1

What is Mathematics About?

Michael Dummett

The two most abstract of the intellectual disciplines, philosophy and mathematics, give rise to the same perplexity: what are they *about*? The perplexity does not arise solely out of ignorance: even the practitioners of these subjects may find it difficult to answer the question. Mathematics presents itself as a science in the general sense in which history is a science, namely as a sector in the quest for truth. Even those least instructed in other sciences, however, have some general idea what it is that those sciences strive to establish the truth about. Historians aim at establishing the truth about what was done by and what happened to human beings in the past; more exactly, to human beings after they had invented writing. Physicists try to discover the general properties of matter under the widest variety of conditions; more generally, of matter and of what it propagates, such as light and heat. But what is it that mathematicians investigate?

An uninformative answer could be given by listing various types of mathematical object and mathematical structure: mathematicians study the properties of natural numbers, real numbers, ordinal numbers, groups, topological spaces, differential manifolds, lattices, *and the like*. Apart from the difficulty of explaining "and the like," such an answer is uninformative because it is given from within: one has to know some mathematics – even if, in some of the cases, only a little – if one is to understand the answer, whereas the sample answers to the questions concerning history and physics could be understood without knowing any history or physics.

Some maintain, nevertheless, that mathematics is a science like any other. The claim is unconvincing prima facie: what is immediately striking about mathematics is how *unlike* any other science it is. It is true that, in the more mathematicized sciences such as physics, there may be elaborate deductions from initial premisses, just as there are in mathematics; but they play a different role. In mathematics, their purpose is to establish theorems, that is, mathematical truths; in physics, they serve to elicit consequences of a theory, which can then be used to make predictions but also to test the theory. The word "theory" is used quite differently in mathematics and in the other sciences. In physics, biology, and so forth, it carried the connotation of a hypothesis; however well established a physical or biological theory, it always remains open to refutation or revision. In mathematics, there is no such connotation. We are all familiar with the idea of observations designed to test – to confirm or refute – the general theory of relativity; but we should be unable to conceive of observations designed to test number theory or group theory.

The most determined effort to represent mathematics as empirical in character was made by John Stuart Mill; but he achieved little more than to point out, what is in any case evident, that mathematics can be *applied* to empirical reality. That, indeed, is a salient feature of mathematics that any

Michael Dummett, What is Mathematics About? *Mathematics and Mind*. Ed. Alexander George. Oxford: Oxford University Press, 1994. 11–26.

philosophical account of it must explain; but it is not to be explained by characterizing mathematics as itself an empirical science. Our very vocabulary indicates the difference. We do not speak of "applying" a physical theory when we draw physical consequences from it, but only when we base some technological innovation upon it. Even someone who accepted all Mill's arguments would have no ground for regarding mathematics as a science like any other; it would still differ markedly from all others. For Mill, the axioms and definitions of mathematics are derived from very general facts apparent to untutored observation; but the theorems are still consequences drawn by deductive reasoning from those axioms and definitions, without further appeal to observation, let alone to refined observations made in artificially created conditions or with the help of sophisticated instruments. Moreover, as Frege pointed out, the mathematical notions whose application Mill was anxious to locate solely in physical reality have in fact far wider application. It is misleading to say that we encounter the natural numbers, for example, in the physical world; for, while physical situations may indeed need to be described by citing a natural number as the number of physical objects of some given kind, non-physical situations may equally need to be described by citing a natural number as the number of non-physical objects of some given kind, for instance as the number of different proofs of the fundamental theorem of algebra, or, indeed, of roots of an equation. The same holds good of sets. These notions are too general for us to locate them in any particular realm of reality; as Frege maintained, they apply within every sector of reality, and the laws governing them hold good, not only of what we find to exist, but of all of which we can frame intelligible thoughts.

If mathematics is not about some particular realm of empirical reality, what, then, *is* it about? Some have wished to maintain that it is indeed a science like any other, or, rather, differing from others only in that its subject-matter is a super-empirical realm of abstract entities, to which we have access by means of an intellectual faculty of intuition analogous to those sensory faculties by means of which we are aware of the physical realm. Whereas the empiricist view tied mathematics too closely to certain of its applications, this view, generally labeled "platonist," separates it too widely from them: it leaves it unintelligible how the denizens of this atemporal, supra-sensible realm could have any connection with, or bearing upon, conditions in the temporal, sensible realm that we inhabit.

Like the empiricist view, the platonist one fails to do justice to the role of proof in mathematics. For, presumably, the supra-sensible realm is as much God's creation as is the sensible one; if so, conditions in it must be as contingent as in the latter. The continuum hypothesis, for example, might *happen* to hold, even though we can apprehend neither its truth nor anything in which its truth is implicit. That there may be mathematical facts that we shall be forever incapable of establishing is a possibility admitted by some mathematicians and philosophers of mathematics, though denied by others. When admitted, however, it is normally admitted on the ground that our inferential powers are limited: there may be consequences of our initial assumptions that we are unable to draw. If these are first-order consequences, we could "in principle" draw them, since they could be elicited by reasoning each step of which was simple; but the proofs might be too long and complex for us ever to be able to hit on them, or even follow them, in practice. If they are second-order consequences, we may be unable even in principle to see that they follow. But, if we take seriously the analogy between our supposed faculty of intuition and our perceptual faculties, there is no reason why there may not be mathematical facts that are in no sense consequences of anything of which we are aware. We may observe a physical object without either perceiving all its features or being able to deduce all of them from what we do perceive; if mathematical structures are merely the inhabitants of another realm of reality, apprehended by us in a manner analogous to our perception of physical objects, there is no reason why the same should not be true of them. There are indeed hypotheses and conjectures in mathematics, as there are in astronomy; but, while both kinds may be refuted by deducing consequences and proving them to be false, the mathematical ones cannot be established simply by showing their consequences to be true. In particular, we cannot argue that the truth of a hypothesis is the only thing that would explain that of one of its verified consequences; there is nothing in mathematics that could be described as inference to the best explanation. Above all, we do not seek, in order to refute or to confirm a hypothesis, a means of refining our intuitive faculties, as astronomers seek to improve their instruments. Rather, if we suppose the hypothesis true, we seek for a *proof* of it, and it remains a mere hypothesis, whose assertion would therefore be unwarranted, until we find one. True, we seek to make our methods of proof ever more explicit and precise. This is not analogous to the improvement of the instruments, however. Methods of proof serve to elicit consequences, not to yield a more extensive evidential base; if the hypothesis is to be established, this must be done, not by testing its consequences, but by exhibiting *it* as a consequence of what we already know. Platonism can no more explain these differences between mathematics and the natural sciences than empiricism can, for both go astray by claiming to discern too close an analogy between them.

A brilliant answer to our question, but one now generally discredited, was given by Gottlob Frege and sustained by Russell and Whitehead. It was, essentially, that mathematics is not about *anything in particular*: it consists, rather, of the systematic construction of complex deductive arguments. Deductive reasoning is capable of eliciting, from comparatively meager premisses and by routes far from immediately obvious, a wealth of often surprising consequences; in mathematics, such routes are explored and the means of deriving those consequences are stored for future use in the form of propositions. Mathematical theorems, on this account, embody deductive subroutines which, once discovered, can be repeatedly used in a variety of contexts.

This answer, generally called the "logicist" thesis, was brilliant because it simultaneously explains various puzzling features of mathematics. It explains its methodology, which involves no observation, but relies on deductive proof. It explains the exalted qualification it demands for an assertion: in other sciences, a high degree of probability ranks as sufficient ground for putting forward a statement as true, but, in mathematics, it must be incontrovertibly *proved*. It explains its generality; it explains our impression of the necessity of its truths; it explains why we are so perplexed to say what it is about. Above all, it explains why mathematics has such manifold applications, and what it is for it to be applied. It allows that mathematical statements are genuinely propositions, true or false, and hence accounts for what is manifestly so, that mathematicians may be interested in determining their truth-values regardless of the uses to which they may be put; at the same time, it explains the content of those propositions as depending on the possibility of applying them, and thus justifies Frege's dictum that it is applicability alone that raises arithmetic from the rank of a game to that of a science. By contrast, Wittgenstein's account of mathematics, which lays even greater stress on application, makes the existence of pure mathematicians a phenomenon for pathology. It will be my purpose in this essay to maintain that the logicist answer, if not the exact truth of the matter, is closer to the truth than any other than has been put forward.

The classic versions of logicism both ran aground on the problem of the existence of mathematical objects, those abstract entities of which mathematical theories, taken at face-value, treat, and, above all, of the elements of the fundamental mathematical domains; the domain of the natural numbers and that of the real numbers. The aim of representing a mathematical theory as a branch of logic is in tension with recognizing it as a theory concerning objects of any kind, as its normal formulation presents it as being: for we ordinarily think of logic as comprising a set of principles independent of what objects the universe may happen to contain. Frege nevertheless believed that the truth of number theory and of analysis demanded the existence of those objects with which, on the face of it, they are concerned; and so he had to justify the belief in their existence, while reconciling it with the purely logical character of arithmetical statements. In trying to achieve this, he ran into actual contradiction. Russell and Whitehead, greatly concerned with the need to avoid contradiction, tried to construct foundations for mathematics in accordance with the more natural conception of logic as independent of the existence of any particular objects: their classes are not genuine objects at all, but mere surrogates, statements about them being explained as a disguised means of talking about properties of objects, properties of such properties, properties of properties of *those* properties, and so on upwards. Frege had never given any good reason for insisting on the genuine existence of mathematical objects; perhaps the only plausible reason lies in the difficulties encountered by Russell and Whitehead in trying to dispense with them. The price they paid for doing so was that, in order to ensure the existence of sufficiently many of their object-surrogates, they had to make assumptions that could not be rated as logical, or even likely to be true. The Axiom of Infinity, saying that there are infinitely many concrete objects, was needed to make sure that the natural numbers did not terminate; the Axiom of Reducibility, saying that there are sufficiently many properties of things of a given type definable without speaking of all properties of those things, was needed to guarantee the completeness of the real-number system.

More recently, Hartry Field has advanced what may be seen as a modification of the logicist thesis. Frege argued that the application of a mathematical theory, outside mathematics or within it, requires, to warrant it, a stronger claim than the consistency of the theory being applied. Suppose that a theorem in one mathematical theory T is proved by appeal to another, auxiliary, theory S. It is then not enough, Frege reasoned, to know that if the theory T is free from contradiction, then so is the combination of T and S: for that would warrant us in claiming no more than that we shall not involve ourselves in contradiction if we accept the theorem, whereas we wanted to be in a position to *assert* that theorem; and for that, Frege held, we must know the auxiliary theory S to be *true*. Field argues that we need claim nothing so strong as truth on behalf of a theory in order to warrant its applications. It we want to show that a mathematical theory S can be legitimately invoked as an auxiliary to some other theory T (which may be a scientific theory or another mathematical one), we need only claim something intermediate between the logical truth of S and its consistency relative to T, namely that the conjunction of T and S is a conservative extension of T. This means that anything expressible in the language of T that could be proved from T together with S could already be proved – perhaps at greater length and with greater difficulty – in the theory T on its own.

Field, too, is concerned with the existence of mathematical objects. He agrees with Frege, as against Russell and Whitehead, that the truth of a mathematical theory demands the existence of the mathematical objects of which it purports to treat: that is his reason for denying the truth of the theory, since he disbelieves in the existence of any such objects.

To give substance to the claim that, when the theory S is added to the theory T, it yields a conservative extension of T, we must be able to formulate T without reference to whatever objects of the theory S are regarded as objectionable; achieving such reformulations is the major part of Field's program. His motivation for seeking to explain the applications of mathematics without recognizing the existence of mathematical objects lies in his general disbelief in abstract objects of any kind. It is on this that criticism has centered: can he really formulate scientific theories without appeal to abstract objects?

For Frege, on the other hand, the error that blocked any reasonable philosophy of mathematics was the failure to recognize that abstract objects may be quite as objective as concrete ones (in his terminology, non-actual objects as actual ones). He characterized abstract objects much in the way that philosophers are disposed to do today, namely as objects lacking causal powers; by "objective" he meant something that is neither a content of consciousness nor created by any mental process. It is a common complaint about abstract objects that, since they have no causal powers, they cannot explain anything, and that the world would appear just the same to us if they did not exist: we can therefore have no ground to believe in their existence. For Frege, such a complaint would reveal a crude misunderstanding. He gave as an example of an object that is abstract but perfectly objective the equator. If you tried to explain to someone who had never heard of it what the equator was, you would certainly have to convey to him that it cannot be seen, that you cannot trip over it, and that you feel nothing when you cross it. If he then objected that everything would be exactly the same if there were no such thing as the equator, and that therefore we can have no reason for supposing it to exist, it would be clear that he had still not understood what sort of object we take the equator to be. What has to be done is to explain to him how the term "the equator" is used in whole sentences: how it is to be determined whether or not someone has crossed the equator, whether some natural feature lies on it or to the north or south of it, and so on. That is all that can be done, and all that needs to be done: if he still persists in objecting, there is nothing we can do but pity him for being in the grip of a misleading picture.

Thus reference to an abstract object is to be understood only by grasping the content of sentences involving such reference, and it is only by specifying the truth-conditions of such sentences that it can be explained what such an object is: it is only in the course of saying something intelligible about an object that we make genuine reference to it. This, indeed, holds good for all objects, concrete or abstract; but, because of their failure to appreciate the dependence of reference upon the context of a proposition, philosophers are tempted to dismiss the object referred to as mythological only when it is abstract, since sentences involving reference to concrete objects include those in which they are indicated by means of demonstrative terms, which is to say that concrete objects can be encountered. For Frege, however, to treat mathematical objects such as numbers as fictitious because abstract is to commit as crude a blunder as to do the same for the equator, one which springs from the same misunderstanding about what referring to an object involves.

In this, Frege was surely right. He did not, however, take his *general* defense of the existence of abstract objects as dispensing us from any work in particular cases, but only as pointing to the kind of work that needed to be done. In each case, we have to specify the truth-conditions of sentences containing terms for objects of the kind in question; in those with which he was concerned, for natural numbers, or cardinal numbers in general, and for real numbers. And this was, for him, a highly problematic task, but one that he believed he could solve. The first lesson of the contradiction was that he was woefully mistaken in that belief.

If we reject Field's all-encompassing nominalism, his program takes on a different aspect. Much of the criticism directed at it falls away, once the task is no longer that of avoiding reference to all abstract objects; it continues to be of interest because it focuses on the problem that defeated both Frege and Russell, of either justifying or explaining away reference to specifically mathematical objects, and that remains a problem even after the general objective of eliminating all reference to abstract objects has been discarded. Field's program then becomes a new strategy for resolving the problem of mathematical objects. Nevertheless, Field envisages the justification of his conservative extension thesis as being accomplished only piecemeal. For each mathematical theory, and each theory to which it is applied, the demonstration is to be carried out specifically for those two theories; no presumption is created by the successful execution of the program in one case that it will work in others. If, for example, it is shown that real analysis yields a conservative extension when adjoined to Newtonian mechanics, real analysis will not have received a general justification as a mathematical theory, but only in application to Newtonian mechanics. Now suppose that some millionaire is converted to Field's philosophy of mathematics and endows an institute to carry out Field's program for all scientific theories and all mathematical theories which find application to them; and suppose that the institute is uniformly successful: it has so far examined every existing scientific theory, and every application of mathematics made within it, and has in each case succeeded in establishing Field's claim. Then it has still not established that claim for future applications of other parts of mathematics to existing theories, nor for applications of mathematics to scientific theories yet to be devised. Long before this stage, however, we should have become dissatisfied with the institute's work. For each mathematical theory, we should surely demand a guarantee that it would always yield a conservative extension when adjoined to any scientific theory, so that it would be justified once and for all; and we should also require an explanation why it demonstrably did yield a conservative extension when adjoined to every known scientific theory. Such an explanation would have to turn on the character of the mathematical theory itself, independently of the particular scientific theories to which it was applied; and it would presumably provide the sought-for guarantee. Without such an explanation, we could hardly suppose that we had reached the fundamental truth of the matter; for it could not very well be a mere fortunate coincidence that the theories devised by the mathematicians just happened to yield conservative extensions when adjoined to the theories developed by physicists and other scientists. It is difficult to think what such a general explanation could be, unless it was that mathematical theories, if not logically true in the strict sense, have some closely related property. Field's thesis is not a single one, but a bundle of numerous particular theses; and, as such, it lacks the generality that is required of an adequate account of the applicability of mathematical theories.

Once we have achieved the required reformulation of the theory T to which some mathematical theory S is to be applied, Field's strategy is to prove a representation theorem for T. To avoid unnecessary detail, I will illustrate this by Field's own preliminary example. Here T is (an adaptation of) Hilbert's axiomatization of Euclidean geometry, while S is the theory of real numbers. T is formulated without reference to numbers of any kind, but with variables ranging only over geometrical objects; it is based on axioms governing primitive predicates expressing properties of and relations between them. In Hilbert's original formulation, there were three sorts of variable, for points, lines (determined by any two distinct points), and planes (determined by any three non-collinear points); Field prefers to conceive it as using variables only over points. In this case, there will be a four-place predicate holding between points x, y, z, and w just in case the line segment xy is congruent to the line segment zw, and a three-place predicate saying that y lies between x and z on some line. Then a partial rendering of Hilbert's representation theorem states that there will be a binary function d from the points in any model of T into the non-negative real numbers such that

$$\mathrm{d}(x,y) = \mathrm{d}(z,w)$$

just in case xy is congruent to zw and

$$\mathrm{d}(x,z) = \mathrm{d}(x,y) + \mathrm{d}(y,z)$$

just in case y is between x and z. By laying down suitable conditions on the distance function d, we could prove a converse, namely that any structure on which was defined a function d satisfying those conditions could be converted into a model of T by explaining segment-congruence and betweenness in the manner just stated.

This helps to explain how real numbers can be used as an auxiliary device for proving results within this particular theory, namely the theory T of Euclidean geometry; it does not, of course, illuminate the uses of real numbers in other applications. Field remarks that the function d is unique only up to multiplication by a positive constant. This reflects the obvious fact that a quantity – here, a distance – does not by itself determine a real number alone, but only in conjunction with a unit. What uniquely determines a real number is a ratio between distances: if we replaced d by a function e of four arguments, giving the ratio of the distance between x and y to that between z and w (where z and w are distinct), we could reformulate the representation theorem so that e would be unique. (We should then require that xy be congruent to zw just in case $x = y$ and $z = w$ or $\mathrm{e}(x, y, z, w) = 1$, and that y be between x and z just in case $y = z$ or $\mathrm{e}(x, z, y, z) = \mathrm{e}(x, y, y, z) + 1$.) We need not do this, of course; it is enough to observe that if real numbers are uniquely determined by ratios between distances, it at once follows that there will be distance functions d obtained from one another by multiplication by positive real numbers. (Any such d will be obtainable from e by setting $\mathrm{d}(x, y) = \mathrm{e}(x, y, a, b)$ for suitable fixed distinct a and b.) Furthermore, real numbers correspond uniquely to ratios, not merely between distances, but between quantities of any one type. Hence, given an adequate analysis, such as that aimed at in measurement theory, of what, in general, constitutes a range of quantities, we have a hope of a general explanation of why there will be a unique mapping of pairs of objects that have these quantities on to the real numbers or on to some subset of them. From these sketchy remarks, it is possible to glimpse how such a general explanation might be made to yield a theorem of which a whole range of corollaries ensuring a representation by means of real numbers would be special cases. We should then have secured the desired generality for explaining the applications of real numbers on Fieldian lines. Such a theorem would encapsulate the general principle for applying the theory of real numbers.

This, however, would do nothing to convince Field of the existence of real numbers. Frege held that real numbers *are* ratios between quantities. Once we have abandoned the superstitious nominalist horror of abstract objects in general, there would be nothing problematic about the existence of real numbers in the context of some empirical theory involving quantities of one or another kind, if they were identified with ratios between those quantities. What real numbers there were would depend upon what quantities there were: there would be no danger of our not having sufficiently many real numbers for our purposes.

The difficulty about mathematical objects thus arises because we want our mathematical theories to be pure in the sense of not depending for the existence of their objects on empirical reality, but yet to satisfy axioms guaranteeing sufficiently many objects for any applications that we may have occasion to make. The significant distinction is not between abstract objects and concrete objects, but between mathematical objects and all others, concrete or abstract. Plenty of abstract objects exist only contingently, the equator, for example: their existence is contingent upon the existence of concrete objects, and upon their behavior or the relations obtaining between them. Ratios between empirically given quantities would be dependent abstract objects of this kind. By contrast, the existence of mathematical objects is assumed to be independent of what concrete objects the world contains.

In order to confer upon a general term applying to concrete objects – the term "star," for example – a sense adequate for its use in existential statements and universal generalizations, we consider it enough that we have a sharp criterion for whether it applies to a given object, and a sharp criterion for what is to count as one such object – one star, say – and what as two distinct ones: a criterion of application and a criterion of identity. The same indeed holds true for a term, like "prime number," applying to mathematical objects, but regarded as defined over an already given domain. It is otherwise, however, for such a mathematical term as "natural number" or "real number" which determines a domain of quantification. For a term of this sort, we make a further demand: namely, that we should "grasp" the domain, that is, the totality of objects to which the term applies, in the sense of being able to circumscribe it by saying what objects, in general, it comprises – what natural numbers, or what real numbers, there are.

The reason for this difference is evident. For any kind of concrete object, or of abstract object whose existence depends upon concrete objects, external reality will determine what objects of that kind there are; but what mathematical objects there are within a fundamental domain of quantification is supposed to be independent of how things happen to be in the world, and so, if it is to be determinate, *we* must determine it. On the face of it, indeed, a criterion of application and a criterion of identity do not suffice to confer determinate truth-conditions on generalizations involving some general term, even when it is a term covering concrete objects: they can only give them a content to be construed as embodying a *claim*. So understood, an existential statement amounts to a claim to be able to give an instance; a universal statement is of the form "Any object to which the term is recognized as being applicable will be found to satisfy such-and-such a further condition." An utterance that embodies a claim is accepted as justified if the one who makes it can vindicate his claim, and rejected as unjustified if he fails to do so; a universally quantified statement is shown to be unjustified if a counterexample comes to light, but is justified only if the speaker can give adequate grounds for the conditional expectation he arouses. The difference between such an utterance and one that carries some definite truth-condition is that the claim relates to what the speaker can do or what reasons he can give, whereas the truth-condition must be capable of being stated independently of his abilities or his knowledge.

This is not, however, how we usually think of quantified statements about empirical objects. We normally suppose that, given that we are clear what has to be true of a celestial object for it to be a star, and when a star observed on one occasion is the same as one observed on another, we need do nothing more to assure definite truth-conditions to statements of the form "There is a star with such-and-such a property" or "All stars have such-and-such a property." This assumption reflects our natural realism concerning the physical universe. Whether this realism about the physical universe is sound, or (as I myself strongly suspect) ought itself to be challenged, is a question not here at issue: what matters in the present context is the contrast between what we standardly take to be needed to secure determinate truth-conditions for statements involving generality in the empirical case and in the mathematical one.

We are, indeed, usually disposed to be quite as firmly resolved that our mathematical statements should have truth-conditions that they determinately either satisfy or fail to satisfy as we are that this should hold good of our empirical statements. This is something that it never occurred to Frege to doubt. He acknowledged the necessity for specifying the truth-conditions of the statements of a mathematical theory; unfortunately, he persuaded himself that the domain of the individual variables could be determined simply by laying down the formation rules of the fundamental terms and fixing the criterion of identity for them, which he did by means of an impredicative specification, and produced an ingenious but fallacious argument to this effect.

Despite his realism about mathematics, even Frege did not think that mathematical reality determined the truth or falsity of statements quantifying over a domain of mathematical objects, without our needing to specify their truth-conditions; and his successors, mindful of the disaster that overtook him, have accepted the need to specify the domain outright, or to form some conception of it, before interpreting the primitive predicates of a theory as applying to elements of that domain. Notoriously, however, we have found little better means of accomplishing this task than Frege did. The characterizations of the domains of fundamental mathematical theories such as the theory of real numbers that we are accustomed to employ usually convince no one that any sharp conception underlies them save those who are already convinced; this leads to an impasse in the philosophy of mathematics where faith opposes incredulity without either possessing the resources to overcome the other. Moreover, this outcome seems intrinsic to the situation. A fundamental mathematical theory, for present purposes, is one from which we originally derive our conception of a totality of the relevant cardinality: it appears evident that we cannot characterize the domain of such a theory without circularity.

What is the way out of this impasse? We may approach this by asking after the error that underlay the assumptions which led Frege into contradiction – not that involved in his fallacious justification of those assumptions, but in the assumptions themselves. We have grown so accustomed to the paradoxes of set theory that we no longer marvel at them; yet their discovery was one of the most profound conceptual discoveries of all time, fully worthy to rank with the discovery of irrational numbers. Cantor saw far more deeply into the matter than Frege did: he was aware, long before, that one cannot simply assume every concept to have an extension with a determinate cardinality. Yet even he did not see all the way: for he made the distinction between concepts that do, and those that do not, have such an extension an absolute one, whereas the depth of the discovery lies in the fact that it is not. Taken as an absolute distinction, it generates irresoluble perplexity. We are thoroughly at home with the conception of transfinite cardinal numbers; but consider what happens when someone is first introduced to that conception. A certain resistance has first to be overcome: to someone who has long been used to finite cardinals, and only to them, it seems obvious that there can only be finite cardinals. A cardinal number, for him, is arrived at by counting; and the very definition of an infinite totality is that it is impossible to count it. This is not a stupid prejudice. The scholastics favored an argument to show that the human race could not always have existed, on the ground that, if it had, there would be no number that would be the number of all the human beings there had ever been, whereas for every concept there must be a number which is that of the objects falling under it. All the same, the prejudice is one that can be overcome: the beginner can be persuaded that it makes sense, after all, to speak of the number of natural numbers. Once his initial prejudice has been overcome, the next stage is to convince the beginner that there are distinct transfinite cardinal numbers: not all infinite totalities have as many members as each other. When he has become accustomed to this idea, he is extremely likely to ask, "How many transfinite cardinals are there?" How should he be answered? He is very likely to be answered by being told, "You must not ask that question." But why should he not? If it was, after all, all right to ask, "How many numbers are there?", in the sense in which "number" meant "finite cardinal," how can it be wrong to ask the same question when "number" means "finite or transfinite cardinal"? A mere prohibition leaves the matter a mystery. If gives no help to say that there are some totalities so large that no number can be assigned to them. We can gain some grasp of the idea of a totality too big to be counted, even at the stage when we think that, if it cannot be counted, it does not have a number; but, once we have accepted that totalities too big to be counted may yet have numbers, the idea of one too big even to have a number conveys nothing at all. And merely to say, "If you persist in talking about the number of all cardinal numbers, you will run into contradiction," is to wield the big stick, not to offer an explanation.

The fact revealed by the set-theoretic paradoxes was the existence of indefinitely extensible concepts – a fact of which Frege did not dream and even Cantor had only an obscure perception. An indefinitely extensible concept is one such that, if we can form a definite conception of a totality all of whose members fall under that concept, we can, by reference to that totality, characterize a larger totality of all whose members fall under it. Russell's concept *class not a member of itself* provides a beautiful example of an indefinitely extensible concept. Suppose that we have conceived of a class *C* all of whose members fall under the concept. Then it would certainly involve a contradiction to suppose *C* to be a member of itself. Hence, by considering the totality consisting of the members of *C* together with *C* itself, we have specified a more inclusive totality than *C* all of whose members fall under the concept *class not a member of itself*. Are we to say, then, that the concept *class not a member of itself* does not have an extension? We must indeed say that, by the nature of the case, we can form no conception of the totality of all objects falling under that concept, even of the totality of all objects of which we can conceive and which we should recognize as falling under that concept. On the other hand, to the question whether it is wrong to suppose that every concept defined over a determinate domain of distinguishable objects has an extension we must answer, "Surely not." Suppose that we have succeeded in specifying, or in clearly conceiving, some determinate domain of distinguishable objects, some or all of which are classes, and over which the membership relation is well defined. Then we must regard it as determinate, for an element of that domain, whether or not it is a class and, if so, whether or not it is a member of itself. A concept whose application to a determinate totality is itself determinate must pick out a determinate subtotality of elements that fall under it; and so the concept *class not a member of itself* must have a definite extension within that domain. All that we are forbidden to suppose is that any class belonging to the domain coincides with the extension of that concept. Frege's mistake thus did not lie in taking the notion of a class, or, more exactly, his notion of a value-range (the extension of a function), to be a logical rather than a mathematical one, as is sometimes said, not even in any straightforward sense, in supposing every function to have an extension; it lay in failing to perceive the notion to be an indefinitely extensible one, or, more generally, in failing to allow for indefinitely extensible concepts at all.

There can be no objection to quantifying over all objects falling under some indefinitely extensible concept, say over everything we should, given an intelligible description of it, recognize as an ordinal number, provided that we do not think of the statements formed by means of such quantification as having determinate truth-conditions; we can understand them only as making claims of the kind already sketched. They will not then satisfy the laws of classical logic, but only the weaker laws of intuitionistic logic. Abandoning classical logic will not, by itself, preserve us from contradiction if we maintain the same assumptions as before; but, since we no longer conceive ourselves to be quantifying over a fully determinate totality, we shall have no motive to do so.

Cantor's celebrated diagonal argument to show that the totality of real numbers is not denumerable has precisely the form of a principle of extension for an indefinitely extensible concept: given any denumerable totality of real numbers, we can define, *in terms of that totality*, a real number that does not belong to it. The argument does not show that the real numbers form a non-denumerable totality unless we assume at the outset that they form a determinate totality comprising all that we shall ever recognize as a real number: the alternative is to regard the concept *real number* as an indefinitely extensible one. It might be objected that no contradiction results from taking the real numbers to form a determinate totality. There is, however, no ground to suppose that treating an indefinitely extensible concept as a definite one will always lead to inconsistency; it may merely lead to our supposing ourselves to have a definite idea when we do not. This hypothesis explains the lameness of our attempts at a characterization of the supposed determinate totality of all real numbers, and relieves us of the embarrassment resulting from the apparent need for such a characterization; for the characterization of an indefinitely extensible concept demands much less than the once-for-all characterization of a determinate totality.

The adoption of this solution has a steep price, which most mathematicians would be unwilling to pay: the rejection of classical methods of argument in mathematics in favor of constructive ones. The prejudices of mathematicians do not constitute an argument, however: the important question for us is whether constructive mathematics is adequate for applications. We have so far assumed a realist view of the physical universe: would this be compatible with a less than fully realist view of mathematics? Not on the face of it; but, having taken the concept *real number*, but not yet that of *natural number*, to be indefinitely extensible, we have not yet attained a fully constructive conception of the real numbers, since they are essentially infinite objects, involving some notion such as that of an infinite sequence. By contrast, each natural number is a finite, that is, finitely describable, object: the totality of natural numbers is therefore of a radically different kind from the totality of real numbers. It does not follow that we may call it a determinate totality. Consider Frege's "proof" that every natural number has a successor: given any initial segment of the natural numbers, from 0 to n, the number of terms of that segment is again a natural number, but one larger than any term of the segment. As Frege presents it, the proof begs the question, since it rests on the assumption that we already have a domain containing the cardinal number of any subset of that domain; but the striking resemblance between this argument and that which showed the indefinite extensibility of the concept *set not a member of itself* suggests a reinterpretation of it as showing the indefinite extensibility of the concept *natural number*. The natural objection is that, when we attain the totality of all natural numbers, the supposed principle of extension ceases to apply, since the number of natural numbers is not itself a natural number.

This, however, is again to assume that we have a grasp of the totality of natural numbers: but do we? Certainly we have a clear grasp of the step from any natural number to its successor; but this is merely the essential principle of extension. The totality of natural numbers contains what, from our standpoint, are enormous numbers, and yet others relatively to which those are minute, and so on indefinitely; do we really have a grasp of such a totality?

A natural response is to claim that the question has been begged. In classing *real number* as an indefinitely extensible concept, we have *assumed* that any totality of which we can have a definite conception is at most denumerable; in classing *natural number* as one, we have assumed that such a totality will be finite. Burden-of-proof controversies are always difficult to resolve; but, in this instance, it is surely clear that it is the other side that has begged the question. It is claiming to be able to convey a conception of the totality of real numbers, without circularity, to one who does not yet have it. We are assuming that the latter does not have, either, a conception of any other totality of the power of the continuum. He therefore does not *assume* as a principle that any totality of which it is possible to form a definite conception is at most denumerable: he merely has as yet no conception of any totality of higher cardinality. Likewise, a conception of the totality of the natural numbers is supposed to be conveyed to one as yet unaware of any but finite totalities; but all that he is given is a principle of extension for passing from any finite totality to a larger one. The fact is that a concept determining an intrinsically infinite totality – one whose infinity follows from the concept itself – simply *is* an indefinitely extensible one; in the long history of mankind's grappling with the notion of infinity, this fact could not be clearly perceived until the set-theoretic paradoxes forced us to recognize the existence of indefinitely extensible concepts. Not all indefinitely extensible concepts are equally exorbitant, indeed; we have been long familiar, from the work initiated by Cantor, with the fact that there is not just one uniform notion of infinity, but a variety of them: but this should not hinder us from acknowledging that every concept with an intrinsically infinite extension belongs to one or another type of indefinitely extensible one.

The recognition of this fact compels us to adopt a thoroughly constructive version of analysis: we cannot fully grasp any one real number, but only to an approximation, although there are important differences in the extent to which we can grasp them, for example between those for which we have an effective method of finding their decimal expansions and those for which we do not. This strongly suggests that the constructive theory of real numbers is *better* adapted to their applications than its classical counterpart; for, although the realist assumption is that every quantity has some determinate magnitude, represented, relatively to some unit, by a real number, it is a commonplace that we can never arrive at that magnitude save to within an approximation.

This partly answers our question how far a realist view of the physical universe could survive the replacement of classical by constructive analysis. On a constructive view of the matter, the magnitude of any quantity, relatively to a unit, may be taken to be given by a particular real number, which we may at any stage determine to a closer approximation by refinement of the measurement process; but no precise determination of it will ever be warranted, nor presumed to obtain independently of our incapacity to determine it. The assumption that it has a precise value, standing in determinate order relations to all rational numbers and known to God if not to us, stems from the realist metaphysics that informs much of our physical theory. This observation does not, however, settle whether the assumption is integral to those theories or a piece of metaphysics detachable from them; and this question cannot be answered without detailed investigation. If it should prove that the applications of any mathematical theory to physics can be adequately effected by a constructive version of that theory, it would follow that realist assumptions play no role in physical theory as such, but merely govern the interpretation we put upon our physical theories; in this case, physics itself might for practical purposes remain aloof from metaphysics. If, on the other hand, it were to prove that constructive mathematics is inadequate to yield the applications of mathematics that we actually make, and that classical mathematics is strictly required for them, it would follow that those realist assumptions do play a significant role in physics as presently understood. That would not settle the matter, of course. There would then be a metaphysical question whether the realist assumptions could be justified; if not, our physics, as well as our mathematics, would call for revision along constructive lines.

But would it not be better to adopt Field's approach, rather than one calling for a revision of practice on the part of the majority of mathematicians? What the answer ought to be if there were any real promise of success for Field's enterprise is hard to say; but there is a simple reason why he has provided none. He proposes to infer the conservativeness of a given mathematical theory with respect to a given physical theory from the relevant representation theorem by means of a uniform argument resting upon the consistency of a version of ZF with Urelemente. Why, then, does he believe ZF to be consistent? Most people do, indeed: but then most people are not nominalists. They believe ZF to be consistent because they suppose themselves in possession of a perhaps hazily conceived intuitive model of the theory; Field can have no such reason. Any such intuitive model must involve a conception of the totality of ordinals less than the first strongly inaccessible one; and no explanation of the term "model" has been offered according to which the elements of a model need not be supposed to exist. The reason offered by Field himself for believing in the consistency of ZF is that "if it weren't consistent someone would have probably discovered an inconsistency in it by now"; he refers to this as inductive knowledge. To have an inductive basis for the conviction, it is not enough to observe that some theories have been discovered to be inconsistent in a relatively short time; it would be necessary also to know, of some theories not discovered to be inconsistent within around three-quarters of a century, that they *are* consistent. Without non-inductive knowledge of the consistency of some comparable mathematical theories there can be no inductive knowledge of the consistency of any mathematical theory. Field's proof of conservativeness therefore rests upon a conviction for which he can claim no ground whatever; one far more extravagant than any belief in the totality of real numbers.

I have argued that it is useless to cast around for new answers to the question what mathematics is about: the logicists already had essentially the correct answer. They were defeated by the problem of mathematical objects because they had incompatible aims: to represent mathematics as a genuine science, that is, as a body of *truths*, and not a mere auxiliary of other sciences; to keep it uncontaminated from empirical notions; and to justify classical mathematics in its entirely, and, in particular, the untrammeled use of classical logic in mathematical proofs. Field wishes to abandon the first, and others argue for abandoning the second: I have argued the abandonment of the third. On some conceptions of logic, it may be protested that this is not a purely logicist account, on the ground that mathematical objects still do not qualify to be called logical objects; but this is little more than a boundary dispute. If the domains of the fundamental mathematical theories are taken to be given by indefinitely extensible concepts, then we have what Frege sought and failed to find: a way of characterizing them that renders our right to refer to them unproblematic while yet leaving the existence of their elements independent of any contingent states of affairs. If the price of this solution to the problem of the basis of those theories is that argumentation within mathematics is compelled to become more cautious than that which classical mathematicians have been accustomed to use, and more sensitive to distinctions to which they have been accustomed to be indifferent, it is a price worth paying, especially if the resulting versions of the theories indeed prove more apt for their applications.