4 First order logic

So, let's get down to business! This chapter starts with a quick overview of the topics we will be treating as belonging to the basics of first-order logic (or predicate logic, quantificational logic, call it what you will: I'll use 'FOL' for short). Then there are some main recommendations for texts covering these topics, followed by some suggestions for parallel and further reading. The chapter ends with some additional comments, mostly responding to frequently asked questions.

A note to philosophers. If you have carefully read and mastered a substantial introductory logic text for philosophers such as Nick Smith's, or even my own, you will already be familiar with (basic versions of) a significant amount of the material covered in this chapter. However, the big change is that you will now begin to see the perhaps familiar material 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 do work through the topics again, now with more mathematical precision.

4.1 FOL overview: basic topics

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 quite complex ways, introducing the kind of obscurities and possible ambiguities we 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, tidied-up, 'formalized' languages. We tackle a stretch of reasoning by first regimenting it into a suitable unambiguous formalized language with tidier logical operators, and then we can evaluate the reasoning once recast into this more transparent form. This way, we have a division of labour. There's the task of working out the intended structure of the original argument as we render it into a clear formal language. Then there's the separate business of assessing the validity of the resulting nicely regimented argument.

In FOL, therefore, we use appropriate formal languages which contain, in particular, (i) tidily disciplined surrogates for the connectives *and*, *or*, *if*, *not* (standardly symbolized $\wedge$, $\vee$, $\rightarrow$ and $\neg$) and (ii) replacements for the ordinary language quantifiers (roughly, using $\forall x$ for *every x is such that ...*, and $\exists y$ for *some y is such that...*). Although the fun really starts once we have the quantifiers in play, it is useful to develop FOL in two corresponding main stages. So:

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

(ii) We next develop full first-order languages which also have the apparatus of quantification – such languages still have very sparse logical apparatus but are in fact rich enough to regiment almost all mathematical reasoning. We can then explore the logic of arguments rendered into such languages.

We can usefully say just a bit more about the stages of this project now – though don't worry, of course, if you don't yet fully grasp the import of every point:

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

(i.b) On the *semantic* side, you need to understand the idea of a *valuation* for a propositional language: we start with an assignment of truth-values, *true* vs *false*, to the atomic components of our languages, and then explain how to evaluate complex sentences involving the connectives by using the 'truth-functional' interpretation of the connectives.

With these basic semantic ideas in play, we can now define the relation of semantic entailment, where the sentences Γ semantically entail the sentence φ when no possible valuation will make all the sentences Γ true without making φ true too. We then get to explore some of the key properties of this entailment relation, and learn how to calculate when the relation holds.

(i.c) Different textbook presentations of (i.a) and (i.b) will be pretty similar; but now the paths fork. For the next topic will be a deductive *proof-system* for propositional logic, and there is a variety of such systems to choose from. For a start – forgive the arm-waving! – we have

1. Old-school 'axiomatic' systems. (Proofs are simple linear sequences of formulas set out e.g. as a single vertical column; each formula in a proof after the initial premisses is either an instance of a logical axiom or follows from what's gone before by one of a couple of basic rules of inference.)

2. Natural deduction done Fitch-style. (We are this time allowed lots of rules of inference and so don't need logical axioms; in particular, we are now allowed to use forms of inference where we make temporary suppositions for the sake of argument – as when we temporarily assume φ, show it entails a contradiction and conclude that *not*-φ. To keep track of which temporary suppositions are in play at a given point of an argument, we use a multi-column layout, with the line of formulas in the proof jumping to and fro. The premisses sit at the top of the home column, and the desired conclusion at the bottom, but we *indent* the argument into a new column each time a new temporary supposition is made, and cancel the indentation when that supposition is dropped again.)

3. Natural deduction done Gentzen-style. (Again, as in informal reasoning, we are allowed lots of rules of inference and don't need logical axioms. But this time, we set out the proof in the form of a root-at-the-bottom-tree. Premisses appear at the top of branches; proof-branches can join, as when we combine φ on one branch and ψ on another to the infer their conjunction $\varphi \wedge \psi$; and the final conclusion is at the root of the tree. Again we are allowed to use inferences which discharge temporary suppositions, but we will need a different device to keep track of them).

4. 'Truth trees' or 'semantic tableaux'. (Instead of arguing directly from given premisses to our desired conclusion, we 'work backwards'. So we start by assuming that the premisses are true and the conclusion is false. Put these assumptions at the top of an upside-down tree. We aim to show that these initial assumptions are contradictory (so indeed, if the given premisses are true, so is the conclusion). We need to set our our working in a downwards-branching tree because we will need to explore possible alternatives at choice points.)

5. Sequent calculi. (Also due to Gentzen, our working is again set out as a a root-at-the-bottom-tree. But this time the 'nodes' of the tree are decorated not with single formulas but with so-called sequents – in one form, these are expressions like $\Gamma : \Delta$ which tell us that if the formulas Γ are all true then at least one of the formulas Δ is true.)

These different types of proof-system (or rather their specific versions for 'classical' propositional logic) are equivalent – meaning that, given some premisses, we can warrant the same conclusions in each system. But they differ considerably in their intuitive appeal and user-friendliness, as well as in some of their more technical features.

In fact, Fitch proofs and truth trees are easiest for beginners to manipulate in order to produce working formal proofs – which is why elementary logic books for philosophers usually introduce one or other (or both) systems. But, as we will see, mathematical logic text books at the next level of sophistication usually focus on the first and third types of system, on axiomatic systems and Gentzen-style natural deduction. True, axiomatic systems in their raw state can initially be pretty horrible to use – but proofs can be expedited once you learn some basic dodges (like the use of the so-called 'Deduction Theorem'). Gentzen-style proofs are by comparison very elegant, though they still take some getting used to compared to Fitch-style proofs which stick rather closer to informal patterns of reasoning. Finally, the sequent calculus really comes into its own in more advanced work on so-called proof theory.

At some point, the educated logician will want to find about *all* these proof styles – at the very least, you should get a general sense of how they respectively work, and come to appreciate the interrelations between them.

(i.d) We now have two different ways of defining what's a deductively good argument in propositional logic. We said that some premisses Γ semantically entail a conclusion φ if every possible valuation which makes Γ true makes φ true. We can now say that some premisses Γ entail the conclusion φ in the proof-system S if there is an S-type derivation of the conclusion φ from the premisses Γ (or a derivation of the sequent $\Gamma\colon \{\varphi\}$). 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 the premisses Γ, then φ really *is* semantically entailed by the premisses 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 semantically entailed by the premisses Γ, then there is indeed some S-derivation of φ from the premisses Γ.

So we will want proofs of soundness and completeness for our favoured proof-system S. These results will hold no terrors! However, in establishing soundness and completeness for propositional logics we encounter useful strategies which can later be beefed-up to give us soundness and completeness results for stronger logics.

(ii.a) Now for FOL (predicate logic, quantificational logic). Again, *syntax* first. And a story needs to be told in particular about why the syntax of the 'quantifier/variable' expressions of generality in formal logic is significantly different from the syntax of 'every' or 'some' in ordinary language. There are also decisions to be made about allowing the use of so-called 'free variables' and/or having 'dummy names' or 'parameters' in the language.

(ii.b) Turning to *semantics*: the first key idea we need is that of a *structure* (a domain of objects equipped with some relations and/or functions). And, crucially, you need to grasp the idea of an interpretation of a language in such a structure; you'll need to understand how such an interpretation generates a unique valuation (a unique assignment of truth-values) for every sentence of the interpreted language. A proper formal semantic story with the bells and whistles needed to cope with quantifiers is non-trivial.

With these semantic ideas to hand, you can again define a relation of semantic entailment, now for FOL expressions. The FOL sentences Γ semantically entail φ when no interpretation in any appropriate structure can make all the sentences Γ true without making φ true too.

You'll again need to know some of the basic properties of this entailment relation.

(ii.c) Next, we will want to explore a proof system for FOL. Corresponding to the five basic types of system we mentioned for propositional logic, you can again encounter five different types of proof system, with their varying attractions. And it should be said that even when we have chosen the basic type we want to work with, there remain other important choices to make – though we should end up each time with an equivalent system which warrants the same arguments. Again, you'll want at some point to find out about all these different styles of proof: but (as we said) we will start by looking at axiomatic and (one flavour of) natural deduction systems.

(ii.d) As with propositional logic, we will want *soundness* and *completeness* proofs which show that our chosen proof systems for FOL don't overshoot (giving us false positives) or undershoot (so we can't derive some semantically valid entailments). And at *this* point, it might be said, the study of FOL becomes really interesting: the completeness proof in particular involves more sophisticated ideas than anything we have met before. It also has intriguing corollaries, as we will see when we move from our initial investigation of FOL to begin so-called 'model theory', about which much more in the next chapter.

4.2 The main recommendations on FOL

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

1. 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'. It should in fact be entirely manageable for self study by philosophers and mathematicians alike (philosophers can skip over a few more mathematical illustrations).

   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 interesting 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.

   Very warmly recommended, then. For the moment, you only *need* read up to and including §7.7 (under two hundred pages). But having got that far, you might as well read the final couple of sections and the Postlude too! The book has brisk solutions to some of the exercises. (A demerit mark, though, to OUP for not publishing C&H more cheaply.)

Next, you should complement C&H by reading the first third of

2. Christopher Leary and Lars Kristiansen's *A Friendly Introduction to Mathematical Logic*** (1st edn by Leary alone, Prentice Hall 2000; inexpensive 2nd edn Milne Library 2015; and now freely available for download).

   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, this time done in axiomatic style. At this stage you could stop reading after §3.2, which means you will be reading just 86 pages.

   Unusually, L&K dive straight into a full treatment of first-order logic without spending an introductory chapter or two on propositional logic. But that happily means (in the present context) that you won't feel that you are labouring through the very beginnings of logic yet 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 C&H very much indeed, and say more about this admirable book in another book note.

For additional motivation for the quantifier/variable notation, see Chapters 26–28 of Peter Smith, *Introduction to Formal Logic*** (2nd edn. downloadable from www.logicmatters.net/ifl). For more explanation of the Gentzen-style proof system treated perhaps just a little quickly by C&H, see Chapter 9, 'Natural Deduction' of Richard Zach et al., *Sets, Logic, Computation*** (downloadable from the Open Logic Project, https://slc.openlogicproject.org).

Next, here's an alternative to the C&H/L&K pairing which is very approachable and can also be warmly recommended:

3. Derek Goldrei's *Propositional and Predicate Calculus: A Model of Argument** (Springer, 2005) is explicitly designed for self-study. Read Chs. 1 to 5 (you could skip §§4.4 and 4.5, 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. So this is a somewhat more 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 very great deal of thought has gone into making this text as helpful as possible. Some may find it occasionally goes a bit too slowly and carefully, though I'd say that this is erring on the right side in an introductory book: if you struggle slightly with the alternative reading, or just want a comfortingly manageable text, you should find this exceptionally accessible. Or you might just warm more to Goldrei's style anyway.

Like L&K, Goldrei uses an axiomatic system (which is one reason why, on balance, I still recommend starting with C&H instead). As with C&H and L&K, I like the tone and approach a great deal.

Supplement Goldrei with a brief glance at a natural deduction proof system, e.g. again as in the Open Logic Project's *Sets, Logic, Computation* Ch. 9.

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

4.3 Some parallel/additional 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 will list – in very rough order of difficulty/sophistication – a small handful of further texts which could well make for useful parallel or additional reading. Then in the following section, I will mention some other books I've been asked about.

I'll begin with a book written by a philosopher for philosophers:

4. 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 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 (as noted before, it is important eventually to understand what is going on in these different kinds of proof-system). So Bostock ranges more widely but less deeply than e.g. Goldrei. 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 therefore be, as intended, particularly accessible to philosophers who haven't done much formal logic before and should, if read in parallel, help ease the transition to coping with the more mathematical style of the books recommended in the last section.

To repeat, 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, here's a short and painless introduction to trees for propositional logic which will also give you the basic idea: Peter Smith, Propositional Truth Trees.

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

5. 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 more conventional textbook presentations of C&H, L&K and/or Goldrei.

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

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

6. Herbert Enderton, *A Mathematical Introduction to Logic* (Academic Press 1972, 2002). This also focuses on a Hilbert-style axiomatic system, and is often regarded as a classic of exposition. However, it does strike me as somewhat more difficult than L&K, so I'm not surprised that students often report finding it a bit challenging *if used by itself as a first text.* Still, it is an admirable and very reliable piece of work which you should be able to cope with well if used as a supplementary second text, e.g. after you have tackled C&H.

   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 this book note.

I'll also mention – though this time with a little hesitation – another much used text which has gone through multiple editions and should also be in any library; it is a useful natural-deduction based alternative to C&H. Later chapters of this book are also mentioned later in this Guide as possible reading for more advanced work, so it could be worth making early acquaintance with . . .

7. Dirk van Dalen, *Logic and Structure** (Springer, 1980; 5th edition 2012). The chapters up to and including §3.2 provide an introduction to FOL via natural-deduction. The treatment *is* often approachable and written with a relatively light touch. However – and this explains my hesitation – 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, 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. (See my more extended review of the whole book.)

Going up a level of mathematical sophistication, we get to absolute classic, short but packed with good things:

8. Raymond Smullyan, *First-Order Logic** (Springer 1968, Dover Publications 1995). This is terse, but those with a taste for mathematical elegance can certainly try 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 delight for the mathematically minded.

Finally, taking things in a new direction, don't be put off by the title of

9. 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 various places in the book there are illustrations of how to implement various algorithms in Prolog. But either you can easily pick up the very small amount of background knowledge about Prolog that's needed to follow everything that is going on (and that's quite a fun thing to do anyway) or you can just skip those implementation episodes.

   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 – yes, a sixth style of proof! – in this exceptionally clearly written book.

This book's emphasis is, then, rather different from most of the other recommended books. So I initially hesitated to mention it here in this Guide. However, I do think that the fresh light thrown on first-order logic makes the detour through this 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.)

4.4 Other treatments?

I will end this chapter by responding to a variety of Frequently Asked Questions, mostly questions raised in response to earlier versions of the Guide. Occasionally, I have to be pretty negative.

Hold on! What about an introductory text with a more proof-theoretic slant? Well, it's 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. Suppose we are working in some formal system and (i) have so far proved α and also proved β; then (ii) we can (rather boringly!) infer the conjunction $\alpha \wedge \beta$ (remember '$\wedge$' means *and*). Now, suppose later in the proof (iii) we appeal to that conjunction $\alpha \wedge \beta$ and infer α. We wouldn't have gone *wrong*; but obviously the detour has been *pointless*, given that at stage (i) we could already prove α. So 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 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, you could look at Jan von Plato's *Elements of Logical Reasoning** (CUP, 2014). This is based on the author's lectures for an introductory 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 any classroom work to round things out, this book might not be as easily accessible as the author intends. But suppose you have already met one system of natural deduction (e.g., as in C&H), and now want to know more about 'proof-theoretic' aspects of this and related systems. Suppose, for example, that you want to know about variant ways of setting up ND systems, about proof-search, about the relation with so-called sequent calculi, etc. Then this book makes a reasonably clear and approachable start. However, my own recommendation would be to concentrate for the moment on the core topics covered by the books mentioned so far, and then later dive into proof theory proper, covered in Chapter **??**.

A blast from the past: What about Mendelson? Somewhat to my surprise, perhaps the most frequent question I used to get asked in response to early versions of the Guide is 'But what about Mendelson, Chs. 1 and 2'? Well, Elliot Mendelson's *Introduction to Mathematical Logic* (Chapman and Hall/CRC 6th edn 2015) was first published when I was a student and the world was a great deal younger. 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 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 battling with. But unsurprisingly, fifty years on, there are many rather more accessible and more amiable alternatives for beginning serious logic. Mendelson's book is a landmark worth visiting one day, but I can't recommend starting there. For a little more about it, see here.

(As an aside, if you *do* want an old-school introduction from roughly the same era, you might more enjoy Geoffrey Hunter, *Metalogic** (Macmillan 1971, University of California Press 1992). This is not groundbreaking in the way e.g. Smullyan's *First-Order Logic* is, nor is it as comprehensive as Mendelson: but it was an exceptionally good textbook from a time when there were few to choose from, and I still regard it with admiration. 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 revisit in §6.1. Unfortunately, the typography – from pre-LaTeX days – isn't at all pretty to look at: this can make the book's pages initially appear rather unappealing. 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!)

Two more standard mathematical logic texts: what about Ebbinghaus, Flum and Thomas? what about Rautenberg? We now turn to 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. But revisiting it, I can't find myself wanting to recommend it as a good place to start, either for philosophers or for mathematicians. 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 mangling some issues of principle and failing totally to capture the elegance that a sequent calculus can have. For more on this book, see this book note.

Wolfgang Rautenberg's *A Concise Introduction to Mathematical Logic* (Springer, 2nd edn. 2006) has some nice touches. But its first hundred pages on FOL are probably *too* concise to serve most readers as a first introduction; and its preferred formal system is not a 'best buy' either. §1.5 on applications of the compactness theorem for propositional logic is adding to your reading list, though.

The very latest thing: what about the Open Logic Text? This is an entirely admirable, collaborative, open-source, enterprise inaugurated by Richard Zach, and very much work in progress. You can download the latest full version from this page.

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." My sense is that the chapters on propositional and quantificational logic – mixed into *Sets, Logic, Computation*** – have been expanded considerably, and are now much more book-like. They still, however, have some rather idiosyncratic features (for example why *start* with the LK system of sequent calculus?). But my main worry is that the discussions probably still go a little too fast to work as a purely stand-alone introduction for initial self-study without the benefit of lecture support. But your mileage may vary. And certainly, these notes could be *very* useful for reinforcing/revision.

Designed for philosophers: Why not The Logic Book? What about Sider? What about Bell, DeVidi and Solomon? 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 doubt that those students 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 with the books mentioned in the last section. 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.

Theodore Sider – a very well-known philosopher – has written a text called *Logic for Philosophy** (OUP, 2010) aimed at philosophers, 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 **??**. 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, I can't at all recommend the first half of this book: I explain why in another book note. True, a potentially attractive additional feature of this part of Sider's book is that it does contain brief discussions about e.g. some non-classical propositional logics, and about descriptions and free logic. But remember all this is being done in just 130 pages, which means that things are whizzing by very fast, so the breadth of Sider's coverage here goes with far too much superficiality. If you want some breadth, Bostock's book is still very much to be preferred.

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 some breadth, you'd again do better sticking with more expansive Bostock.

And I have *still* only touched on a small proportion of books we could mention! In particular, the other Big Books on Mathematical Logic of course have treatments of the relevant material: see the Guide's Appendix for more detailed discussion of some options. However, I think none of these Big Books hit quite the right level and pace to make them ideal for a *first* encounter with basic FOL, at least for many readers.

But I don't want to finish on a negative note: so finally . . .

Puzzles galore: What about some of Smullyan's other books? I have already warmly recommended Smullyan's 1968 classic *First-Order Logic*. He went on to write some classic and very accessible texts on Gödel's theorem and on recursive functions, which we'll be mentioning later. But as well as these, Smullyan wrote many 'puzzle' based-books aimed at a wider audience, including the justly famous 1981 *What is the Name of This Book?** (Dover Publications reprint, 2011).

More recently, he wrote *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, e.g. 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 from 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 more conventional. The first 170 pages are relevant to FOL. A rather uneven read, it seems to me, but again perhaps an engaging supplement to the texts recommended above.