7 Set theory, less naively

In Chapter 2, we touched on some elementary concepts and constructions involving sets. We now go further into set theory, though still not beyond the beginnings that any logician really ought to know about. In Part III of the Guide we will return to cover more advanced topics like 'large cardinals', proofs of the consistency and independence of the Continuum Hypothesis, and a lot more besides. But this present chapter concentrates on some core basics.

7.1 Elements of set theory: an overview

Even more, perhaps, than previous overviews, this one may well fall squarely between two stools, being too elementary for mathematicians and not explanatory enough for those readers whose background is purely in philosophy. So if, for one reason or the other, you find that these preliminary remarks aren't particularly helpful for you, then do simply skip on to the next section which gives the main reading recommendations.

(a) If you have not already done so, you now want to get a really firm grip on the key facts about the 'algebra of sets' (concerning unions, intersections, complements and how they interact). You also need to know, inter alia, the basics about powersets, about encoding pairs and other finite tuples using unordered sets, and about Cartesian products, the extensional treatment of relations and functions, the idea of equivalence classes, and how to treat infinite sequences as sets (see Chapter 2).

(b) Moving on, one fundamental early role for set theory was "putting the theory of real numbers, and classical analysis more generally, on a firm foundation". What does this involve?

It only takes a finite amount of data to fully specify a particular natural number. Similarly for integers and rational numbers. But not so for real numbers. As is very familiar, a real can rendered e.g. by an infinite sequence of ever-closer rational approximations, but these need never terminate. So in theorizing about real numbers we are tangling with the infinite. Set theory gives us a framework for reasoning about such non-finite data. How?

Assume, for the moment, that do we already have the rational numbers to hand, and let's now define the idea of a sequence of ever-closer rational approximations more carefully. A *Cauchy sequence*, then, is an infinite sequence of rationals $s_1, s_2, s_3, \ldots$ which *converges* – i.e. the differences $|s_m - s_n|$ are as small as we want, once we get far enough along the sequence. More carefully, take any $\epsilon > 0$ however small, then for some k, $|s_m - s_n| < \epsilon$ for all $m, n > k$. Now say that two Cauchy sequences $s_1, s_2, s_3, \ldots$ and $s'_1, s'_2, s'_3, \ldots$ are *equivalent* if their members eventually get arbitrarily close – i.e. when we take any $\epsilon > 0$ however small, then for some k, $|s_n - s'_n| < \epsilon$ for all $n > k$. Cauchy identifies real numbers with equivalence classes of Cauchy sequences. So, for Cauchy, $\sqrt{2}$ would be the equivalence class containing any sequence of rationals like 1.4, 1.41, 1.414, 1.4142, 1.41421, …, i.e. rationals whose squares approach 2.

Alternatively, dropping the picture of sequential approach, we can identify a real number with a *Dedekind cut*, defined as a (proper, non-empty) subset C of the rationals which (i) is downward closed – i.e. if $q \in C$ and $q' < q$ then $q' \in C$ – and (ii) has no largest member. For example, take the negative rationals together with the positive ones whose square is less than two: these form a cut. Dedekind (more or less) identifies the positive irrational $\sqrt{2}$ with the cut we just defined.

Assuming some set theory, we can now show that – whether defined as cuts on the rationals or defined as equivalence classes of Cauchy sequences of rationals – the real numbers have the desired properties either way. Assuming our set theory is consistent, the resulting theory of the reals can be shown to be consistent too.

We can then go on define functions between real numbers in terms of sets of ordered tuples of reals. I won't spell this out further here. However, you should get to know something of how the overall story goes, and also get some sense of what assumptions about sets are needed for the story to work to give us a basis for reconstructing classical real analysis. (You will need a number of levels of sets: sets of rationals, and sets of sets of rationals, and sets of sets of sets, and up a few more levels depending on the details.)

(c) Now, as far as construction of the reals and the foundations of analysis are concerned, we could take the requisite set theory – the apparatus of infinite sets, infinite sequences, equivalence classes and the rest – as describing a *superstructure* sitting on top of a given prior basic universe of rational numbers already governed by a prior suite of numerical laws. However, we don't need to do this. For we can in fact *already* construct the rationals and simpler number systems within set theory itself.

For the naturals, pick any set you like and call it '0'. And then consider e.g. the sequence of sets $0; \{0\}; \{\{0\}\}; \{\{\{0\}\}\}; \ldots$. Or alternatively, consider the sequence $0; \{0\}; \{0, \{0\}\}; \{0, \{0\}, \{0, \{0\}\}\}; \{0, \{0\}, \{0, \{0\}\}, \{0, \{0\}, \{0, \{0\}\}\}\}; \ldots$ where at each step after the first we extend the sequence by taking the set of all the sets we have so far. Either sequence then has the structure of the natural-number series. There is a first member; every member has a unique successor (which is distinct from it); different members have different successors; the sequence never circles around and starts repeating. So such a sequence of sets will do as a representation, implementation, or model of the natural numbers (call it what you will).

Let's not get hung up about the best way to describe the situation; we will simply say we have constructed a natural number sequence. And elementary reasoning about sets will show that the familiar arithmetic laws about natural numbers apply to numbers as just constructed (including e.g. the principle of arithmetical induction).

Once we have a natural number sequence we can go on to construct the integers from it in various ways. Here's one. Informally, any integer equals $m - n$ for some natural numbers m, n (to get a negative integer, take $n > m$). So, first shot, we can treat an integer as an ordered pair of natural numbers. But since $m - n = m' - n'$ for lots of m', n', choosing a particular pair of natural numbers to represent an integer involves an arbitrary choice. So, a neater second shot, we can treat an integer as an equivalence class of ordered pairs of natural numbers (where the pairs $\langle m, n\rangle$ and $\langle m', n'\rangle$ are equivalent in the relevant way when $m + n' = m' + n$). Again the usual laws of integer arithmetic can then be proved from basic principles about sets.

Similarly, once we have constructed the integers, we can construct rational numbers in various ways. Informally, any rational equals p/q for integers p, q, with $q \neq 0$. So, first shot, we can treat a rational numbers as a particular ordered pair of integers. Or to avoid making a choice between equivalent renditions, we can treat a rational as an equivalence class of ordered pairs of integers.

We again needn't go further into the details here, though – at least once in your mathematical life! – you will want to see them worked through enough to confirm that these can constructions can indeed all be done. The point we want to emphasize now is simply this: once we have chosen an initial object to play the role of 0 – the empty set is the conventional choice – and once we have a set-building operation which we can iterate sufficiently often, and once we can form equivalence classes from among sets we have already built, we can construct sets to do the work of natural numbers, integers and rationals in standard ways. Hence, we don't need a theory of the rationals prior to set theory before we can go on to construct the reals: *the whole game can be played inside pure set theory.*

(d) Another theme. It is an elementary idea that two sets are equinumerous (have the same cardinality) just if we can match up their members one-to-one, i.e. when there is a one-to-one correspondence, a bijection, between the sets. It is easy to show that the set of even natural numbers, the set of primes, the set of integers, the set of rationals are all *countably* infinite in the sense of being equinumerous with the set of natural numbers.

By contrast, as we noted in §2.1, a simple argument shows that the set of infinite binary strings is not countably infinite. Two corollaries:

1. An infinite binary string can be thought of as representing a set of natural numbers, namely the set which contains n if and only if the n-th digit in the string is 1; and different strings represent different sets of naturals. Hence the powerset of the natural numbers, i.e. the set of subsets of the naturals, is also not countably infinite.

2. An infinite binary string can equally well be thought of as representing a real number between 0 and 1 in binary; and different strings represent

different reals. So the set of real numbers between 0 and 1 is not countably infinite either – hence neither is the set of all the real numbers.

And now a famous question arises – easy to ask, but (it turns out) extraordinarily difficult to answer. Take an infinite collection of real numbers. It could be equinumerous with the set of natural numbers (like, for example, the set of *real* numbers 0, 1, 2, ...). It could be equinumerous with the set of all the real numbers (like, for example, the set of irrational numbers). But are there any infinite sets of reals of intermediate size (so to speak)? – can there be an infinite subset of real numbers that can't be put into one-to-one correspondence with just the natural numbers and can't be put into one-to-one correspondence with all the real numbers either? Cantor conjectured that the answer is 'no'; and this negative answer is known as the Continuum Hypothesis.

Efforts to confirm or refute the Continuum Hypothesis were a major driver in early developments of set theory. We now know the problem is a profound one – the standard axioms of set theory can't settle the hypothesis one way or the other. Is there some attractive and natural additional axiom which will settle the matter? I'll not give a spoiler here! – but exploration of this question takes us way beyond the initial basics of set theory.

(e) The argument that the power set of the naturals isn't equinumerous with the set of naturals can be generalized. Cantor's Theorem tells us that a set is *never* equinumerous with its powerset.

Note, there *is* a bijection between the set A and the set of singletons of elements of A; in other words, there *is* a bijection between A and *part* of its powerset $\mathcal{P}(A)$. But we've just seen that there is no bijection between A and the *whole* of $\mathcal{P}(A)$. Intuitively then, A is smaller in size than $\mathcal{P}(A)$, which will in turn be smaller than $\mathcal{P}(\mathcal{P}(A))$, etc. We now want to develop this intuitive idea of one set's having a smaller cardinal size than another into a competent general theory about relative cardinal size.

(f) Let's pause to consider the emerging picture.

Starting perhaps from some given urelements – elements which don't themselves have members – we can form sets of them, and then sets of sets, sets of sets of sets, and so on and on: and at each new level, we accumulate more and more sets formed from the urelements and/or the sets formed at earlier levels. At each level, more and more sets are formed. In particular, once we have an infinite number of entities at one level, we get an even greater infinity of entities at the next as we form powersets, and so on up.

Now, for purely mathematical purposes such as reconstructing analysis, it seems that a we only need a single non-membered base-level entity, and it is tidy to think of this as the empty set. So for internal mathematical purposes, we can take the whole universe of sets to contain only 'pure' sets (when we look at the members of members of ... members of sets, we find nothing other than more sets). But what if we want to be able to apply set-theoretic apparatus in talking about e.g. widgets or wombats or (more seriously!) space-time points? Then it might seem that we will want the base level of non-membered elements to be populated with those widgets, wombats or space-time points as the case might be. However, it seems that we can always *code* for widgets, wombats or space-time points using some kind of numbers, and we can treat those numbers as sets. So our set-theory-for-applications can *still* involve only pure sets. That's why typical introductions to set theory either explicitly restrict themselves to talking about pure sets, or after officially allowing the possibility of urelements promptly ignore them.

(g) Lots of questions arise. Here are two:

1. First, how far can we iterate the 'set of' operation – how high do these levels upon levels of sets-of-sets-of-sets-of-... stack up? Once we have the natural numbers in play, we only need another dozen or so more levels of sets in which to reconstruct 'ordinary' mathematics: but now we are embarked on set theory for its own sake, how far can we go up the hierarchy of levels?

2. Second, at a particular level, how many sets do we get at that level? And a prior question, how do we 'count' the members of infinite sets?

   With finite sets, we not only talk about their relative sizes (larger or smaller), but actually count them and give their absolute sizes by using finite cardinal numbers. These finite cardinals are the natural numbers, which we have learnt can be identified with particular sets. We now want similarly to have a story about the infinite case; we not only want an account of relative infinite sizes but also a theory about infinite cardinal numbers apt for giving the size of infinite collections. Again these infinite cardinals will be identified with particular sets. But how can this story go?

It turns out that to answer both these questions, we need a new notion, the idea of infinite ordinal numbers. We can't say very much about this here, but some arm-waving pointers might be useful.

(h) Let's start rather naively. Here are the familiar natural numbers, but re-sequenced with the evens in their usual order before the odds in *their* usual order:

$$0, 2, 4, 6, \ldots, 1, 3, 5, 7, \ldots.$$

If we use '$\sqsubset$' to symbolize the order-relation here, then $m \sqsubset n$ just in case either (i) m is even and n is odd or else (ii) m and n have the same parity and $m < n$. Note that $\sqsubset$ is a *well-ordering* in the standard sense that it is a linear order and, for any numbers we take, one will be the $\sqsubset$-least.

Now, if we march through the naturals in their new $\sqsubset$-ordering, checking off the first one, the second one, the third one, etc., where does the number 7 come in the order? Plainly, we cannot reach it in any finite number of steps: it comes, in a word, *transfinitely* far along the $\sqsubset$-sequence. So if we want a position-counting number (officially, an *ordinal* number) to tally how far along our well-ordered sequence the number 7 is located, we will need a transfinite ordinal. We will have to say something like this: We need to march through all the even numbers, which here occupy positions arranged exactly like all the natural numbers in their natural order. And then we have to go on another 4 steps. Let's use 'ω' to indicate the length of the sequence of natural numbers in their natural order, and we'll call a sequence structured like the naturals in their natural order an ω-sequence. The evens in their natural order can be lined up one-to-one with the naturals in order, so form another ω-sequence. Hence, to indicate how far along the re-sequenced numbers we find the number 7, it is then tempting to say that it occurs at $\omega + 4$-th place.

And what about the whole sequence, evens followed by odds? How long is it? How might we count off the steps along it, starting 'first, second, third, ...'? After marching along as many steps as there are natural numbers in order to treck through the evens, then – pausing only to draw breath – we have to march on through the odds, again going through positions arranged like all the natural numbers in their natural ordering. So, we have two ω-sequences, put end to end. It is very natural to say that the positions in the whole sequence are tallied by a transfinite ordinal we can denote $\omega + \omega$.

Here's another example. There are familiar maps for coding ordered pairs of natural numbers by a single natural: take, for example, the function which maps m, n to $[m, n] = 2^m(2n + 1) - 1$. And consider the following ordering on these 'pair-numbers' $[m, n]$:

$$[0, 0], [0, 1], [0, 2], \ldots, [1, 0], [1, 1], [1, 2], \ldots, [2, 0], [2, 1], [2, 2], \ldots, \ldots$$

If we now use '$\prec$' to indicate this order, then $[m, n] \prec [m', n']$ just in case either (i) $m < m'$ or else (ii) $m = m'$ and $n < n'$. (This type of ordering is standardly called *lexicographic*: in the present case, compare the dictionary ordering of two-letter words drawn from an infinite alphabet.) Again, $\prec$ is a well-ordering on the natural numbers.

Where does $[5, 3]$ come in this sequence? Before we get to this 'pair' there are already five blocks of the form $[m, 0], [m, 1], [m, 2], \ldots$ for fixed m, each as long as the naturals in their usual order, first the block with $m = 0$, then the block with $m = 1$, and three more blocks, each ω long; and then we have to count another four steps along, tallying off $[5, 0], [5, 1], [5, 2], [5, 3]$. So it is inviting to say we have to count along to the $\omega \cdot 5 + 4$-th step in the sequence to get to the 'pair' $[5, 3]$.

And what about the whole sequence of 'pairs'? We have blocks ω long, with the blocks themselves arranged in a sequence ω long. So this time it is tempting to say that the positions in the whole sequence of 'pairs' are tallied by a transfinite ordinal we can indicate by $\omega \cdot \omega$.

We can continue. Suppose we re-arrange the natural numbers into a new well-ordering like this: take all the numbers of the form $2^l \cdot 3^m \cdot 5^n$, ordered by ordering the triples $\langle l, m, n\rangle$ lexicographically, followed by the remaining naturals in their normal order. We tally positions in *this* sequence by the transfinite ordinal $\omega \cdot \omega \cdot \omega + \omega$. And so it goes.

Note by the way that we have so far been considering just (re)orderings of the familiar set of natural numbers – the sequences are equinumerous, and have the same infinite *cardinal* size; but the well-orders are tallied by different infinite *ordinal* numbers. Or so we want to say.

But is this sort of naive talk of transfinite ordinals really legitimate? Well, it was one of Cantor's great and lasting achievements to show that we can indeed make perfectly good sense of all this.

Now, in Cantor's work the theory of transfinite ordinals is already entangled with his nascent set theory. Von Neumann later cemented the marriage by giving the canonical treatment of ordinals in set theory. And it is via this treatment that students now typically first encounter the arithmetic of transfinite ordinals, some way into a full-blown course about set theory. This approach can, unsurprisingly, give the impression that you have to buy into quite a lot of set theory in order to understand even the basics about ordinals and their arithmetic. However, not so. Our little examples so far are of recursive (re)orderings of the natural numbers – i.e. a computer can decide, given two numbers, which way round they come in the ordering. There is a whole theory of recursive ordinals which talks about how to tally the lengths of such (re)orderings of the naturals, which has important applications e.g. in proof theory. And these tame beginnings of the theory of transfinite ordinals needn't entangle us with the kind of rather wildly infinitary and non-constructive ideas characteristic of modern set theory.

(i) However, here we *are* concerned with set theory, and so our next topic will naturally be von Neumann's very elegant implementation of ordinals in set theory as the 'hereditarily transitive sets'. The basic idea is to define a particular well-ordered sequence of sets – call them the $\text{ordinals}_{\text{vN}}$ – and show that any well-ordered collection of objects, however long the ordering, will have the same type of ordering as an initial segment of these $\text{ordinals}_{\text{vN}}$. So we can use the $\text{ordinals}_{\text{vN}}$ as a universal measuring scale against which to tally the length of any well-ordering.

And at this point, I'll have to leave it to you to explore the details of the construction of the $\text{ordinals}_{\text{vN}}$ in the recommended readings. But once we have them available, we can say more about the way that the universe of sets is structured; we can take the levels to be indexed by $\text{ordinals}_{\text{vN}}$ (and then assume that for every ordinal there is a corresponding level of the universe).

We can also now define a scale of cardinal size. We noted that well-orderings of different ordinal length can be equinumerous; different $\text{ordinals}_{\text{vN}}$ can have the same cardinality. So von Neumann's next trick is to define a cardinal number to be the first ordinal (in the well-ordered sequence of ordinals) in a family of equinumerous ordinals. Again this neat idea we'll have leave for the moment for later exploration. However – and this is an important point – to get this to all work out as we want, in particular to ensure that we can assign any two non-equinumerous sets respective cardinalities κ and λ such that either $\kappa < \lambda$ or $\lambda < \kappa$, we will need the Axiom of Choice. (This is something to keep looking out for in beginning set theory: where do we start to need to appeal to some Choice principle?)

(j) We are perhaps already rather past the point where scene-setting remarks at this level of arm-waving generality can be very helpful. Time to dive into the details! But one final important observation before you start.

The themes we have been touching on can and perhaps should initially be presented in a relatively informal style. But something else that also belongs here near the beginning of your first forays into set theory is an account of the development of axiomatic ZFC (Zermelo-Fraenkel set theory with Choice) as the now standard way of formally regimenting set theory. As you will see, different books take different approaches to the question of just *when* it is best to start getting more rigorously axiomatic, formalizing our set-theoretic ideas.

Now, there's a historical point worth noting, which explains something about the shape of the standard axiomatization. You'll recall from the remarks in §2.1(b) that a set theory which makes the assumption that *every* property has an extension will be inconsistent. So Zermelo set out in an epoch-making 1908 paper to lay down what he thought were the basic assumptions about sets that mathematicians actually *needed*, while not overshooting and falling into such contradictions. His axiomatization was not, it seems, initially guided by a positive conception of the universe of sets so much as by the desire to keep safe and not assume too much. But in the 1930s, Zermelo himself and especially Gödel came to develop the conception of sets as a hierarchy of levels (with new sets always formed from objects at lower levels, so never containing themselves, and with no end to the levels where we form more sets from what we have accumulated so far, so we never get to a paradoxical set of all sets). This cumulative hierarchy is described and explored in the standard texts. Once this conception is in play, it does invite a more direct and explicit axiomatization as a story about levels and sets formed at levels: however, it was only much later that this positively motivated axiomatization gets spelt out, particularly in what has come to be called Scott-Potter set theory. Most text books stick for their official axioms to the Zermelo approach, hence giving what looks to be a rather unmotivated selection of axioms whose attraction is that they all look reasonably modest and separately in keeping with the hierarchical picture, so unlikely to get us into trouble. In particular the initial recommendations below take this conventional line.

7.2 Main recommendations on set theory

This present chapter is, as advertised, just about the basics of set theory. Even here, however, there are is a very large number of books to choose from, so an annotated Guide will (I hope!) be particularly welcome.

But first, if you want a more expansive 35pp. overview of basic set theory, with considerably more mathematical detail and argument, I think the following chapter (the best in the book?) works pretty well:

1. Robert S. Wolf, *A Tour Through Mathematical Logic* (Mathematical Association of America, 2005), Ch. 2, 'Axiomatic Set Theory'.

And let me me mention again an introduction to set-theoretic ideas which I noted in §2.2, which you may have skipped past then.

2. Cambridge lecture notes by Tim Button have become incorporated into *Set Theory: An Open Introduction*** (2019) tinyurl.com/opensettheory, and this short book is one of the most successful outputs from the Open Logic Project. Its earlier chapters in particular are extremely good, and are very clear on the conceptual motivation for the iterative conception of sets and its relation to the standard ZFC axiomatization. However, things get a bit patchier as the book progresses: later chapters on ordinals, cardinals, and choice, get rather tougher, and might work better (I think) as parallel readings to the more expansive main recommendations I'm about to make. But very well worth looking at.

However, Button can't get into enough detail in his brisk notes, so most readers will want to look instead at one or other of the first two of the following admirable 'entry level' treatments which cover rather more in a bit more depth but still very accessibly:

3. Derek Goldrei, *Classic Set Theory* (Chapman & Hall/CRC 1996). The author taught at the Open University, and wrote specifically for students engaged in remote learning: his book has the friendly subtitle 'For guided independent study'. The result as you might expect – especially if you looked at Goldrei's FOL text mentioned in §3.3 – is exceptionally clear, and it is indeed admirably well-structured for independent self-teaching. Moreover, it is rather attractively written (as set theory books go!). The coverage is very much as as outlined in our overview. And one particularly nice feature is the way the book (unusually?) spends enough time motivating the idea of transfinite ordinal numbers before turning to their now conventional implementation in set theory.

4. Herbert B. Enderton's, *The Elements of Set Theory* (Academic Press, 1977) forms a trilogy along with the author's *Logic* and *Computability* which we have already mentioned in earlier chapters.

   This book again has exactly the coverage we need at this stage. But more than that, it is particularly clear in marking off the informal development of the theory of sets, cardinals, ordinals etc. (guided by the conception of sets as constructed in a cumulative hierarchy) from the ensuing formal axiomatization of ZFC. It is also particularly good and non-confusing about what is involved in (apparent) talk of classes which are too big to be sets – something that can mystify beginners. It is written with a certain lightness of touch and proofs are often presented in particularly well-signposted stages. The last couple of chapters perhaps do get a bit tougher, but overall this really is quite exemplary exposition.

Also starting from scratch, we find two further excellent books which are rather less conventional in style:

5. Winfried Just and Martin Weese, *Discovering Modern Set Theory I: The Basics* (American Mathematical Society, 1996). This covers overlapping ground to Goldrei and Enderton, but perhaps more zestfully and with a little more discussion of conceptually interesting issues. At some places, it is more challenging – the pace can be a bit uneven.

   I like the style a lot, and think it works very well. I don't mean the occasional (slightly laboured?) jokes: I mean the in-the-classroom feel of the way that proofs are explored and motivated, and also the way that teach-yourself exercises are integrated into the text. The book is evidently written by enthusiastic teachers, and the result is very engaging. (The story continues in a second volume.)
6. Yiannis Moschovakis, *Notes on Set Theory* (Springer, 2nd edition 2006). This also takes a slightly more individual path through the material than Goldrei and Enderton, with occasional bumpier passages, and with glimpses ahead. But to my mind, this is very attractively written, and again nicely complements and reinforces what you'll learn from the more conventional books.

Of these two pairs of books, I'd rather strongly advise reading one of the first pair and then one of the second pair.

I will add two more firm recommendations at this level. The first might come as a bit of surprise, as it is something of a 'blast from the past'. But we shouldn't ignore old classics – they can have a lot to teach us even when we have read the more recent books, and this is very illuminating:

7. Abraham Fraenkel, Yehoshua Bar-Hillel and Azriel Levy, *Foundations of Set-Theory* (North-Holland, originally 1958; but you want the revised 2nd edition 1973): Chapters 1 and 2 are the immediately relevant ones. Both philosophers and mathematicians should appreciate the way this puts the development of our canonical ZFC set theory into some context, and also discusses alternative approaches. Standard textbooks can present our canonical theory in a way that makes it seem that ZFC has to be the One True Set Theory, so it is worth understanding more about how it was arrived at and where some choice points are. This book really is attractively readable, and should be very largely accessible at this early stage. I'm not myself an enthusiast for history for history's sake: but it is very much worth knowing the stories that unfold here.

Now, as I noted in the initial overview section, one thing that every set-theory novice now acquires is the picture of the universe of sets as built up in a hierarchy of stages or levels, each level containing all the sets at previous levels plus new ones (so the levels are cumulative). It is significant that, as Fraenkel et al. makes clear, the picture wasn't firmly in place from the beginning. But the hierarchical conception of the universe of sets is brought to the foreground in

8. Michael Potter, *Set Theory and Its Philosophy* (OUP, 2004). For philosophers and for mathematicians concerned with foundational issues this surely is a 'must read', a unique blend of mathematical exposition (mostly about the level of Enderton, with a few glimpses beyond) and extensive conceptual commentary. Potter is presenting not straight ZFC but a very attractive variant due to Dana Scott whose axioms more directly encapsulate the idea of the cumulative hierarchy of sets. It has to be said that there are passages which are harder going, sometimes because of the philosophical ideas involved, and sometimes because of occasional expositional compression. However, if you have already read a set theory text from the main list, you should have no problems.

7.3 Some parallel/additional reading on standard ZFC

There are so many good set theory books with different virtues, many by very distinguished authors, that I should certainly pause to mention some more.

Let me begin by mentioning a bare-bones, introductory book, a level or so down in coverage from what we really want here, but which some might find a helpful preliminary read:

9. Paul Halmos, *Naive Set Theory** (1960: republished by Martino Fine Books, 2011). The purpose of this famous book, Halmos says in his Preface, is "to tell the beginning student ... the basic set-theoretic facts of life, and to do so with the minimum of philosophical discourse and logical formalism". He proceeds pretty naively in the second sense we identified in §2.1(b) .True, he tells us about some official axioms as he goes along, but he doesn't explore the development of set theory inside a resulting formal theory. This is informally written in an unusually conversational style for a maths book, concentrating on the motivation for various concepts and constructions. Some might warm to this classic (though ignore the remarks in the Preface about set theory being 'pretty trivial stuff'!).

Next, here are four introductory books at the right sort of level, listed in order of publication; each has many things to recommend it to beginners. Browse through to see which might suit your interests:

10. D. van Dalen, H.C. Doets and H. de Swart, *Sets: Naive, Axiomatic and Applied* (Pergamon, 1978). The first chapter covers the sort of elementary (semi)-naive set theory that any mathematician needs to know, up to an account of cardinal numbers, and then takes a first look at the paradox-avoiding ZF axiomatization. This is very attractively and illuminatingly done. (Or at least, the conceptual presentation is attractive – sadly, and a sign of its time of publication, the book seems to have been photo-typeset from original pages produced on electric typewriter, and the result is visually not attractive at all.)

    The second chapter carries on the presentation axiomatic set theory, with a lot about ordinals, and getting as far as talking about higher infinities, measurable cardinals and the like. The final chapter considers some applications of various set theoretic notions and principles. Well worth seeking out, if you don't find the typography off-putting.

11. Karel Hrbacek and Thomas Jech, *Introduction to Set Theory* (Marcel Dekker, 3rd edition 1999). This eventually goes a bit further than Enderton or Goldrei (more so in the 3rd edition than earlier ones), and you could – on a first reading – skip some of the later material. Though do look at the final chapter which gives a remarkably accessible glimpse ahead towards large cardinal axioms and independence proofs. Recommended if you want to consolidate your understanding by reading a second presentation of the basics and want then to push on just a bit.

    Jech is a major author on set theory whom we'll encounter again, and Hrbacek once won a AMA prize for maths writing. So, unsurprisingly, this is a very nicely put together book.

12. Keith Devlin, *The Joy of Sets* (Springer, 1979: 2nd edn. 1993). The opening chapters of this book are remarkably lucid and attractively written. The opening chapter explores 'naive' ideas about sets and some set-theoretic constructions, and the next chapter introduces axioms for ZFC pretty gently (indeed, non-mathematicians could particularly like Chs 1 and 2, omitting §2.6). Things then speed up a bit, and by the end of Ch. 3 – some 100 pages into the book – we are pretty much up to the coverage of Goldrei's much longer first six chapters, though Goldrei says more about (re)constructing classical maths in set theory. Some will prefer Devlin's fast-track version. (The rest of the book then covers non-introductory topics in set theory, of the kind we take up again in Part III.)

13. Judith Roitman, *Introduction to Modern Set Theory*** (Wiley, 1990: a 2011 version is available at tinyurl.com/roitmanset. This relatively short, and very engagingly written, book manages to cover quite a bit of ground – we've reached the constructible universe by p. 90 of the downloadable pdf version, and there's even room for a concluding chapter on 'Semi-advanced set theory' which says something about large cardinals and infinite combinatorics. A few quibbles aside, this could make excellent revision material as Roitman is particularly good at highlighting key ideas without getting bogged down in too many details.

Those four books all aim to cover the basics in some detail. The next two books are much shorter, and are differently focused.

14. A. Shen and N. K. Vereshchagin, *Basic Set Theory* (American Mathematical Society, 2002). This is just over 100 pages, and mostly about ordinals. But it is very readable, with 151 'Problems' as you go along to test your understanding. Potentially *very* helpful by way of revision/consolidation.

15. Ernest Schimmerling, *A Course on Set Theory* (CUP, 2011) is perhaps slightly mistitled, if ‘course’ suggests a comprehensive treatment. This is just 160 pages long, starting off with a brisk introduction to ZFC, ordinals, and cardinals. But then the author explores applications of set theory to other areas of mathematics such as topology, analysis, and combinatorics, in a way that will be particularly interesting to mathematicians. An engaging supplementary read at this level.

Applications of set theory to mathematics are also highlighted in a book in the LMS Student Text series which is worth mentioning here:

16. Krzysztof Ciesielski, *Set Theory for the Working Mathematician* (CUP, 1997). This eventually touches on advanced topics in the set theory. But the earlier chapters introduce some basic set theory, which is then put to work in e.g. constructing some strange real functions. So this might well appeal to mathematicians who know some analysis, who could indeed tackle Chs 6 to 8 on the basis of other introductions.

7.4 Further conceptual reflection on set theories

(a) Michael Potter’s *Set Theory and Its Philosophy* must be the starting point for philosophical reflections about set theory. In particular, he gives a good account of how our standard set theory emerges from a certain hierarchical conception of the universe of sets as built up in stages. There is also now an excellent more recent exploration of the conceptual basis of set theory in

17. Luca Incurvati, *Conceptions of Set and the Foundations of Mathematics* (CUP, 2020). Incurvati gives more by way of a careful defence of the hierarchical conception of sets and also an unusually sympathetic critique of some rival conceptions and the set theories which they motivate. Knowledgeable and readable.

Rather differently, if you haven’t tackled their book in working on model theory, you will want to look at

18. Tim Button and Sean Walsh’s *Philosophy and Model Theory** (OUP, 2018). Now see especially §1.B (on first-order vs second-order ZFC), Ch. 8 (on models of set theory), and perhaps Ch. 11 (more on Scott-Potter set theory).

(b) I will leave further philosophical commentary until the Part III chapter on more advanced set theory – except that I will here mention a short piece by Penelope Maddy, which takes us right back to our starting point when we introduced set theory as giving us a ‘foundation’ for real analysis. But what does that really mean? Maddy starts by noting “It’s more or less standard orthodoxy these days that set theory ... provides a foundation for classical mathematics. Oddly enough, it’s less clear what ‘providing a foundation’ comes to.” Her opening pages then give a particularly clear and crisp account of what might be meant by talk of foundations in this context. It is *very* well worth reading for orientation:

19. Penelope Maddy, 'Set-theoretic foundations', in A. Caicedo et al., eds., *Foundations of Mathematics* (AMS, 2017), available at tinyurl.com/maddy-found See §1 in particular.

7.5 A little more history

As already shown in the recommended book by Fraenkel, Bar-Hillel and Levy, the history of set theory is a long and tangled story, fascinating in its own right and conceptually illuminating too. José Ferreirós has an impressive book *Labyrinth of Thought: A History of Set Theory and its Role in Modern Mathematics* (Birkhäuser 1999). But that's more than most readers are likely to want. But you will find some of the headlines here, worth chasing up especially if you didn't read the book by Fraenkel et al.:

20. José Ferreirós, 'The Early Development of Set Theory', *The Stanford Encyclopaedia of Philosophy*, available at [sep-devset].

This article has references to many more articles, like Kanimori's fine piece on 'The mathematical development of set theory from Cantor to Cohen'. But you might to need to be on top of rather more set theory before getting to grips with *that*.

7.6 Postscript: Other treatments?

What else is there? A classic introduction is given by Patrick Suppes, *Axiomatic Set Theory** (vast Nostrand 1960, republished by Dover 1972). Clear and straightforward as far as it goes: but there are better alternatives now. There is also a classic book by Azriel Levy with the inviting title *Basic Set Theory** (Springer 1979, republished by Dover 2002). However, while this is still 'basic' in the sense of not dealing with topics like forcing, this *is* quite an advanced-level treatment of the set-theoretic fundamentals. So let's return to it in Part III.

András Hajnal and Peter Hamburger have a book *Set Theory* (CUP, 1999) which is also in the LMS Student Text series. They nicely bring out how much of the basic theory of cardinals, ordinals, and transfinite recursion can be developed in a semi-informal way, before introducing a full-fledged axiomatized set theory. But I think Enderton or van Dalen et al. do this better. The second part of this book is on more advanced topics in combinatorial set theory.

George Tourlakis's *Lectures in Logic and Set Theory, Volume 2: Set Theory* (CUP, 2003) has been recommended to me a number of times. Although this is the second of two volumes, it is a stand-alone text. Indeed Tourlakis goes as far as giving a 100 page outline of the logic covered in the first volume as the long opening chapter in this volume. Assuming you have already studied FOL, you can initially skip this chapter, consulting if/when needed. That still leaves over 400 pages on basic set theory, with long chapters on the usual axioms, on the Axiom of Choice, on the natural numbers, on order and ordinals, and on cardinality. (The final chapter on forcing should be omitted at this stage, and strikes me as less clear than what precedes it.)

As the title suggests, Tourlakis aims to retain something of the relaxed style of the lecture room, complete with occasional asides and digressions. And as the page length suggests, the pace is quite gentle and expansive, with room to pause over questions of conceptual motivation etc. However, there is a certain quite excessive and unnecessary formalism that many (most?) will find off-putting, and which slows things right down. Simple constructions and results therefore take a *very* long time to arrive. We don't meet the von Neumann ordinals for three hundred pages, and we don't get to Cantor's theorem on the uncountability of $\mathcal{P}(\omega)$ until p. 455! So while this book might be worth dipping into for some of the motivational explanations, I can't myself recommend it overall.

Finally, I'll mention another more recent text from the same publisher, Daniel W. Cunningham's *Set Theory: A First Course* (CUP, 2016). But this doesn't strike me as a particularly friendly introduction. As the book progresses, it turns into pages of old-school Definition/Lemma/Theorem/Proof with rather too little commentary; key ideas seem often to be introduced in a phrase, without much discursive explanation. Readers who care about the logical niceties will also raise their eyebrows at the author's over-causal way with use and mention, or e.g. the too-typically hopeless passage about replacing variables with values on p. 14. And this isn't just being pernickety: what exactly are we to make of the claim on p. 31 that a class is "any collection of the form $\{x\colon \varphi(x)\}$"? Not recommended to logicians of a sensitive disposition!