1

Introduction

1.1 Primitive notions

The word 'definition' has come to have a dangerously reassuring sound, owing no doubt to its frequent occurrence in logical and mathematical writings.

– Willard van Orman Quine[1]

1.1.1 Definitions – avoiding circularity

Some elementary observations have profound corollaries. Here is an example. Suppose finitely many points are distributed in some space and each point is joined to a number of other points by arrows, forming a complex directed network. We choose a point at random and trace a path, following the direction of the arrows, spoilt for choice at each turn. No matter how skillfully we traverse the network, and no matter how large the network is, we are forced at some stage to return to a point we have already visited. Every road eventually becomes part of a loop, in fact many loops.

A dictionary is a familiar example of such a network. Represent each word by a point and connect it via outwardly pointing arrows to each of the words used in its definition. We see that a dictionary is a dense minefield of circular definitions. In practice it is desirable to make these loops as large as possible, but this is a tactic knowingly founded on denial. The union of all such loops forms the core of the artificial language world of the dictionary, every word therein definable in terms of the loop members. So not only does this language core contain circular definitions, it is made of them! (In an appropriate tribute to the problem, the dictionary on my desk has some very short loops in the region containing the words *meaning* and *sense*.)

Any idealistic attempt to populate a dictionary with a finite number of words defined finitely and exclusively in terms of one another leads either to

[1] 'Two dogmas of Empiricism', in Quine [**172**].

the humiliation and dissatisfaction of circularity or, what is the better of the two options, to the nomination of some notions as 'primitive'. These primitive elements by nature cannot be defined in terms of other words, yet they form the building blocks of all other entries in the dictionary.

No dictionary of natural language is organized in this way, and nor should we expect it to be. Lexicographers skillfully try to evade circularity by using certain observable phenomena, the sensory and emotional meat of daily existence, and abstractions thereof, as their primitive reference points. As the observables are given fairly crude finitistic word descriptions such efforts will always fail for the reason just given, but the points at which a definition struggles to avoid circularity signpost those areas where the dictionary begins to step outside the boundaries of its original function; it has never pretended to be, and never will be, a self-contained account of the meaning of all things, and it is not a philosophical text. Perhaps it is not surprising that extremist philosophies adopting the position that 'all is language' tend to fall into a curious form of nihilism, turning all philosophical discussion into petty exercises in obscurantist creative writing.

By a long process of generalization and analogy, natural language reaches for high concepts far removed from the concrete stuff that lies at its foothills. This process is intuitive – new concepts are described as and when they are needed. Of course our dictionary example is just a convenient somewhat artificial one – however, the same finitistic obstruction applies to any attempt to chase meaning. How is it that a finite universe can harbour sense?

Most accounts of mathematics rely on an intuitive base of instantly recognizable Platonic objects, for example the classical number systems and various geometric notions. That these are 'instantly recognizable' (however this may be interpreted) is perhaps surprising given that none of them exactly match the comparatively crude real worldly things that helped to inspire them; but this act of abstraction, an effortless ability to simplify by assigning an ideal object to a large number of perceived objects is, to our great benefit, the way we have evolved to 'make sense of what we sense' – we simply couldn't think without it. However if we want to describe these familiar structures without deferring to vague and unreliable intuitions we have no choice but to embrace the 'primitive element' approach. An unworkable alternative which is only marginally better than circularity is to admit an infinite regress of definitions – a bottomless pit where each notion is defined in terms of lower notions.

Stretching the analogy to breaking point, by fixing a set of primitive elements we create an artificial platform spanning a cross-section of the bottomless pit. From this base we can look up (what properties follow from our basic assumptions?) and we can look down (what, at a deeper level, is capable of describing all of the properties we have chosen?). In a sense the two directions represent different aspects of mathematical versus logical enquiry, although we mustn't take this naive picture too seriously.

The crucial question is: what should we take to be the primitive elements of all of mathematics? As nineteenth century ideas moved into the twentieth, it was generally agreed that the primitive notions should be axioms governing

Cambridge University Press
978-1-107-05831-6 - The Logic of Infinity
Barnaby Sheppard
Excerpt
More information

'aggregates' together with a means of describing or isolating objects within such aggregates (sets/classes and the idea of membership). This was the birth of set theory.

Remarks

1. There is a natural way to modify directed networks which relieves them of all loops. The basic idea is in essence to repeatedly contract each minimal loop to a single vertex. However, in order not to lose too much information, we have to be careful not to combine two adjacent loops in one operation. To this end we repeat the following simple process. Suppose our directed network has labelled vertices. To each loop of minimal size we associate a new vertex whose label is the set of all labels of the vertices of the loop it represents. We then delete all vertices belonging to these minimal loops (which means we also delete any edges that are deprived of one or two end vertices). Next we connect the newly created vertices to the rest of the network in the natural way: the outgoing edges are the outgoing edges of the original loop vertices (in the case where minimal cycles are joined to one another by an edge, or share edges or vertices, some of these edges will be joined to one of the new vertices); and the incoming edges are the incoming edges of the original loop vertices. We then repeat the construction on the new network. Eventually all loops will vanish and we will be left with a tree with some set-labelled vertices. In extreme cases we might end up with just one vertex with a complicated label of nested sets. Performing this construction on the network associated with a good sized dictionary would create an amusing toy model for what might be called a lexicographic universe; one gets something of a burning curiosity to see this process in action (at least to determine the first few stages – what are the initial few minimal cycles?).

2. The common practice of mathematical abstraction is this: Take some intuitive object. Determine some properties of this object (express these properties in set theoretic/algebraic terms). Then consider the class of all objects satisfying these properties. Of course this class will include a model of the intuitive object you started with, but it might include a host of other objects too, possibly some surprises. The extremal objects in this set might reveal something new about the original object. To what extent can the original object be isolated within this larger class of models; is there a further set of properties which completely characterizes it among its cousins? This is a simple idea, but it is remarkable how powerful it can be. We will see some examples of this later.

A chain of definitions must be an infinite regress, an eventual cycle or it must terminate in an undefined primitive notion. Rejecting the first two options we must choose the primitive notions of the language of mathematics. Can all of mathematics be described in terms of a theory of aggregates?

1.1.2 Intuition and its dangers

Fed by constant sensory bombardment from the outset, and through total immersion in our environment, we develop an intuitive, albeit often misguided, grasp of some fundamental notions. Exposed to a tidal wave of examples we build soft conceptual structures on softer foundations and form a picture of the physical and logical universe that is childishly simplistic yet, provided we don't question our ideas too closely, free from alarms and surprises. Such a thin grasp of the world about us is adequate for the vast majority of human activities. Once the sensory parameters are shifted, perhaps by space (the weirdness of the subatomic world, or at the other extreme of magnification, clusters of galaxies mere specks of dust) or by time (superslow motion, all living creatures now at a geological pace, or superfast where we watch continents drift and observe the evolution of species) we are immediately alienated, the intuitions moulded by our mesoscopic upbringing of little use.

Without the ability to quickly construct a rough understanding of some notions we would be helpless. This basic intelligence is essential in order to gain a foothold in any subject; whether it is by forming a vague picture or by drawing a rough analogy, we need some foundation on which to hang our thoughts. Later enquiry may change the initial picture, sometimes drastically, but nevertheless the ability to concoct an immediate (fuzzy) impression when presented with a new concept is crucial.

The reality is that we, short-lived impatient biological beasts, work on unfamiliar complicated stimuli from the 'top down', dissecting pieces as required. We tend to stop further scrutiny when we feel satisfied that the new entity has been explained in terms of ideas we think we already understand. What qualifies as 'satisfaction' is a matter of taste and experience. If, on the other hand, we are working from the bottom up, building the foundations of a subject, we are most content when the primitive assumptions are few in number, self-evident and consistent to the best of our knowledge. A system of assumptions is consistent if from it we cannot deduce both a statement and its negation. The qualification 'to the best of our knowledge' may seem like an unsatisfactory appeal to the arbitrary limits of human mathematical ability, however, as Gödel famously demonstrated (see Appendix C), there are many systems for which it is impossible to prove consistency without appealing to principles outside the system.

It can be very rewarding to analyze notions that we take for granted. Such analysis often uncovers surprises and we realize what little we understood in the first place. Sometimes what we pretend to be obvious is far from clear, and through the microscope we perceive magnificent strangeness. It is partly the *apparent* surreal quality of some results in mathematics which initially draws many people to it. These exotic ideas, when they first appear, seem to live far from the utilitarian regions of science, yet they have a curious habit of barging their way to the forefront of physics.

Many are drawn to mathematics simply for the pleasure of total immersion in a consistent but otherworldly universe. Stanisław Ulam acknowledges this escapist aspect of the subject in his autobiography *Adventures of a Mathematician*.[1]

Eventually, in the mature stages of such critical investigation, we come to regard the newly uncovered world as normal. All too often, mimicking the progress of science itself, we must abandon or more often demote to crude approximations and special cases our preconceptions and gradually develop a new intuition based on a more refined point of view. The ability, and more importantly, the willingness, to constantly question one's beliefs and understanding seems essential for any sort of intellectual progress.

Remarks

1. On reflection it is quite alarming how many of our intuitions are based on pure guesswork, or even prejudice, founded on almost nothing at all. Even the most elementary facts about the physical world can be counterintuitive at first; here I am especially thinking of the pre-Galilean misconception that heavier bodies fall faster than lighter bodies. Of course a well designed experiment will swiftly correct this (simply ensure that both objects experience the same air resistance). But there is no need to climb the steps of the Tower of Pisa; a moment's thought can help to shatter the false belief: imagine two bodies falling and joining one another, or the reverse, a single body breaking up into two smaller pieces as it falls. The notion of a body momentarily changing its rate of acceleration as it splits or coalesces is clearly unnatural. In particular consider the moment when two large bodies touch at a point; are we to believe that the united body will suddenly start to plummet even faster owing to this tiny point of contact, and then equally suddenly decelerate as the contact is broken? Alternatively consider connecting a light and heavy body together with a length of string. Is the lighter body somehow expected to reduce the speed of descent of the heavier body, as if it were a parachute, even though their combined weight exceeds the weight of each component part? From these considerations alone one would conjecture that all bodies experiencing the same air resistance, and in particular all bodies in a vacuum, fall at the same rate, and experiment verifies this. Finding the right way to think about something is the difficult part.

[1] Ulam **[217]**, p. 120. Ulam compares some mathematical practices with drug use, or with absorption in a game of chess, which some mathematicians embrace as a means of avoiding the events of a world from which they wish to escape!

2. From a certain extremely reductionist viewpoint, what Eugene Wigner famously described as 'the unreasonable effectiveness of mathematics in the natural sciences'[1] is not quite as unreasonable as it might seem at first glance. That is to say, if one imagines a bottom-up description of the physical Universe, starting with some primitive elements governed by some simple relations, then there is no surprise at all that the emergent macroscopic world that grows out of this will be a mathematical one. From this point of view every mathematical result, no matter how esoteric, that is describable in the underlying 'logic' of the Universe has the potential to describe something physical. There is perhaps still some room for mystery in that many of the macroscopic mathematical relations in the Universe are so surprisingly tractable (inverse square laws, beautiful symmetry and so forth) and so elegantly related to one another. It is remarkable that the symbolically expressed imaginative fantasies of mathematics can have such concrete applications. If the Universe is not mathematical, what else could it possibly be?[2]

Intuition is reliable only in the limited environment in which it has evolved. Unable to abandon its prejudices completely, we must constantly question what appears to be obvious, often revealing conceptual problems and hidden paradoxes. One intuitive notion which is ultimately paradoxical is that of arbitrary collections.

1.1.3 Arbitrary collections

One notion that we take for granted, to the point of blissful ignorance, has been mentioned already: finiteness. To give a sound definition of finiteness is a surprisingly sticky problem, but it is not intractable. We shall come to its formal definition later, but for now we will have to settle for the intuitive idea and take on trust that it can be formalized.

Finiteness is a property of certain arbitrary collections of objects. Without reference to the sophisticated notion of 'number' or 'counting' how might we compare such aggregates? We shall give an answer to this shortly. The objects we have in mind, in this naive introduction at least, are free to be anything we care to imagine, physical objects or, most often, abstract notions. We deliberately delay mentioning 'numbers' for reasons which will soon become clear; we will model numbers as certain specific sets of objects. By an *element* of a given collection we mean one of the objects in the collection.

[1] *Communications on Pure and Applied Mathematics*, **13**, 1–14 (1960). A copy of the article is easily found online.

[2] There is of course a huge literature on this subject. For an interesting take on the relationship between mathematics and the empirical sciences from a *fictionalist* point of view, see Leng [**137**].

It would be liberating to have at our disposal a list of synonyms for the word 'collection'. However, as is often the case in mathematics, the most attractive possible labels have very particular meanings attached to them already ('group', 'category', 'set' and 'class', for example, are all taken). One can be fairly confident that any alternative to 'aggregate' has acquired a technical definition in some branch of mathematics. This can cause interpretative difficulties in certain expositions, but in practice context dictates meaning.

We will be seeing much talk of 'sets' and 'classes'. Roughly speaking, a set is a 'well-behaved' collection formed from other sets according to certain rules. Each finite collection is a set, as are many infinite collections, including the basic number systems we are about to introduce, with the exception of the collection of all ordinal numbers and the collection of all cardinal numbers.[1] Most set theories include an axiom stating that a collection modelling the natural numbers is a set, or one can prove this from more general axioms, and further axioms give rise to infinitely many other infinite sets. A set theory with no infinite sets is obviously possible, but such a setting is clearly not the right environment for the subject of this book, and to adopt such an unnecessarily restrictive theory would also mean having to reject vast amounts of beautiful (and very useful) mathematics.

We can think of classes as collections of objects which share a common property. All sets are classes, but there are classes which are not sets. These sprawling monsters, called *proper classes*, are banished from the safer world of sets because their presence gives rise to unpleasant consequences: the paradoxes. If we were to allow proper classes to be sets then the theory would face an internal contradiction and come crashing down, collapsing under the weight of its own ambition. In fact, because of this delicate divide between sets and proper classes, the notion of 'set' is rather subtle and deserves more discussion; 'set' is treated as a primitive term in set theory, just as 'line' and 'point' are primitive terms in axiomatic geometry. Exactly what is meant by a set – that is, its principal interpretation; how it differs from the intuitive notion of a collection – is very difficult to answer briefly, and it is this difficulty that makes the subject of set theory peculiar at its outset. Most other branches of mathematics can define their central concepts early in the development and with relative ease; group theory begins (along with motivating examples) with the definition of a group; topology begins (again with motivating examples) with the definition of a topological space; measure theory soon gives the definition of a measure and so on. Set theory, on the other hand, starts with 'set variables', 'set' being a primitive term loaded with intuitive interpretation, the *naive* theory of which turns out to be fundamentally flawed. In more explicit terms, which will be clarified later, most mathematical theories grow from a plethora of examples, the common features of which dictate the axioms of the underlying theory (each example comprising a *model* of the theory). The historical development of set theory did not fully conform to this familiar pattern; the axioms came before

[1]When such statements are made I am tacitly making reference to the set theory known as Zermelo–Fraenkel set theory (ZF), which we will meet later, and which is the main theory considered in this book.

all genuine models, only the successes and failures of naive set theory serving as a guide. Every abstract theory begins, like set theory, with a set of constants together with variables, relation symbols and axioms (see Section 2.3), but it is quite rare to establish such a system without having a collection of models, even vague intuitive ones, in mind.

It was the twentieth century clarification of the notions of theory and model which was to bring the question 'what is a set?' into sharper focus. Any meaningful answer to the question must be postponed until an explicit model of the theory has been exhibited, and once this is done the answer is fairly banal (a set, in a given model, is simply a member of the model).[1] Prior to the construction of concrete models of set theory, at the theory-building stage, the question should instead be 'what properties do we want sets to have?'. Our answer to this question is to present a short list of properties (the non-logical axioms) which correspond to the properties we imagine intuitive sets – mathematical aggregates – to have, based on our prior experience with mathematics.

Much of the early development of set theory concerns itself with the task of determining when a class is and when it is not a set, or at least in which ways sets can be combined to form other sets. I will use 'set' and 'class' quite casually in this introduction with the hope that the reader understands that not all classes are sets, and that the scope of some notions such as size (i.e. 'number of elements') is restricted to sets and has no meaning for proper classes. At the centre of set theory is the powerful idea of regarding as a single entity any collection of objects.

We exhibit small finite sets by listing their elements inside braces, e.g. $\{a, b, 4, 7\}$. Larger finite sets with an easily discernible pattern are listed in the informal manner $\{0, 2, 4, \ldots, 96, 98, 100\}$, and a similar conceit can be used for some infinite sets. We can be more precise and write $\{n : n \text{ is an even integer no less than 0 and no greater than 100}\}$ for the latter set. More generally we use $\{x : \phi\}$, meaning 'the class of all x satisfying the property ϕ', where ϕ is some statement. A casual use of this simple device leads to profound difficulties which will be discussed later.

In the presentation of a set the order in which its elements are displayed is of no consequence, and any repetition of an element is redundant. It is important to understand the distinction between a and $\{a\}$; the latter is a set comprising a single element, namely the set a, while the former is the set a itself, which may have many elements.

[1] Analogously it is meaningless to ask, when presented with the theory of groups (see Subsection 2.4.1), 'what is an element of a group?'. This is a question for an individual *model* of group theory, i.e. a particular group, to answer, and the answer, depending on the model, could be that it is a permutation, a rotation or translation in space, an invertible matrix, an integer, a function, or indeed any other set theoretically describable mathematical object. In axiomatic theories of geometry we postulate the properties of objects suggestively called 'points' and 'lines', but in describing models of the theory one needn't take these elements to be the familiar points and lines as they are understood in the conventional sense. It can be advantageous to depart from the usual intuitive objects which motivated the axioms. Indeed, this kind of imaginative departure was the conceptual leap required to produce concrete models of non-Euclidean geometries (see Subsection 2.4.2).

REMARKS

1. Jumping ahead of ourselves a little, we can give a precise set theoretic definition of a natural number, and so we can collect all such numbers together to form a class. Do we want this collection to be a set, or should it remain a proper class? After a century of very close scrutiny no contradiction has been shown to follow from the assumption that the class of all natural numbers is a set, so it would seem to be needlessly restrictive not to allow it to be a set. We shall come back to this later.

2. It should be stressed that ZF is just one among many approaches to set theory. Some alternatives closely resemble ZF, others are wildly different. Indeed, there are approaches to the foundations of mathematics that are not 'set theoretic' at all. However, I think ZF is perhaps the most natural place to begin; and once its ideas have been absorbed it is easy to consider variations and deviations from the norm.

It is surprisingly difficult to rigorously define finiteness. In studying this problem we begin to ask what it means for two sets to have the same number of objects. What is 'number'?

1.1.4 Equipollence

What do we mean when we say that two sets are of the same 'size'? If these sets are finite then the notion is something we seem to be able to cope with at a very early stage in our cognitive development. We might proceed as follows. Let us assume that we have been presented with two finite sets of objects. We then simultaneously take one object from each set, continuing to remove pairs until at least one of the sets has been exhausted. If at this terminal stage there still remain objects in front of us then the pair of sets we started with clearly had a different 'number of elements'. Essentially we have just paired off objects, that is, we say that two sets are of the same size if we can pair off the elements of the first set with the elements of the second set in a one-to-one fashion.

As this is such a crucial concept we need to introduce a little terminology in order to make the notion precise.[1] A *function* between two sets A and B is a mapping associating each element of A with some element of B. If we call the function f then the statement that f maps the element a of A to the element b

[1] At this stage we are denoting arbitrary sets by upper case roman letters. Later we adopt the convention that sets are to be denoted by lower case letters and classes by upper case letters.

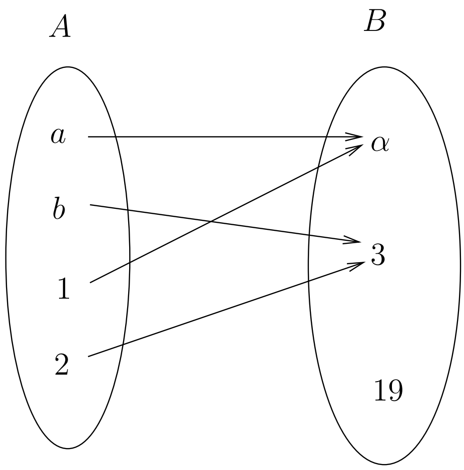

Figure 1.1 A simple-minded depiction of a function mapping the finite set $A = \{a, b, 1, 2\}$ to the finite set $B = \{\alpha, 3, 19\}$, the details of the mapping illustrated by a collection of arrows indicating which elements of B are associated with each element of A.

of B is denoted $f(a) = b$. One might regard, and in fact later *define*, a function f from A to B as a collection $\mathbf{F}$ of ordered pairs $(x, f(x))$ where each element of the set A appears as the first entry of exactly one of the ordered pairs in $\mathbf{F}$ and where $f(x)$ is in B for all x in A. The notation $f : A \to B$ is used as a shorthand for the statement 'f is a function from A to B'. We call A the *domain* of f and B the *codomain* of f. The notion of function is one of the most basic and widespread in mathematics.

It is commonplace to picture a function between arbitrary finite sets of objects as a collection of arrows joining elements of the domain to their targets in the codomain (as in Figure 1.1). If the function maps x to y and we wish to suppress the name of the function itself we often use the notation $x \mapsto y$.

We need to consider certain types of function. A function $f : A \to B$ is *injective* (or *one-to-one*) if no two different elements of A are mapped to the same element of B, that is to say $f(x) = f(y)$ implies $x = y$. The function illustrated in Figure 1.1 is not injective because both elements a and 1 of A map to the same element of B, namely α. A function is *surjective* (or *onto*) if each element of B is the image of at least one element of A, that is for every b in B there exists an a in A such that $f(a) = b$. The function illustrated in Figure 1.1 is not surjective because no element of A maps to the element 19 of B. A function is *bijective*, or is a *bijection*, if it is both injective and surjective.

In the case when A and B are finite sets the intuitive statement that A and B have the same number of elements is perfectly captured by the rigorous statement that there exists a bijection $A \to B$. We extend this particular manifestation of number to all sets, finite or otherwise, and say that sets C and D have the same *cardinality* if there exists a bijection $C \to D$. We interpret this correspondence as meaning that C and D have the same 'number' of elements,