundness Theorem.

Right-to-left. I prove the contrapositive, i.e. that if $\Omega \nvdash \mathcal{A}$ then $\Omega \nvDash \mathcal{A}$. So suppose $\Omega \nvdash \mathcal{A}$. Then $\Omega \vdash \neg\mathcal{A}$, by Lemma 6.10. So $\Omega \vDash \neg\mathcal{A}$, by the Soundness Theorem. But in that case, $\Omega \nvDash \mathcal{A}$, since otherwise we would have that $\Omega \vDash (\mathcal{A} \land \neg\mathcal{A})$, contradicting Lemma 6.8.   ∎

What Lemma 6.11 tells us, in essence, is the following: if you pin down what is happening with atomic sentences and their negations, then you determine all the semantic and proof-theoretic consequences; moreover, these consequences *coincide*. (It is no wonder, then, that SimpleSearch worked precisely by breaking down sentences into pathways through TFL-proof-skeletons that consist of atomic sentences and their negations.) And my strategy for my second proof of the Completeness Theorem is to *lift* this result up into the more general case.

Here is the idea, in a bit more detail. Let Δ be any sentences such that $\Delta \nvdash \bot$. I will first generate some jointly polarising sentences, Ω, from Δ. Since Ω are jointly polarising,

there is a unique valuation which makes them all true, and, in turn, this makes all of Δ true. This will prove Lemma 6.2, from which the Completeness Theorem follows (see §6.1).

At this level of generality, the strategy is very similar to our first proof of the Completeness Theorem. But there is a difference. To generate the jointly polarising sentences from Δ, I shall appeal to the simple fact that *either* Δ are jointly contrary *or* Δ are not jointly contrary. In saying this, I am not relying on the existence of algorithm (such as SimpleSearch) to determine which possibility obtains; I am simply appealing to an instance of excluded middle (in the metalanguage). And to show that the jointly polarising sentences tautologically entail every sentence among Δ, I shall lean upon our very general observations about jointly polarising sentences.

Without further ado, then, here is the proof.

Proof of Lemma 6.2 without an algorithm. Suppose $\Delta \nvdash \bot$; let $\mathcal{A}_1, \mathcal{A}_2, ..., \mathcal{A}_m$ be all the atomic sentences that occur as subsentences in any of the sentences among Δ.[7] I generate my jointly polarising sentences by a recursive definition:

> *Base clause.* If $\Delta, \mathcal{A}_1 \nvdash \bot$, then let Ω_1 be $\mathcal{A}_1$. Otherwise, let Ω_1 be $\neg\mathcal{A}_1$
> *Recursion clause.* For each number n such that $1 < n < m$:
>
> - If $\Delta, \Omega_n, \mathcal{A}_{n+1} \nvdash \bot$, then let Ω_{n+1} be all of Ω_n and $\mathcal{A}_{n+1}$
> - Otherwise, let Ω_{n+1} be all of Ω_n and $\neg\mathcal{A}_{n+1}$

I now make an important claim about this recursive procedure:

> **Claim.** $\Delta, \Omega_n \nvdash \bot$, for all $1 \leq n \leq m$.

I shall prove this Claim by (ordinary) induction.

> *Base case.* If $\Delta, \mathcal{A}_1 \nvdash \bot$, then clearly $\Delta, \Omega_1 \nvdash \bot$. So suppose instead that $\Delta, \mathcal{A}_1 \vdash \bot$. In this case, Ω_1 is just $\neg\mathcal{A}_1$. Moreover, in this case there is a TFL-proof ending in '$\bot$' with assumptions among $\Delta, \mathcal{A}_1$. Manipulating this TFL-proof and employing the rule $\neg$I, we obtain $\Delta \vdash \neg\mathcal{A}_1$. And, since $\Delta \nvdash \bot$, we obtain $\Delta, \neg\mathcal{A}_1 \nvdash \bot$. So, again, $\Delta, \Omega_1 \nvdash \bot$.
> *Induction case.* Fix n and suppose, for induction, that $\Delta, \Omega_n \nvdash \bot$. Now suppose that $\Delta, \Omega_n, \mathcal{A}_{n+1} \vdash \bot$. Then, as before we can show that $\Delta, \Omega_n, \neg\mathcal{A}_{n+1} \nvdash \bot$.

It follows from the Claim that $\Omega_m \nvdash \bot$. Consequently, there is no sentence such that both it and its negation are among Ω_m. So Ω_m jointly satisfy both (P1) and (P2), and hence Ω_m are jointly polarising.

Now, observe that, for any sentence $\mathcal{B}$ among Δ, we have $\Delta, \Omega_m \vdash \mathcal{B}$ (just using the rule R). Since $\Delta, \Omega_m \nvdash \bot$, it must be that $\Delta, \Omega_m \nvdash \neg\mathcal{B}$. *A fortiori*, $\Omega_m \nvdash \neg\mathcal{B}$. So $\Omega_m \vdash \mathcal{B}$ by Lemma 6.10, and hence $\Omega_m \vDash \mathcal{B}$, by the Soundness Theorem. Generalising: $\Omega_m \vDash \mathcal{B}$, for all $\mathcal{B}$ among Δ.

Finally: by Lemma 6.8, there is a valuation that makes all of Ω_m true. And since $\Omega_m \vDash \mathcal{B}$, for every $\mathcal{B}$ among Δ, this valuation makes all of Δ true. ∎

---

[7]The fact that there are only finitely many such atomic sentences is guaranteed by my assumption (flagged in §6.2, and still in play) that Δ are finite.

This proof of Lemma 6.2 is much more abstract than the proof I offered in §6.2. It provides very little information about the valuation that makes all of Δ true, only that one exists. Of course, at the point where we consider our definition of each Ω_{n+1}, we could use SimpleSearch to *test* for joint-contrariness. But there is no *need* for us to do this. As a result, the algorithm-free proof is crisper and brisker than the proof that proceeded via an algorithm. (Indeed, *despite* all the work in §6.2, I have left an *awful* lot for you to fill in!) So we have a nice and brisk proof of Lemma 6.2. And, of course, since we have a proof of that, we have a proof of the Completeness Theorem itself.

6.4   The superior generality of the abstract proof

The more abstract proof has a further advantage over its algorithmic rival: it is a bit more *general*. I shall explain why in this section, but I should caution that this material is for enthusiasts only.

In §§6.2–6.3, I assumed that we were only given *finitely* many sentences. But what happens when we relax this assumption? Before attempting to answer that question, we should take a little time to make sure we know what it would *mean* to say that $\Gamma \vDash C$, or that $\Gamma \vdash C$, when Γ are *infinite*.

To say that $\Gamma \vDash C$ is to say that there is no valuation that makes all the sentences among Γ true whilst making C false. Thus far, I have encouraged you to think about valuations in terms of truth tables. If Γ are infinite, this will probably be inappropriate, since that would require you to believe in an *infinitely large* truth table. Nonetheless, the use of truth tables is only an heuristic. And when Γ are infinite, it still makes sense to ask whether there is a valuation which makes all of them true (and also C false).

It might, however, be somewhat mind-boggling to consider a claim that $\Gamma \vdash C$ when Γ are *infinite*. After all, every TFL-proof is finitely long: we number each of the lines (with natural numbers) and every proof has a final line. And we could hardly start a TFL-proof by writing down *infinitely many assumptions*, and then drawing inferences from them, since we would never even get to the end of the process of writing down the assumptions!

If we want to make sense of the idea that one might be able to prove something from *infinitely* many assumptions, then we need to conceive of $\Gamma \vdash C$ as telling us that there is a TFL-proof, all of whose initial assumptions are *among* Γ, and which ends with C on no further assumptions. Otherwise put: to say that $\Gamma \vdash C$ is to say that there is some a finitely long TFL-proof which starts with only finitely many of the assumptions among Γ, and which ends with C on no further assumptions.

Understanding '⊢' in this way does not affect the proof of the Soundness Theorem at all (do spend a few minutes confirming this!). However, *both* of our proofs of the Completeness Theorem depended upon the assumption that we were working with only finitely many sentences.

This is clear in the case of the algorithmic proof offered in §6.2. When we are concerned with Completeness, we are asking whether it follows from the fact that $\Gamma \vDash C$, that $\Gamma \vdash C$. If we want to deploy SimpleSearch to address this, then we shall need to input Γ together with $\neg C$ into SimpleSearch. But if Γ are infinite, SimpleSearch will not terminate. Indeed, we will not even get past Step 1 of SimpleSearch, since we will never finish writing *all* of Γ and $\neg C$ as assumptions!

A similar problem arises for the more abstract proof of Lemma 6.2 offered in §6.3. In the course of that proof, I assumed that only finitely many atomic sentences occur as subsentences in the sentences among Δ, namely, $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_m$, and I used this assumption to select my jointly polarising sentences, Ω_m. But if Δ are infinite, it might be that infinitely many atomic sentences occur as subsentences among Δ, and my proof will break down.

Fortunately, this last problem can be repaired. In making our repair, the crucial insight is that every TFL-proof is only ever finitely long. This insight will allow us to prove Lemma 6.2 in full generality.

Proof of Lemma 6.2, in the fully general case. Suppose $\Delta \nvdash \bot$; let $\mathcal{A}_1, \mathcal{A}_2, \ldots$, be all the atomic sentences that occur as subsentences in any of the sentences among Δ. We define:

> *Base clause.* If $\Delta, \mathcal{A}_1 \nvdash \bot$, then let Ω_1 be $\mathcal{A}_1$. Otherwise, let Ω_1 be $\neg\mathcal{A}_1$
> *Recursion clause.* For each number $n > 1$:
> - If $\Delta, \Omega_n, \mathcal{A}_{n+1} \nvdash \bot$, then let Ω_{n+1} be all of Ω_n and $\mathcal{A}_{n+1}$
> - Otherwise, let Ω_{n+1} be all of Ω_n and $\neg\mathcal{A}_{n+1}$

Just as before, we can prove by induction that:

> **Claim 1.** $\Delta, \Omega_n \nvdash \bot$, for all n.

Now let Ω be all and only the sentences that are to be found among some (any) Ω_n. I claim:

> **Claim 2.** $\Delta, \Omega \nvdash \bot$

To establish this claim, suppose that $\Delta, \Omega \vdash \bot$. Since every TFL-proof is only finitely long, there must be finitely many sentences, $\pm\mathcal{A}_{i_1}, \pm\mathcal{A}_{i_2}, \ldots, \pm\mathcal{A}_{i_k}$, among Ω such that:

$$\Delta, \pm\mathcal{A}_{i_1}, \pm\mathcal{A}_{i_2}, \ldots, \pm\mathcal{A}_{i_k} \vdash \bot$$

But then all of *these* sentences, $\pm\mathcal{A}_{i_1}, \pm\mathcal{A}_{i_2}, \ldots, \pm\mathcal{A}_{i_k}$ are among some Ω_n. (Specifically, they are in Ω_n, when n is the largest number among $i_1, \ldots, i_k$.) So $\Delta, \Omega_n \vdash \bot$, which contradicts Claim 1.

It follows from Claim 2 that Ω are jointly polarising. The remainder of the proof now proceeds exactly as before. ∎

In short, pursuing our more abstract proof-strategy allows us to offer a 'completely general' proof of the Completeness Theorem. What is more, we can obtain an important result as an immediate consequence.

Theorem 6.12. Compactness Theorem. Let Δ be any sentences. Suppose that, whenever you take only finitely many sentences of Δ, those finitely many sentences are jointly consistent. Then Δ themselves are jointly consistent.

Proof. We prove the contrapositive. Suppose that Δ are jointly inconsistent. By the (completely general) Completeness Theorem, $\Delta \vdash \bot$. Since every TFL-proof is only finitely long, there are finitely many sentences, Δ_0, among Δ such that $\Delta_0 \vdash \bot$. By the Soundness Theorem, Δ_0 are jointly inconsistent. ∎

Practice exercises

A. Justify the following principles, which are appealed to in the proof of Lemma 6.10 and in the more abstract proof of Lemma 6.2.

1. If $\Gamma \vdash C$ then $\Gamma \vdash \neg\neg C$
2. If $\Gamma \nvdash C \land D$, then either $\Gamma \nvdash C$ or $\Gamma \nvdash D$
3. If $\Gamma \vdash \neg C$ or $\Gamma \vdash \neg D$, then $\Gamma \vdash \neg(C \land D)$
4. If $\Gamma \vdash \neg C$ and $\Gamma \vdash \neg D$, then $\Gamma \vdash \neg(C \lor D)$.
5. If $\Gamma \nvdash \bot$, then either $\Gamma, \mathcal{A} \nvdash \bot$ or $\Gamma, \neg\mathcal{A} \nvdash \bot$.
6. If $\Gamma \vdash \mathcal{A}$ and $\Gamma \nvdash \bot$ then $\Gamma, \mathcal{A} \nvdash \bot$.
7. If $\Gamma, \mathcal{A} \nvdash \bot$, then $\Gamma \nvdash \neg\mathcal{A}$
8. If $\Gamma, \neg\mathcal{A} \vdash \bot$, then $\Gamma \vdash \mathcal{A}$

B. Work through Cases 5 and 6 of the proof of Lemma 6.10.

Tim Button is a Senior Lecturer, and Fellow of St John's College, at the University of Cambridge. He mostly works on meta(meta)physics, logic, mathematics, and language. More at www.nottub.com.