3 First-order logic

Now let's get down to business! This chapter starts with an overview of classical first-order logic (FOL) – meaning standard propositional and predicate logic – which is at the core of any mathematical logic course.

At this level, the most obvious difference between various treatments of FOL is in the choice of proof-system: so I will very briefly comment on two options. Then I highlight the main self-study recommendations. These are followed by some suggestions for parallel and further reading. After a short historical section, this chapter ends with a postscript commenting on some other books, mostly responding to frequently asked questions.[1]

3.1 FOL: a general overview

FOL deals with deductive reasoning that turns on the use of 'propositional connectives' like *and*, *or*, *if*, *not*, and on the use of 'quantifiers' like *every*, *some*, *no*. But in ordinary language (and even in informal mathematics) these logical operators work in surprisingly complex ways, introducing the kind of obscurities and possible ambiguities we certainly want to avoid in logically transparent arguments. What to do?

From the time of Aristotle, logicians have used a 'divide and conquer' strategy that involves introducing restricted tightly-regimented languages. For Aristotle, his regimented language was a fragment of very stilted Greek; for us, our regimented languages are artificial formal constructions. Either way, the plan is that we tackle a stretch of reasoning by first reformulating it in a suitable regimented language with much tidier logical operators, and then we can evaluate the reasoning once recast into this more well-behaved form. This way, we have a division of labour. First, we clarify the intended structure of the original argument by rendering it into an unambiguous regimented/formalized language.

[1] A note to philosophers. If you *have* carefully read a substantial introductory logic text for philosophers such as Nicholas Smith's, or even my own, you will already be familiar with (versions of) a fair amount of the material covered in this chapter. However, you will now begin to see topics being re-presented in the sort of mathematical style and with the sort of rigorous detail that you will necessarily encounter more and more as you progress in logic. You do need to start feeling entirely comfortable with this mode of presentation at an early stage. So it is well worth working through even rather familiar topics again, this time with more mathematical precision.

Second, there's the separate business of assessing the validity of the resulting regimented argument.

In FOL, then, we use appropriate formal languages which contain, in particular, tidily-disciplined surrogates for the propositional connectives *and*, *or*, *if*, *not* (standardly symbolized $\wedge$, $\vee$, $\rightarrow$ and $\neg$), plus replacements for the ordinary language quantifiers (roughly, using $\forall\mathsf{x}$ for *every x is such that* ..., and $\exists\mathsf{y}$ for *some y is such that*...).

Although the fun really starts once we have the quantifiers in play, it is very helpful to develop FOL in two main stages:

(a) Typically, we start by introducing propositional languages whose built-in logical apparatus comprises just the propositional connectives, and discuss the propositional logic of arguments framed in these languages. This gives us a very manageable setting in which to first encounter a whole range of logical concepts and strategies.

(b) We then move on to develop the syntax and semantics of richer formal languages which add the apparatus of so-called first-order quantification, and explore the logic of arguments rendered into such languages.

So let's have some more detail about stages (a) and (b).

(a.i) We first look at the *syntax* of propositional languages, defining what count as the well-formed formulas of such languages. If you have already encountered languages of this kind, you will now get to know how to actually prove various things about them that seem obvious and that you perhaps previously took for granted – for example, that 'bracketing works' to avoid ambiguities, so every well-formed formula has a unique unambiguous parsing.

(a.ii) On the *semantic* side, we need the idea of a *valuation* for a propositional language. We start from an assignment of truth-values, *true* vs *false*, to the atomic formulas (the basic components of our languages). This assignment of values to atoms then fixes the truth-values of complex sentences involving the connectives. Here we depend on the 'truth-functional' interpretation of the connectives; and we can use truth-tables to display how the truth-values of complex formulas are fully determined by the truth-values of their atomic parts.

So any formula, atomic or complex, is treated as being either definitely true or definitely false – one or the other, but not both – on any particular valuation: this is the core assumption distinctive of *classical semantics*.

We can then define the key semantic relation of *(tautological) entailment*, where a set of sentences Γ from a propositional language (tautologically) entails the sentence φ when any valuation of the relevant atoms which makes all the sentences in Γ true makes φ true too. We will explore some of the key properties of this semantic entailment relation, and learn how to calculate whether the relation holds.

(a.iii) Different textbook presentations of stages (a.i) and (a.ii) tend to be very similar; but now the path forks. For the next topic will be a *formal deductive* *system* in which we can construct step-by-step derivations of conclusions from premisses in propositional logic. There is a variety of such systems to choose from, and I'll mention half-a-dozen in §3.2. Different proof systems for classical propositional logic will (as you'd expect) be equivalent – meaning that, given some premisses, we can derive the same conclusions in each system. However, the systems do differ considerably in their intuitive appeal and user-friendliness, as well as in some of their more technical features. Note, though: apart from looking at a few illustrative examples, we won't be much interested in producing lots of derivations *inside* a chosen proof system; the focus will be more on establishing results *about* the systems.

In due course, the educated logician will want to learn at least a little about the various styles of proof system – at the minimum, you should eventually get a sense of how they respectively work, and come to appreciate the interrelations between them. But here – as is usual when starting out at this level – we look at *axiomatic* and *natural deduction* systems in particular (I say more about these in the next section).

(a.iv) At this point, then, we will have two quite different ways of defining what makes for a deductively good argument in propositional logic:

> We said that a set of premisses Γ *tautologically entails the conclusion* φ if every possible valuation which makes Γ all true makes φ true (that's a semantically defined idea).
>
> We can now also say that Γ *yields the conclusion* φ *in your chosen proof-system* S if there is an S-type derivation of the conclusion φ from premisses in Γ.

Of course, we want these two approaches to fit together. We want our favoured proof-system S to be *sound* – it shouldn't give false positives. In other words, if there is an S-derivation of φ from Γ, then φ really *is* tautologically entailed by Γ. We also want our favoured proof-system S to be *complete* – we want it to capture all the correct semantic entailment claims. In other words, if φ is tautologically entailed by the set of premisses Γ, then there is indeed some S-derivation of φ from premisses in Γ.

So we will want to establish both the soundness and the completeness of our favoured proof-system S for propositional logic (axiomatic, natural deduction, whatever). Now, these two results will hold no terrors! However, in establishing soundness and completeness for propositional logics you will encounter some useful strategies which can later be beefed-up to give us soundness and completeness results for stronger logics.

(b.i) Having warmed up with propositional logic, we turn to full FOL so we can also deal with arguments whose validity depends on their quantificational structure (starting with the likes of our old friend 'Socrates is a man; all men are mortal; hence Socrates is a mortal'!).

We need to introduce appropriate formal languages with quantifiers (more precisely, with first-order quantifiers: the next chapter explains the contrast with second-order quantifiers). So *syntax* first. And while the syntax of propositional logic is quite straightforward, a story needs to be told about why standard formal expressions of generality are structured as they are.

Consider the simple ordinary-language sentence (i) 'Socrates is wise'. And now note that we can replace the name in (i) with the quantifier expression 'everyone', giving us another sentence (ii) 'Everyone is wise'. Similarly, we can replace the name 'Juliet' in (iii) 'Romeo loves Juliet' with the quantifier expression 'someone' to get the equally grammatical (iv) 'Romeo loves someone'. In FOL, however, while we might render (i) as simply Ws, (ii) will get rendered by something like $\forall\mathsf{xWx}$ (roughly, everyone x is such that x is wise). Similarly if (iii) is rendered Lrj, then (iv) gets rendered by something like $\exists\mathsf{xLrx}$ (roughly, someone x is such that Romeo loves x). But why? It helps to understand the rationale for this departure from the syntactic patterns of ordinary language and the use of the apparently more complex 'quantifier/variable' device for expressing generalizations.

(b.ii) Turning to *semantics*: the first key idea we need is that of a *structure*, a (non-empty) domain of objects equipped with some properties, relations and/or functions. Here, we treat properties etc. extensionally. In other words, we can think of a property as a set of objects from the domain, a binary relation as a set of pairs from the domain, and so on. Compare our remarks on naive set theory in §2.1; though, heeding the point of §2.4, we can arguably take the talk of sets here in a non-committal way.

Then, crucially, you need to grasp the idea of an *interpretation* of a language in such a structure. Names are interpreted as denoting objects in the domain; a one-place predicate gets assigned a property, i.e. a set of objects from the domain (its extension – intuitively, the objects it is true of); a two-place predicate gets assigned a binary relation; and so on.

Such an interpretation of the elements of a first-order language then generates a valuation (a unique assignment of truth-values) for every sentence of the interpreted language. How does it do that? Well, a predicate-name sentence like Ws will be true just if the object denoted by s is in the extension of W; a sentence like Lrj is true if the ordered pair of the objects denoted by r and j is in the extension of L; and so on. That's easy, and the propositional connectives behave as in propositional logic. But adding to the formal semantic story to explain how the interpretation of a language fixes the valuations of the quantified sentences requires important new ideas. (There are in fact a number of slightly different though ultimately equivalent ways of spinning that last stage of the story: you need to get your head round one of them.)

We can now introduce the idea of a *model* for a set of sentences, i.e. an interpretation in a structure which makes all the sentences true together. And we can then again define a semantic relation of *entailment*, this time for FOL sentences. A set of sentences Γ semantically entails φ when any interpretation in any structure which makes all the sentences in Γ true also makes φ true too – for short, when any model for Γ is a model for φ. You'll again need to know some of the basic properties of this entailment relation. (For example, note that if Γ has no model, then – on this definition – Γ semantically entails φ for any φ at all, including any contradiction.)

(b.iii) We next need to explore a proof system for FOL. And you can again encounter six main types of proof system, with their varying attractions and drawbacks. To repeat, you'll want at some future point to find out at least something about all these styles of proof. But we will principally be looking here at axiomatic systems and at one kind of natural deduction.

(b.iv) As with propositional logic, we will want to show that our chosen proof system for FOL is sound and doesn't overshoot (giving us false positives) and is complete and doesn't undershoot (leaving us unable to derive some semantically valid entailments). In other words, if S is our FOL proof system, Γ a set of sentences, and φ a particular sentence, we need to show:

> If there is an S-proof of φ from premisses in Γ, then Γ does indeed semantically entail φ. (Soundness)
>
> If Γ semantically entails φ, then there is an S-proof of φ from premisses in Γ. (Completeness)

The completeness theorem, by the way, comes in two versions, a weaker version where Γ is restricted to having only finitely many members, and a crucial stronger version which allows Γ to be infinite. And it is at *this* point, it might well be said, that the study of FOL becomes really interesting: in particular, establishing strong completeness involves rather more sophisticated ideas than anything we have met before.

(b.v) Later chapters will continue the story along various paths; here it is worth quickly mentioning just one immediate corollary of completeness.

Proofs in formal systems are finitely long; so a proof of φ from Γ can only call on a *finite* number of premisses in Γ. But the (strong) completeness theorem for FOL allows Γ to have an *infinite* number of members. This combination of facts immediately implies the *compactness* theorem for sentences of FOL languages:

(c) If every finite subset of Γ has a model, so does Γ.[2]

This compactness theorem has numerous applications in model theory.

3.2 A little more about types proof-system

(a) You will be familiar with the informal idea of an axiomatized theory. We are given some axioms and some deductive apparatus is presupposed. Then the theorems of the theory are whatever can be derived from the axioms. Similarly:

[2] That's equivalent to the claim that if (i) Γ *doesn't have a model*, then there is a finite subset $\Delta \subseteq \Gamma$ such that (ii) Δ *has no model*. Suppose (i). That implies that Γ semantically entails a contradiction. So by completeness we can derive a contradiction from Γ in your favourite proof system. That proof will only use a finite collection of premisses $\Delta \subseteq \Gamma$. But if Δ proves a contradiction, then by soundness, Δ semantically entails a contradiction, which can only be the case if (ii).

> In an *axiomatic* logical system, we adopt some basic logical truths as axioms, and then explicitly specify the allowed rules of inference, usually just very simple ones such as the 'modus ponens' rule for the conditional: 'from A and $A \to C$, we can infer C'.
>
> A proof from some given premisses to a conclusion then has the simplest possible structure. It is just a sequence of sentences – each one is either (i) one of the premisses, or (ii) one of the logical axioms, or (iii) follows from earlier sentences in the proof by one of the rules of inference – with the whole sequence ending with the target conclusion.
>
> A logical theorem of the system is a sentence that can be proved from the logical axioms alone (without appeal to any further premisses).

Informal deductive reasoning, however, is not relentlessly linear like this. We do not require that each proposition in a proof (other than a given premiss or a logical axiom) has to follow from what's gone before. Rather, we often step sideways (so to speak) to make some new temporary assumption, 'for the sake of argument'. For example, we may say 'Now suppose that A is true'; we go on to show that, given what we've already established, this extra supposition leads to a contradiction; we then drop or 'discharge' the temporary supposition and conclude that *not*-A. That's how a reductio ad absurdum argument works. For another example, we may again say 'Suppose that A is true'; this time we go on to show that we can now derive C; we then discharge the temporary supposition and conclude that *if* A*, then* C). That's how we often argue for a conditional proposition.

Noting this,

> A *natural deduction* system of logic aims to formalize patterns of reasoning now including those where we argue by making and then later discharging temporary assumptions. So, for example, as well as the simple modus ponens rule for the conditional '$\to$', there will be a rule along the lines of 'if we can infer C from the assumption A, we can drop the assumption A and conclude $A \to C$'.

Evidently, in a natural deduction system, we will need some way of keeping track of which temporary assumptions are in play and for how long. Two particular ways of doing this are popular.

> Gerhard Gentzen in his doctoral thesis of 1933 sets proofs out as trees – with premisses or temporary assumptions at the top of branches, the conclusion at the root of the tree – and uses a system for explicitly tagging temporary assumptions and the inference moves where they get dropped.
>
> Later, a multi-column layout was popularized by Frederick Fitch in his classic 1952 logic text, *Symbolic Logic: an Introduction*. This time,

> proofs snake from column to column, moving a column to the right when a new temporary assumption is made, and moving back a column to the left when the assumption is dropped again.

Either way, a theorem of the natural deduction system will be a proposition which is provable without appeal to any premisses (all assumptions made en route are discharged).

(b) That introduces us – in an initial arm-waving way! – to the first three of the commonly used kinds of proof-system:

1. Old-school axiomatic systems.
2. Natural deduction done Gentzen-style.
3. Natural deduction done Fitch-style.
4. 'Truth trees' or 'semantic tableaux'.
5. Sequent calculi.
6. Resolution calculi.

I won't give any details here about the last three. But the easiest kinds of system for non-mathematical beginners to work with are (3) Fitch-style deductions and (4) truth trees – which is why the majority of elementary logic books for philosophers introduce one or other (or both) systems. By contrast, (5) the sequent calculus (in its most interesting form) really comes into its own in more advanced work in proof theory, while (6) resolution calculi are perhaps of particular concern to computer scientists interested in automating theorem proving. Introductory mathematical logic text books, however, usually focus on either (1) axiomatic or (2) Gentzen-style proof systems.

True, axiomatic systems in their raw state can be pretty horrible to use – but they can be made a bit less painful once you learn some basic dodges (like the use of the so-called 'Deduction Theorem'). By comparison, Gentzen-style systems are much more attractive – there is a reason they are called *natural* deduction systems!

I should add, though, that even once you've picked your favoured general *type* of proof-system to work with from (1) to (6), there are many more choices to be made before landing on a specific system of that type. For example, F. J. Pelletier and Allen Hazen published a survey of logic texts aimed at philosophers which use natural deduction systems (tinyurl.com/pellhazen): they note that no less than thirty texts use a variety of Fitch-style system – and rather remarkably *no two of these have exactly the same system of rules for FOL*! Moral? Don't get too hung up on the finest details of a particular textbook's proof-system; it is the guiding ideas that matter, together with the proofs *about* the chosen proof-system (such as the soundness and completeness theorems).

3.3 Main recommendations for reading on FOL

A preliminary reference. In my elementary logic book I do carefully explain the 'design brief' for the languages of FOL, spelling out the rationale for the quantifier-variable notation. For some, this might be helpful parallel reading when working through your chosen main text(s), at the point when that notation is introduced:

1. Peter Smith, *Introduction to Formal Logic*** (2nd edn), Chapters 26–28. Downloadable from logicmatters.net/ifl.

Unsurprisingly, there is a *very* long list of texts which cover FOL. But the whole point of this Guide is to choose. So here are my top recommendations, starting with one-and-a-third books which, taken together, make an excellent introduction:

2. Ian Chiswell and Wilfrid Hodges, *Mathematical Logic* (OUP 2007). This nicely written text is very approachable. It is written by mathematicians primarily for mathematicians. However, it is only one notch up in actual difficulty from some introductory texts for philosophers like mine or Nick Smith's, though – as its title might suggest – it does have a notably more mathematical 'look and feel'. Philosophers can skip over a few of the more mathematical illustrations; while depending on background, mathematicians should be able to take this book at pace.

   The briefest headline news is that authors explore a Gentzen-style natural deduction system. But by building things up in three stages – so after propositional logic, they consider an important fragment of first-order logic before turning to the full-strength version – they make e.g. proofs of the completeness theorem for first-order logic unusually comprehensible. For a more detailed description see my book note on C&H, tinyurl.com/CHbooknote.

   Very warmly recommended, then. For the moment, you only *need* read up to and including §7.6 (under two hundred pages). But having got that far, you might as well read the final few sections and the Postlude too! The book has brisk solutions to some of the exercises.

Next, you should complement C&H by reading the first third of the following excellent book:

3. Christopher Leary and Lars Kristiansen's *A Friendly Introduction to Mathematical Logic*** (1st edn by Leary alone, Prentice Hall 2000; 2nd edn Milne Library 2015). Downloadable at tinyurl.com/friendlylogic.

   There is a great deal to like about this book. Chs. 1–3, in either edition, do indeed make a friendly and helpful introduction to FOL. The authors use an axiomatic system, though this is done in a particularly smooth way. At this stage you could stop reading after the beginning of §3.3 on compactness, which means you will be reading just 87 pages.

   Unusually, L&K dive straight into a treatment of first-order logic without spending an introductory chapter or two on propositional logic: in a sense, as you will see, they let propositional logic look after itself. But this happily means (in the present context) that you won't feel that you are labouring through the very beginnings of logic one more time than is really necessary – this book therefore dovetails very nicely with C&H.

Again written by mathematicians, some illustrations of ideas can presuppose a smattering of background mathematical knowledge; but philosophers will miss very little if they occasionally have to skip an example (and the curious can always resort to Wikipedia, which is quite reliable in this area, for explanations of some mathematical terms). The book ends with extensive answers to exercises.

I like the overall tone of L&K very much indeed, and say more about this admirable book in another book note, tinyurl.com/LKbooknote.

As an alternative to the C&H/L&K pairing, the following slightly more conventional book is also exceptionally approachable:

4. Derek Goldrei's *Propositional and Predicate Calculus: A Model of Argument* (Springer, 2005) is explicitly designed for self-study. Read up to the end of §6.1 (though you could skip §§4.4 and 4.5 for now, leaving them until you turn to elementary model theory).

   While C&H and the first third of L&K together cover overlapping material twice, Goldrei – in a comparable number of pages – covers very similar ground once, concentrating on a standard axiomatic proof system. So this is a relatively gently-paced book, allowing Goldrei to be more expansive about fundamentals, and to give a lot of examples and exercises with worked answers to test comprehension along the way. A great amount of thought has gone into making this text as clear and helpful as possible. Some may find it occasionally goes a bit too slowly, though I'd say that this is erring on the right side in an introductory book for self-teaching: if you want a comfortingly manageable text, you should find this particularly accessible. As with C&H and L&K, I like Goldrei's tone and approach a great deal.

   But since Goldrei uses an axiomatic system throughout, do eventually supplement with at least a brief glance at a Gentzen-style natural deduction proof system.

These three main recommended books, by the way, have all had very positive reports over the years from student users.

3.4 Some parallel and slightly more advanced reading

The material covered in the last section is so very fundamental, and the alternative options so very many, that I really do need to say at least something about a few other books. So in this section I list – in rough order of difficulty/sophistication – a small handful of further texts which could well make for useful additional or alternative reading. Then in the final section of the chapter, I will mention some other books I've been asked about.

I'll begin a notch or two down in level from the texts we have looked at so far, with a book written by a philosopher for philosophers. It should be particularly accessible to non-mathematicians who haven't done much formal logic before, and should help ease the transition to coping with the more mathematical style of the books recommended in the last section.

5. David Bostock, *Intermediate Logic* (OUP 1997). From the preface: "The book is confined to . . . what is called first-order predicate logic, but it aims to treat this subject in very much more detail than a standard introductory text. In particular, whereas an introductory text will pursue just one style of semantics, just one method of proof, and so on, this book aims to create a wider and a deeper understanding by showing how several alternative approaches are possible, and by introducing comparisons between them." So Bostock ranges more widely than the books I've so far mentioned; he does indeed usefully introduce you to tableaux ('truth trees') *and* an Hilbert-style axiomatic proof system *and* natural deduction *and* even a sequent calculus as well. Indeed, anyone could profit from at least a quick browse of his Part II to pick up the headline news about the various approaches.

   Bostock eventually touches on issues of philosophical interest such as free logic which are not often dealt with in other books at this level. Still, the discussions mostly remain at much the same level of conceptual/mathematical difficulty as e.g. my own introductory book. He proves completeness for tableaux in particular, which I always think makes the needed construction seem particularly natural.

   Intermediate Logic should indeed be particularly helpful for those feeling the need for a book that bridges a felt gap between an introductory logic course for philosophers and full-on mathematical logic texts.

As noted, unlike our main recommendations, Bostock does discusses tableaux ('truth trees'). If you are a philosopher, you may well have already encountered them in your introductory logic course. If not, as an alternative to Bostock,

6. My elementary introduction to truth trees for propositional logic available at tinyurl.com/proptruthtrees will quickly give you the basic idea in an accessible way.

Next, back to the level we want, even though it is giving a second bite to an author we've already met, I must mention a rather different discussion of FOL:

7. Wilfrid Hodges, 'Elementary Predicate Logic', in the *Handbook of Philosophical Logic*, Vol. 1, ed. by D. Gabbay and F. Guenthner, (Kluwer 2nd edition 2001). This is a slightly expanded version of the essay in the first edition of the *Handbook* (read that earlier version if this one isn't available), and is written with Hodges's usual enviable clarity and

verve. As befits an essay aimed at philosophically minded logicians, it is full of conceptual insights, historical asides, comparisons of different ways of doing things, etc., so it very nicely complements the textbook presentations of C&H, L&K and/or Goldrei.

Read at this stage the very illuminating first twenty short sections.

Next, here's a much-used text which has gone through multiple editions and should be in any library; it is a very useful natural-deduction based alternative to C&H. Later chapters of this book are also mentioned later in this Guide as possible reading on further topics, so it could be worth making early acquaintance with

8. Dirk van Dalen, *Logic and Structure* (Springer, 1980; 5th edition 2012). The early chapters up to and including §3.2 provide an introduction to FOL via Gentzen-style natural deduction. The treatment *is* often approachable and written with a relatively light touch. However, it has to be said that the book isn't without its quirks and flaws and inconsistencies of presentation (though perhaps you have to be an alert and rather pernickety reader to notice and be bothered by them). Still, having said that, the coverage and general approach is good.

   Mathematicians should be able to cope readily. I suspect, however, that the book would occasionally be tougher going for philosophers if taken from a standing start – which is another reason why I have recommended beginning with C&H instead. For more on this book, see tinyurl.com/dalenlogic.

As a follow up to C&H, I just recommended L&K's *Friendly Introduction* which uses an axiomatic system. As an alternative to that, here is an older (and, in its day, much-used) text which should also be very widely available:

9. Herbert Enderton, *A Mathematical Introduction to Logic* (Academic Press 1972, 2002). This also focuses on an axiomatic system, and is often regarded as a classic of exposition. However, it does strike me as somewhat less approachable than L&K, so I'm not surprised that students do quite often report finding this book a bit challenging *if used by itself as a first text*.

   However, this is an admirable and very reliable piece of work which most readers should be able to cope with well if used as a supplementary second text, e.g. after you have tackled C&H. And stronger mathematicians might well dive into this as their first preference.

   Read up to and including §2.5 or §2.6 at this stage. Later, you can finish the rest of that chapter to take you a bit further into model theory. For more about this classic, see tinyurl.com/enderlogicnote.

I should also certainly mention the outputs from the Open Logic Project. This is an entirely admirable, collaborative, open-source, enterprise inaugurated by Richard Zach, and continues to be work in progress. You can freely download the latest full version and various sampled ‘remixes’ from tinyurl.com/openlogic. In an earlier version of this Guide, I said that “although this is referred to as a textbook, it is perhaps better regarded as a set of souped-up lecture notes, written at various degrees of sophistication and with various degrees of more book-like elaboration.” But things have moved on: the mix of chapters on propositional and quantificational logic in the following has been expanded and developed considerably, and the result is much more book-like:

10. Richard Zach and others, *Sets, Logic, Computation*** (Open Logic). There’s a lot to like here (Chapters 5 to 13 are the immediately relevant ones for the moment). In particular, Chapter 9 could make for very useful supplementary reading on natural deduction. Chapter 8 tells you about a sequent calculus (a slightly odd ordering!). And Chapter 10 on the completeness theorem for FOL should also prove a very useful revision guide.

    My sense is that overall these discussions probably will still go somewhat too briskly for some readers to work as a stand-alone introduction for initial self-study without the benefit of lecture support, which is why this doesn’t feature as one of my principal recommendations in the previous section: however, your mileage may vary. And certainly, chapters from this project could/should be *very* useful for reinforcing what you have learnt elsewhere. You can download *SLC* from tinyurl.com/slcopen.

So much, then, for reading on FOL running on more or less parallel tracks to the main recommendations in the preceding section. I’ll finish this section by recommending two books that push the story on a little. First, an absolute classic, short but packed with good things:

11. Raymond Smullyan, *First-Order Logic** (Springer 1968, Dover Publications 1995). This is terse, but those with a taste for mathematical elegance can certainly try its Parts I and II, just a hundred pages, after the initial recommended reading in the previous section. This beautiful little book is the source and inspiration of many modern treatments of logic based on tree/tableau systems. Not always easy, especially as the book progresses, but a real delight for the mathematically minded.

And second, taking things in a new direction, don’t be put off by the title of

12. Melvin Fitting, *First-Order Logic and Automated Theorem Proving* (Springer, 1990, 2nd end. 1996). This is a wonderfully lucid book by a renowned expositor. Yes, at a number of places in the book there are illustrations of how to implement algorithms in Prolog. But either you can easily pick up the very small amount of background knowledge that’s needed to follow everything that is going on (and that’s quite fun) or you can in fact just skip lightly over those implementation episodes while still getting the principal logical content of the book.

As anyone who has tried to work inside an axiomatic system knows, proof-discovery for such systems is often hard. Which axiom schema should we instantiate with which wffs at any given stage of a proof? Natural deduction systems are nicer. But since we can, in effect, make any new temporary assumption we like at any stage in a proof, again we still need to keep our wits about us if we are to avoid going off on useless diversions. By contrast, tableau proofs (a.k.a. tree proofs) can pretty much write themselves even for quite complex FOL arguments, which is why I used to introduce formal proofs to students that way (in teaching tableaux, we can largely separate the business of getting across the idea of formality from the task of teaching heuristics of proof-discovery). And because tableau proofs very often write themselves, they are also good for automated theorem proving. Fitting explores both the tableau method and the related so-called resolution method which we mentioned as, yes, a sixth style of proof!

This book's approach is, then, rather different from most of the other recommended books. However, I do think that the fresh light thrown on first-order logic makes the slight detour through this extremely clearly written book *vaut le voyage*, as the Michelin guides say. (If you don't want to take the full tour, however, there's a nice introduction to proofs by resolution in Shawn Hedman, *A First Course in Logic* (OUP 2004): §1.8, §§3.4–3.5.)

3.5 A little history (and some philosophy too)

(a) Classical FOL is a powerful and beautiful theory. Its treatment, in one version or another, is always the first and most basic component of modern textbooks or lecture courses in mathematical logic. But how did it get this status?

The first system of formalized logic of anything like the contemporary kind – Frege's system in his *Begriffsschrift* of 1879 – allows higher-order quantification in the sense explained in the next chapter (and Frege doesn't identity FOL as a subsystem of distinctive interest). The same is true of Russell and Whitehead's logic in their *Principia Mathematica* of 1910–1913. It is not until Hilbert and Ackermann in their rather stunning short book *Mathematical Logic* (original German edition 1928) that FOL is highlighted under the label 'the restricted predicate calculus'. Those three books all give axiomatic presentations of logic (though notationally very different from each other): axiomatic systems similar enough to the third are still often called 'Hilbert-style systems'

(b) As an aside, it is worth noting that the axiomatic approach reflects a broadly shared philosophical stance on the very nature of logic. Thus Frege thinks of logic as a science, in the sense of a body of truths governing a special subject matter (they are fundamental truths governing logical operations such as negation, conditionalization, quantification, identity). And in *Begriffsschrift* §13, he extols the general procedure of axiomatizing a science to reveal how a bunch of laws hang together: 'we obtain a small number of laws [the axioms] in which . . . is included, though in embryonic form, the content of all of them'. So it is not surprising that Frege takes it as appropriate to present logic axiomatically too.

In a rather different way, Russell also thought of logic as a science; he thought of it as in the business of systematizing the most general truths about the world. A special science like chemistry tells us truths about certain kinds of constituents of the world and certain of their properties; for Russell, logic tells us absolutely general truths about *everything*. If you think like *that*, treating logic as (so to speak) the most general science, then of course you'll again be inclined to regiment logic as you do other scientific theories, ideally by laying down a few 'basic laws' and then showing that other general truths follow.

Famously, Wittgenstein in the *Tractatus* reacted radically against Russell's conception of logic. For him, logical truths are *tautologies*. They are not deep ultimate truths about the most general, logical, structure of the universe; rather they are *empty* claims in the sense that they tell us nothing informative about how the world is: they merely fall out as byproducts of the meanings of the basic logical particles.

That last idea can be developed in more than one way. But one approach is Gentzen's in the 1930s. He thought of the logical connectives as getting their meanings from how they are used in inference (so grasping their meaning involves grasping the inference rules governing their use). For example, grasping 'and' involves grasping, inter alia, that (i) from A *and* B you can (of course!) derive A. Similarly, grasping the conditional involves grasping, inter alia, that (ii) a derivation of the conclusion C from the temporary supposition A warrants an assertion of *if* A *then* C. But now consider this little two-step derivation:

> Suppose for the sake of argument that A *and* B; then we can derive A – by the rule (i) which partly fixes the meaning of 'and'.
>
> And given that little suppositional inference, the rule (ii) which partly gives the meaning of 'if' entitles us to drop the supposition and conclude *if* A *and* B*, then* A.

In other words, the inference rules (i) and (ii) enable us to derive that logical truth 'for free' (from no remaining assumptions): it's a theorem of a formal system with those rules.

If this is right, and if the point generalizes, then we don't have to see such logical truths as reflecting deep facts about the logical structure of the world (whatever that could mean): logical truths fall out just as byproducts of the inference rules whose applicability is, in some sense, built into the very meaning of e.g. the connectives and the quantifiers.

It is a nice question how far we should buy that sort of de-mystifying story about the nature of logical truth. But whatever your judgement on that, there surely *is* something odd about thinking with Frege and Russell that a systematized logic is primarily aiming to regiment a special class of ultra-general truths. Isn't logic at bottom about good and bad reasoning practices, about what makes for a good proof? Shouldn't its prime concern be the correct styles of valid inference? So shouldn't a formalized logic highlight *rules of valid proof-building* (perhaps as in a natural deduction system) rather than stressing *logical truths*?

(c) Back to the history of the technical development of logic. An obvious starting place is with the clear and judicious

13. William Ewald, 'The Emergence of First-Order Logic', *Stanford Encyclopaedia*, tinyurl.com/emergenceFOL.

If you want rather more, the following is also readable and very helpful:

14. José Ferreiros, 'The Road to Modern Logic – an Interpretation', *Bulletin of Symbolic Logic* 7 (2001): 441–484, tinyurl.com/roadtologic.

And for a longer, though rather bumpier, read – you'll probably need to skim and skip! – you could also try dipping into this more wide-ranging piece:

15. Paolo Mancosu, Richard Zach and Calixto Badesa, 'The Development of Mathematical Logic from Russell to Tarski: 1900–1935' in Leila Haaparanta, ed., *The History of Modern Logic* (OUP, 2009, pp. 318–471): tinyurl.com/developlogic.

3.6 Postscript: Other treatments?

I will end this chapter by responding to a variety of Frequently Asked Questions raised in response to earlier versions of the Guide (often of the form "But why haven't you recommended X?"). So (a) I look, mostly critically, at some books aimed at philosophers. Then (b) I comment on a few well-known mathematical logic texts. Finally (c) I mention some fun extras.

(a.i) *Designed for philosophers: What about Sider? What about Bell, DeVidi and Solomon?* Theodore Sider has written a text called *Logic for Philosophy** (OUP, 2010) which I've repeatedly been asked to comment on. The book in fact falls into two halves. The second half (about 130 pages) is on modal logic, and I will return to that in Chapter 10. The first half of the book (almost exactly the same length) is on propositional and first-order logic, together with some variant logics, so is very much on the topic of this chapter. But while the coverage of modal logic is quite good, the treatment of FOL isn't. I certainly can't recommend the first half of this book: I explain why in a book note, tinyurl.com/siderbook.

A potential alternative to Bostock at about the same level, and which can initially look promising, is John L. Bell, David DeVidi and Graham Solomon's *Logical Options: An Introduction to Classical and Alternative Logics* (Broadview Press 2001). This book covers a lot pretty snappily – for the moment, just Chapters 1 and 2 are relevant – and some years ago I used it as a text for second-year seminar for undergraduates who had used my own tree-based book for their first year course. But many students found the exposition too terse, and I found myself having to write very extensive seminar notes. If you want a book at this intermediate level, you'd do much better sticking with the more expansive Bostock.

(a.ii) *What about The Logic Book?* Many US philosophers have had to take courses based on *The Logic Book* by Merrie Bergmann, James Moor and Jack Nelson (first published by McGraw Hill in 1980; a sixth edition was published – at a quite ludicrous price – in 2013). I rather doubt that those students over the years have much enjoyed the experience! This is a large book, over 550 pages, starting at about the level of my introductory book, and going as far as metalogical results like a full completeness proof for FOL, so its coverage overlaps considerably with the main recommendations of §3.3. But while reliable enough, it all strikes me, like some other readers who have commented, as *very* dull and laboured, and often rather unnecessarily hard going. You can certainly do better.

(a.iii) *What about an introductory text with a more proof-theoretic slant?* It is true that there are some questions about systems of FOL which can be tackled at a quite introductory level, yet which aren't addressed by any of the readings so far mentioned. For a simple example, suppose we are working from given premisses in a formal proof system and (i) have so far derived A and also derived B; then (ii) we can (rather boringly!) infer the conjunction $A \wedge B$ (remember '$\wedge$' means *and*). Now, suppose later in the same proof (iii) we appeal to that conjunction $A \wedge B$ to derive A. We wouldn't have gone *wrong*; but obviously we have gone on a pointless detour, given that at stage (i) we have already derived A. There is evident interest in the question of how to eliminate such detours and other pointless digressions from proofs. Gentzen famously started the ball rolling in his discussions of how to 'normalize' proofs in one of his natural deduction systems, and he showed how normalization results can be used to derive other important properties of the proof systems.

For a first encounter with this sort of topic, the place to look is Jan von Plato's *Elements of Logical Reasoning** (CUP, 2014). This is based on the author's lectures for a general introductory FOL course. But a lot of material is touched on in a relatively short compass as von Plato talks about a range of different natural deduction and sequent calculi. I suspect that, without prior knowledge of ND and without classroom work to round things out, this book might not be as accessible as the author intends. Still, if you already have encountered ND systems and want to know about variant ways of setting up such systems, about proof-search, about the relation with so-called sequent calculi, etc., then the first half of this book makes a pretty clear start. However, my own recommendation would be to concentrate for the moment on the core topics in FOL covered by the books we have mentioned previously, leaving it to later to dive into proof theory in more detail, as covered in Chapter 9.

(b.i) *A blast from the past: What about Mendelson?* Turning to mainstream mathematical logic books, perhaps the most frequent question I used to get asked in response to early versions of the Guide was 'But what about Mendelson, Chs. 1 and 2'? Well, Elliott Mendelson's *Introduction to Mathematical Logic* (originally van Nostrand 1964; and now Chapman and Hall/CRC, 6th edn. 2015). The book was I think the first modern textbook of its type (so immense credit to Mendelson for that), and I no doubt owe my whole career to it – it got me through tripos!

It seems that some others who learnt using the book are in their turn still using it to teach from. But let's not get too sentimental! It has to be said that the book in its first incarnation was often brisk to the point of unfriendliness, and the basic look-and-feel of the book hasn't changed a great deal as it has run through successive editions. Mendelson's presentation of axiomatic systems of logic are quite tough going, and as the book progresses in later chapters through formal number theory and set theory, things if anything get somewhat less reader-friendly. Which certainly *doesn't* mean the book won't repay working through. But quite unsurprisingly, over fifty years on, there are many rather more accessible and more amiable alternatives for beginning serious logic. Mendelson's book is a landmark well worth visiting one day, but I can't recommend starting there (especially for self-study). For a little more about it, tinyurl.com/mendelsonlogic.

(If you *do* want an old-school introduction from the same era, you might more enjoy Geoffrey Hunter, *Metalogic** (Macmillan 1971, University of California Press 1992). This is not as comprehensive as Mendelson: but it was an exceptionally good textbook from a time when there were few to choose from. Read Parts One to Three at this stage. And if you are finding it rewarding reading, then do eventually finish the book: it goes on to consider formal arithmetic and proves the undecidability of first-order logic, topics we consider in Chapter 6. Unfortunately, the typography – from pre-LaTeX days – isn't at all pretty to look at. But in fact the treatment of an axiomatic system of logic is extremely clear and accessible. It might be worth blowing the dust off your library's copy!)

(b.ii) *Five more recent mathematical logic texts: What about Ebbinghaus, Flum and Thomas? Hedman? Hinman? Rautenberg? Kaye?* We start with H.-D. Ebbinghaus, J. Flum and W. Thomas, *Mathematical Logic* (Springer, 2nd edn. 1994). This is the English translation of a book first published in German in 1978, and appears in a series 'Undergraduate Texts in Mathematics', which indicates the intended level. The book is often warmly praised and is (I believe) quite widely used in Germany. There is a lot of material here, often covered well. But revisiting the book, I can't find myself wanting to recommend it as a good place to *start*. The core material on the syntax and semantics of first-order logic in Chs 2 and 3 is presented more accessibly and more elegantly elsewhere. And the treatment of a sequent calculus Ch. 4 strikes me as poor, with the authors (by my lights) mangling some issues of principle and certainly failing to capture the elegance that using a sequent calculus can bring. For more on this book, see tinyurl.com/EFTbooknote.

Shawn Hedman's *A First Course in Logic* (OUP, 2004) is subtitled 'An Introduction to Model Theory, Proof Theory, Computability and Complexity'. So there is no lack of ambition in the coverage! The treatment of basic FOL is patchy, however. It is pretty clear on semantics, and the book can be recommended to more mathematical readers for its treatment of more advanced model-theoretic topics (see §5.3 in this Guide). But Hedman offers a peculiarly ugly not-so-natural deductive system. As already noted, however, he *is* good on so-called resolution proofs. For more about what does and what doesn't work in Hedman's book, see [tinyurl.com/hedmanbook.

Peter Hinman's *Fundamentals of Mathematical Logic* (A. K. Peters, 2005) is a massive 878 pages, and as you'd expect covers a great deal. Hinman is, however, not really focused on deductive systems for logic, which don't make an appearance until over two hundred pages into the book. And most readers will find this book pretty tough going. This is not, then, the place to start with FOL. However, the first three chapters of the book certainly contain some interesting supplementary material once you have got hold of the basics from elsewhere, and could appeal to mathematicians. For more about what does and what doesn't work in Hinman's book, see tinyurl.com/hinmanbook.

Wolfgang Rautenberg's *A Concise Introduction to Mathematical Logic* (Springer, 2nd edn. 2006) has some nice touches. But I suspect that its first hundred pages on FOL are rather *too* concise to serve most readers as an initial introduction; and its preferred formal system is not a 'best buy' either. Can be recommended as good revision material, though.

Finally, Richard Kaye is the author of a particularly attractively written 1991 classic on models of Peano Arithmetic (we will meet this in Part III). So I had high hopes for his later *The Mathematics of Logic* (CUP 2007). "This book", he writes, "presents the material usually treated in a first course in logic, but in a way that should appeal to a suspicious mathematician wanting to see some genuine mathematical applications. . . . I do not present the main concepts and goals of first-order logic straight away. Instead, I start by showing what the main mathematical idea of 'a completeness theorem' is, with some illustrations that have real mathematical content." So the reader is taken on a mathematical journey starting with König's Lemma (I'm not going to explain that here!), and progressing via order relations, Zorn's Lemma (an equivalent to the Axiom of Choice), Boolean algebras, and propositional logic, to completeness and compactness of first-order logic. Does this very unusual route work as an introduction? I am not at all convinced. It seems to me that the journey is made too bumpy and the road taken is far too uneven in level for this to be appealing as an early trip through first-order logic. But if you *already* know a fair amount of this material from more conventional presentations, the different angle of approach in this book could be very interesting and illuminating.

(c) *Puzzles galore: What about some of Smullyan's other books?* Let me finish on a warmly positive note. I have already strongly recommended Raymond Smullyan's 1968 classic *First-Order Logic*. Smullyan went on to write some classic texts on Gödel's theorem and on recursive functions, which we'll be mentioning later. But as well as these, he also wrote many 'puzzle'-based books aimed at a wider audience, including e.g. the rightly famous *What is the Name of This Book?** (Dover Publications reprint of 1981 original, 2011) and *The Gödelian Puzzle Book** (Dover Publications , 2013).

Smullyan has also written *Logical Labyrinths* (A. K. Peters, 2009). From the blurb: "This book features a unique approach to the teaching of mathematical logic by putting it in the context of the puzzles and paradoxes of common language and rational thought. It serves as a bridge from the author's puzzle books to his technical writing in the fascinating field of mathematical logic. Using the logic of lying and truth-telling, the author introduces the readers to informal reasoning preparing them for the formal study of symbolic logic, from propositional logic to first-order logic, … The book includes a journey through the amazing labyrinths of infinity, which have stirred the imagination of mankind as much, if not more, than any other subject."

Smullyan starts, then, with puzzles, of this kind: you are visiting an island where there are Knights (truth-tellers) and Knaves (persistent liars) and then in various scenarios you have to work out what's true given what the inhabitants say about each other and the world. And, without too many big leaps, he ends with first-order logic (using tableaux), completeness, compactness and more. This is no substitute for standard texts, but – for those with a taste for being led up to the serious stuff via sequences of puzzles – an entertaining and illuminating supplement.

(Smullyan's later *A Beginner's Guide to Mathematical Logic** (Dover Publications, 2014) is rather more conventional. The first 170 pages are relevant to FOL. A rather uneven read, it seems to me; but again an engaging supplement to the main texts recommended above.)