Philosophy of mathematics

Selected readings

SECOND EDITION

Edited by
Paul Benacerraf
STUART PROFESSOR OF PHILOSOPHY
PRINCETON UNIVERSITY

Hilary Putnam
WALTER BEVERLY PEARSON PROFESSOR OF
MODERN MATHEMATICS AND MATHEMATICAL LOGIC
HARVARD UNIVERSITY

1983

Cambridge University Press
Cambridge
London New York New Rochelle
Melbourne Sydney

ordinary classical method assumed by the platonist for specifying truth-conditions for the sentences of a given language, viz. that we are first given a domain of objects and know what it is to assign an object in this domain to a free variable (i.e. to treat the variable as in effect denoting that object), here requires modification. We must be told what process has to take place in order to assign a term to a particular one of the objects in the domain, as well as under what conditions an atomic sentence formed with such a term is true or false. The point is quite independent of the issue on which much of the debate between platonists and constructivists has been concentrated, namely whether, when the domain has been fixed and the truth-conditions of atomic sentences laid down, those of quantified sentences automatically follow. The present point, if taken seriously, would, at least in the case of non-denumerable domains, lead quite a long way in the direction of constructivism.

What, then, about those intuitions by which we agreed that the platonists were rightly impressed?—what is their status at the end of this peregrination? Once we have looked full in the face the nature of the task of supplying a given range of sentences with a determinate sense—whether by stipulating conditions of truth and falsity or conditions of provability—we shall, I think, be prepared to treat the notion of an intuitive model as we do that of an intuitive proof, in the sense of one lacking in full rigour. An intuitive model is a half-formed conception of how to determine truth-conditions for a given class of sentences. It is not an ultimate guarantee of consistency, nor the product of a special faculty of acquiring mathematical understanding. It is merely an idea in the embryonic stage, before we have succeeded in the laborious task of bringing it to birth in a fully explicit form. That is how all important ideas form, and the task of bringing them to birth is perhaps the most difficult and interesting of all intellectual tasks. Intuition is not a special source of ineffable insight: it is the womb of articulated understanding.

14. The Philosophical Basis of Intuitionistic Logic (1973)

THE QUESTION WITH which I am here concerned is: What plausible rationale can there be for repudiating, within mathematical reasoning, the canons of classical logic in favour of those of intuitionistic logic? I am, thus, not concerned with justifications of intuitionistic mathematics from an eclectic point of view, that is, from one which would admit intuitionistic mathematics as a legitimate and interesting form of mathematics alongside classical mathematics: I am concerned only with the standpoint of the intuitionists themselves, namely that classical mathematics employs forms of reasoning which are not valid on any legitimate way of construing mathematical statements (save, occasionally, by accident, as it were, under a quite unintended reinterpretation). Nor am I concerned with exegesis of the writings of Brouwer or of Heyting: the question is what forms of justification of intuitionistic mathematics will stand up, not what particular writers, however eminent, had in mind. And, finally, I am concerned only with the most fundamental feature of intuitionistic mathematics, its underlying logic, and not with the other respects (such as the theory of free choice sequences) in which it differs from classical mathematics. It will therefore be possible to conduct the discussion wholly at the level of elementary number theory. Since we are, in effect, solely concerned with the logical constants—with the sentential operators and the first-order quantifiers—our interest lies only with the most general features of the notion of a mathematical construction, although it will be seen that we need to consider these in a somewhat delicate way.

Any justification for adopting one logic rather than another as the logic for mathematics must turn on questions of *meaning*. It would be impossible to contrive such a justification which took meaning for granted, and represented the question as turning on knowledge or certainty. We are certain of the truth of a statement when we have conclusive grounds for it and are certain that the grounds which we have *are* valid grounds for it and *are* conclusive. If classical arguments for mathematical statements are called in question, this cannot possibly be because it is thought that we are, in general, unable to tell with certainty whether an argument is classically valid, unless it is also intuitionistically valid: rather, it must be that what is being put in doubt is whether arguments which are valid by classical but not by intuitionistic criteria are absolutely valid, that is, whether they really do conclusively establish their conclusions as true. Even if it were held that classical arguments, while not in general absolutely valid, nevertheless always conferred a high probability on their conclusions, it would be wrong to characterise the motive for employing only intuitionistic arguments as lying in a desire to attain knowledge in place of mere probable opinion in mathematics, since the very thesis that the use of classical arguments did not lead to knowledge would represent the crucial departure from the classical conception, beside which the question of whether or not one continued to make use of classical arguments as mere probabilistic reasoning is comparatively insignificant. (In any case, within standard intuitionistic mathematics, there is no reason whatever why the existence of a classical proof of it should render a statement probable, since if, e.g., it is a statement of analysis, its being a classical theorem does not prevent it from being intuitionistically disprovable.)

So far as I am able to see, there are just two lines of argument for repudiating classical reasoning in mathematics in favour of intuitionistic reasoning. The first runs along the following lines. The meaning of a mathematical statement determines and is exhaustively determined by its *use*. The meaning of such a statement cannot be, or contain as an ingredient, anything which is not manifest in the use made of it, lying solely in the mind of the individual who apprehends that meaning: if two individuals agree completely about the use to be made of the statement, then they agree about its meaning. The reason is that the meaning of a statement consists solely in its rôle as an instrument of communication between individuals, just as the powers of a chess-piece consist solely in its rôle in the game according to the rules. An individual cannot communicate what he cannot be observed to communicate: if one individual associated with a mathematical symbol or formula some mental content, where the association did not lie in the use he made of the symbol or formula, then he could not convey that content by means of the symbol or formula, for his audience would be unaware of the association and would have no means of becoming aware of it.

The argument may be expressed in terms of the *knowledge* of meaning,

i.e. of understanding. A model of meaning is a model of understanding, i.e. a representation of what it is that is known when an individual knows the meaning. Now knowledge of the meaning of a particular symbol or expression is frequently verbalisable knowledge, that is, knowledge which consists in the ability to state the rules in accordance with which the expression or symbol is used or the way in which it may be replaced by an equivalent expression or sequence of symbols. But to suppose that, in general, a knowledge of meaning consisted in verbalisable knowledge would involve an infinite regress: if a grasp of the meaning of an expression consisted, in general, in the ability to *state* its meaning, then it would be impossible for anyone to learn a language who was not already equipped with a fairly extensive language. Hence that knowledge which, in general, constitutes the understanding of the language of mathematics must be implicit knowledge. Implicit knowledge cannot, however, meaningfully be ascribed to someone unless it is possible to say in what the manifestation of that knowledge consists: there must be an observable difference between the behaviour or capacities of someone who is said to have that knowledge and someone who is said to lack it. Hence it follows, once more, that a grasp of the meaning of a mathematical statement must, in general, consist of a capacity to use that statement in a certain way, or to respond in a certain way to its use by others.

Another approach is via the idea of learning mathematics. When we learn a mathematical notation, or mathematical expressions, or, more generally, the language of a mathematical theory, what we learn to do is to make use of the statements of that language: we learn when they may be established by computation, and how to carry out the relevant computations, we learn from what they may be inferred and what may be inferred from them, that is, what rôle they play in mathematical proofs and how they can be applied in extra-mathematical contexts, and perhaps we learn also what plausible arguments can render them probable. These things are all that we are shown when we are learning the meanings of the expressions of the language of the mathematical theory in question, because they are all that we can be shown: and, likewise, our proficiency in making the correct use of the statements and expressions of the language is all that others have from which to judge whether or not we have acquired a grasp of their meanings. Hence it can only be in the capacity to make a correct use of the statements of the language that a grasp of their meanings, and those of the symbols and expressions which they contain, can consist. To suppose that there is an ingredient of meaning which transcends the use that is made of that which carries the meaning is to suppose that someone might have learned all that is directly taught when the language of a mathematical theory is taught to him, and might then behave in every way like someone who understood that language, and yet not actually understand it, or understand it only incorrectly. But to suppose this is to make meaning ineffable, that is, in principle incommunicable. If this is possible, then no one individual ever has a guarantee that he is understood by any other individual; for all he knows, or can ever know, everyone else may attach to his words or to the symbols which he employs a meaning quite different from that which he attaches to them. A notion of meaning so private to the individual is one that has become completely irrelevant to mathematics as it is actually practised, namely as a body of theory on which many individuals are corporately engaged, an enquiry within which each can communicate his results to others.

It might seem that an approach to meaning which regarded it as exhaustively determined by use would rule out any form of revisionism. If use constitutes meaning, then, it might seem, use is beyond criticism: there can be no place for rejecting any established mathematical practice, such as the use of certain forms of argument or modes of proof, since that practice, together with all others which are generally accepted, is simply constitutive of the meanings of our mathematical statements, and we surely have the right to make our statements mean whatever we choose that they shall mean. Such an attitude is one possible development of the thesis that use exhaustively determines meaning: it is, however, one which can, ultimately, be supported only by the adoption of a holistic view of language. On such a view, it is illegitimate to ask after the content of any single statement, or even after that of any one theory, say a mathematical or a physical theory; the significance of each statement or of each deductively systematised body of statements is modified by the multiple connections which it has, direct and remote, with other statements in other areas of our language taken as a whole, and so there is no adequate way of understanding the statement short of knowing the entire language. Or, rather, even this image is false to the facts: it is not that a statement or even a theory has, as it were, a primal meaning which then gets modified by the interconnections that are established with other statements and other theories; rather, its meaning simply consists in the place which it occupies in the complicated network which constitutes the totality of our linguistic practices. The only thing to which a definite content may be attributed is the totality of all that we are, at a given time, prepared to assert; and there can be no simple model of the content which that totality of assertions embodies; nothing short of a complete knowledge of the language can reveal it.

Frequently such a holistic view is modified to the extent of admitting a class of observation statements which can be regarded as more or less directly registering our immediate experience, and hence as each carrying a determinate individual content. These observation statements lie, in Quine's famous image of language, at the periphery of the articulated structure formed by all the sentences of our language, where alone experience impinges. To these peripheral sentences, meanings may be ascribed in a more or less straightforward manner, in terms of the observational stimuli which prompt assent to and dissent from them. No comparable model of meaning is available for the sentences which lie further towards the interior of the structure: an understanding of them consists solely in a grasp of their place in the structure as a whole and their interaction with its other constituent sentences. Thus, on such a view, we may accept a mathematical theory, and admit its theorems as true, only because we find in practice that it serves as a convenient substructure deep in the interior of the complex structure which forms the total theory: there can be no question of giving a representation of the truth-conditions of the statements of the mathematical theory under which they may be judged individually as acceptable, or otherwise, in isolation from the rest of language.

Such a conception bears an evident analogy with Hilbert's view of classical mathematics; or, more accurately, with Boole's view of his logical calculus. For Hilbert, a definite individual content, according to which they may be individually judged as correct or incorrect, may legitimately be ascribed only to a very narrow range of statements of elementary number theory: these correspond to the observation statements of the holistic conception of language. All other statements of mathematics are devoid of such a content, and serve only as auxiliaries, though psychologically indispensable auxiliaries, to the recognition as correct of the finitistic statements which alone are individually meaningful. The other mathematical statements are not, on such a view, devoid of significance: but their significance lies wholly in the rôle which they play within the mathematical theories to which they belong, and which are themselves significant precisely because they enable us to establish the correctness of finitistic statements. Boole likewise distinguished, amongst the formulas of his logical calculus, those which were interpretable from those which were uninterpretable: a deduction might lead from some interpretable formulas as premisses, via uninterpretable formulas as intermediate steps, to a conclusion which was once more interpretable.

The immediately obvious difficulty about such a manner of construing a mathematical, or any other, theory is to know how it can be justified. How can we be sure that the statements or formulas to which we ascribe a content, and which are derived by such a means, are true? The difference between

Hilbert and Boole, in this respect, was that Hilbert took the demand for justification seriously, and saw the business of answering it as the prime task for his philosophy of mathematics, while Boole simply ignored the question. Of course, the most obvious way to find a justification is to extend the interpretation to all the statements or formulas with which we are concerned, and, in the case of Boole's calculus, this is very readily done, and indeed yields a great simplification of the calculus. Even in Hilbert's case, the consistency proof, once found, does yield an interpretation of the infinitistic statements, though one which is relative to the particular proof in which they occur, not one uniform for all contexts. Without such a justification, the operation of the mechanism of the theory or the language remains quite opaque to us; and it is because the holist is oblivious of the demand for justification, or of the unease which the lack of one causes us, that I said that he is to be compared to Boole rather than to Hilbert. In his case, the question would become: With what right do we feel an assurance that the observation statements deduced with the help of the complex theories, mathematical, scientific and otherwise, embedded in the interior of the total linguistic structure, are true, when these observation statements are interpreted in terms of their stimulus meanings? To this the holist attempts no answer, save a generalised appeal to induction: these theories have 'worked' in the past, in the sense of having for the most part yielded true observation statements, and so we have confidence that they will continue to work in the future.

The path of thought which leads from the thesis that use exhaustively determines meaning to an acceptance of intuitionistic logic as the correct logic for mathematics is one which rejects a holistic view of mathematics, and insists that each statement of any mathematical theory must have a determinate individual content. A grasp of this content cannot, in general, consist of a piece of verbalisable knowledge, but must be capable of being fully manifested by the use of the statement: but that does not imply that every aspect of its existing use is sacrosanct. An existing practice in the use of a certain fragment of language is capable of being subjected to criticism if it is impossible to systematise it, that is, to frame a model whereby each sentence carries a determinate content which can, in turn, be explained in terms of the use of that sentence. What makes it possible that such a practice may prove to be incoherent and therefore in need of revision is that there are different aspects to the use of a sentence; if the whole practice is to be capable of systematisation in the present sense, there must be a certain harmony between these different aspects. This is already apparent from the holistic examples already cited. One aspect of the use of observation statements lies in the propensities we have acquired to assent to and dissent from them under certain types of stimuli; another lies in the possibility of deducing them by means of non-observational statements, including highly theoretical ones. If the linguistic system as a whole is to be coherent, there must be harmony between these two aspects: it must not be possible to deduce observation statements from which the perceptual stimuli require dissent. Indeed, if the observation statements are to retain their status as observation statements, a stronger demand must be made: of an observation statement deduced by means of theory, it must hold that we can place ourselves in a situation in which stimuli occur which require assent to it. This condition is thus a demand that, in a certain sense, the language as a whole be a conservative extension of that fragment of the language containing only observation statements. In just the same way, Hilbert's philosophy of mathematics requires that classical number theory, or even classical analysis, be a conservative extension of finitistic number theory.

For utterances considered quite generally, the bifurcation between the two aspects of their use lies in the distinction between the conventions governing the occasions on which the utterance is appropriately made and those governing both the responses of the hearer and what the speaker commits himself to by making the utterance: schematically, between the *conditions for* the utterance and the *consequences of* it. Where, as in mathematics, the utterances with which we are concerned are *statements*, that is, utterances by means of which assertions can be effected, this becomes the distinction between the grounds on which the statement can be asserted and its inferential consequences, the conclusions that can be inferred from it. Plainly, the requirement of harmony between these in respect of some type of statement is the requirement that the addition of statements of that type to the language produces a conservative extension of the language; i.e., that it is not possible, by going via statements of this type as intermediaries, to deduce from premisses not of that type conclusions, also not of that type, which could not have been deduced before. In the case of the logical constants, a loose way of putting the requirement is to say that there must be a harmony between the introduction and elimination rules; but, of course, this is not accurate, since the whole system has to be considered (in classical logic, for example, it is possible to infer a disjunctive statement, say by double negation elimination, without appeal to the rule of disjunction introduction). An alternative way of viewing the dichotomy between the two principal aspects of the use of statements is as a contrast between *direct* and *indirect* means of establishing them. So far as a logically complex statement is concerned, the introduction rules governing the logical constants occurring in the statement display the most direct means of establishing the statement, step by step in accordance with its logical structure; but the statement may be accepted on the basis of a complicated deduction which relies also on elimination rules, and we require a harmony which obtains only if a statement that has been indirectly established always could (in some sense of 'could') have been established directly. Here again the demand is that the admission of the more complex inferences yield a conservative extension of the language. When only introduction rules are used, the inference involves only statements of logical complexity no greater than that of the conclusion: we require that the derivation of a statement by inferences involving statements of greater logical complexity shall be possible only when its derivation by the more direct means is in some sense already possible.

On any molecular view of language—any view on which individual sentences carry a content which belongs to them in accordance with the way they are compounded out of their own constituents, independently of other sentences of the language not involving those constituents—there must be some demand for harmony between the various aspects of the use of sentences, and hence some possibility of criticising or rejecting existing practice when it does not display the required harmony. Exactly what the harmony is which is demanded depends upon the theory of meaning accepted for the language, that is, the general model of that in which the content of an individual sentence consists; that is why I rendered the above remarks vague by the insertion of phrases like 'in some sense'. It will always be legitimate to demand, of any expression or form of sentence belonging to the language, that its addition to the language should yield a conservative extension; but, in order to make the notion of a conservative extension precise, we need to appeal to some concept such as that of truth or that of being assertible or capable in principle of being established, or the like; and just which concept is to be selected, and how it is to be explained, will depend upon the theory of meaning that is adopted.

A theory of meaning, at least of the kind with which we are mostly familiar, seizes upon some one general feature of sentences (at least of assertoric sentences, which is all we need be concerned with when considering the language of mathematics) as central: the notion of the content of an individual sentence is then to be explained in terms of this central feature. The selection of some one such feature of sentences as central to the theory of meaning is what is registered by philosophical dicta of the form, 'Meaning is ...'—e.g., 'The meaning of a sentence is the method of its verification', 'The meaning of a sentence is determined by its truth-conditions', etc. (The slogan 'Meaning is use' is, however, of a different character: the 'use' of a sentence is not, in this sense, a *single* feature; the slogan simply restricts the *kind* of feature that may legitimately be appealed to as constituting or determining meaning.) The justification for thus selecting some one single feature of sentences as central—as being that in which their individual meanings consist—is that it is hoped that every other feature of the use of sentences can be derived, in a uniform manner, from this central one. If, e.g., the notion of truth is taken as central to the theory of meaning, then the meanings of individual expressions will consist in the manner in which they contribute to determining the truth-conditions of sentences in which they occur; but this conception of meaning will be justified only if it is possible, for an arbitrary assertoric sentence whose truth-conditions are taken as known, to describe, in terms of the notion of truth, our actual practice in the use of such a sentence; that is, to give a general characterisation of the linguistic practice of making assertions, of the conditions under which they are made and the responses which they elicit. Obviously, we are very far from being able to construct such a general theory of the use of sentences, of the practice of speaking a language; equally obviously, it is likely that, if we ever do attain such an account, it will involve a considerable modification of the ideal pattern under which the account will take a quite general form, irrespective of the individual content of the sentence as given in terms of whatever is taken as the central notion of the theory of meaning. But it is only to the extent that we shall eventually be able to approximate to such a pattern that it is possible to give substance to the claim that it is in terms of some *one* feature, such as truth or verification, that the individual meanings of sentences and of their component expressions are to be given.

It is the multiplicity of the different features of the use of sentences, and the consequent legitimacy of the demand, given a molecular view of language, for harmony between them, that makes it possible to criticise existing practice, to call in question uses that are actually made of sentences of the language. The thesis with which we started, that use exhaustively determines meaning, does not, therefore, conflict with a revisionary attitude to some aspect of language: what it does do is to restrict the selection of the feature of sentences which is to be treated as central to the theory of meaning. On a platonistic interpretation of a mathematical theory, the central notion is that of truth: a grasp of the meaning of a sentence belonging to the language of the theory consists in a knowledge of what it is for that sentence to be true. Since, in general, the sentences of the language will not be ones whose truth-value we are capable of effectively deciding, the condition for the truth of such a sentence will be one which we are not, in general, capable of recognising as obtaining whenever it obtains, or of getting ourselves into a position in which we can so recognise it. Nevertheless, on the theory of meaning which underlies platonism, an individual's grasp of the meaning of such a sentence consists in his knowledge of what the condition is which has to obtain for the sentence to be true, even though the condition is one which he cannot, in general, recognise as obtaining when it does obtain.

This conception violates the principle that use exhaustively determines meaning; or, at least, if it does not, a strong case can be put up that it does, and it is this case which constitutes the first type of ground which appears to exist for repudiating classical in favour of intuitionistic logic for mathematics. For, if the knowledge that constitutes a grasp of the meaning of a sentence has to be capable of being manifested in actual linguistic practice, it is quite obscure in what the knowledge of the condition under which a sentence is true can consist, when that condition is not one which is always capable of being recognised as obtaining. In particular cases, of course, there may be no problem, namely when the knowledge in question may be taken as verbalisable knowledge, i.e. when the speaker is able to *state*, in other words, what the condition is for the truth of the sentence; but, as we have already noted, this cannot be the general case. An ability to state the condition for the truth of a sentence is, in effect, no more than an ability to express the content of the sentence in other words. We accept such a capacity as evidence of a grasp of the meaning of the original sentence on the presumption that the speaker understands the words in which he is stating its truth-condition; but at some point it must be possible to break out of the circle: even if it were always possible to find an equivalent, understanding plainly cannot in general consist in the ability to find a synonymous expression. Thus the knowledge in which, on the platonistic view, a grasp of the meaning of a mathematical statement consists must, in general, be implicit knowledge, knowledge which does not reside in the capacity to state that which is known. But, at least on the thesis that use exhaustively determines meaning, and perhaps on any view whatever, the ascription of implicit knowledge to someone is meaningful only if he is capable, in suitable circumstances, of fully manifesting that knowledge. (Compare Wittgenstein's question why a dog cannot be said to expect that his master will come home next week.) When the sentence is one which we have a method for effectively deciding, there is again no problem: a grasp of the condition under which the sentence is true may be said to be manifested by a mastery of the decision procedure, for the individual may, by that means, get himself into a position in which he can recognise that the condition for the truth of the sentence obtains or does not obtain, and we may reasonably suppose that, in this

position, he displays by his linguistic behaviour his recognition that the sentence is, respectively, true or false. But, when the sentence is one which is not in this way effectively decidable, as is the case with the vast majority of sentences of any interesting mathematical theory, the situation is different. Since the sentence is, by hypothesis, effectively undecidable, the condition which must, in general, obtain for it to be true is not one which we are capable of recognising whenever it obtains, or of getting ourselves in a position to do so. Hence any behaviour which displays a capacity for acknowledging the sentence as being true in all cases in which the condition for its truth can be recognised as obtaining will fall short of being a full manifestation of the knowledge of the condition for its truth: it shows only that the condition can be recognised in certain cases, not that we have a grasp of what, in general, it is for that condition to obtain even in those cases when we are incapable of recognising that it does. It is, in fact, plain that the knowledge which is being ascribed to one who is said to understand the sentence is knowledge which transcends the capacity to manifest that knowledge by the way in which the sentence is used. The platonistic theory of meaning cannot be a theory in which meaning is fully determined by use.

If to know the meaning of a mathematical statement is to grasp its use; if we learn the meaning by learning the use, and our knowledge of its meaning is a knowledge which we must be capable of manifesting by the use we make of it: then the notion of *truth*, considered as a feature which each mathematical statement either determinately possesses or determinately lacks, independently of our means of recognising its truth-value, cannot be the central notion for a theory of the meanings of mathematical statements. Rather, we have to look at those things which are actually features of the use which we learn to make of mathematical statements. What we actually learn to do, when we learn some part of the language of mathematics, is to recognise, for each statement, what counts as establishing that statement as true or as false. In the case of very simple statements, we learn some computation procedure which decides their truth or falsity: for more complex statements, we learn to recognise what is to be counted as a proof or a disproof of them. That is the practice of which we acquire a mastery: and it is in the mastery of that practice that our grasp of the meanings of the statements must consist. We must, therefore, replace the notion of truth, as the central notion of the theory of meaning for mathematical statements, by the notion of *proof*: a grasp of the meaning of a statement consists in a capacity to recognise a proof of it when one is presented to us, and a grasp of the meaning of any expression smaller than a sentence must consist in a knowledge of the way in which its presence in a sentence contributes to determining what is to count as a proof of that sentence. This does not mean that we are obliged uncritically to accept the canons of proof as conventionally acknowledged. On the contrary, as soon as we construe the logical constants in terms of this conception of meaning, we become aware that certain forms of reasoning which are conventionally accepted are devoid of justification. Just because the conception of meaning in terms of proof is as much a molecular, as opposed to holistic, theory of meaning as that of meaning in terms of truth-conditions, forms of inference stand in need of justification, and are open to being rejected as unjustified. Our mathematical practice has been disfigured by a false conception of what our understanding of mathematical theories consisted in.

This sketch of one possible route to an account of why, within mathematics, classical logic must be abandoned in favour of intuitionistic logic obviously leans heavily upon Wittgensteinian ideas about language. Precisely because it rests upon taking with full seriousness the view of language as an instrument of social communication, it looks very unlike traditional intuitionist accounts, which, notoriously, accord a minimum of importance to language or to symbolism as a means of transmitting thought, and are constantly disposed to slide in the direction of solipsism. However, I said at the outset that my concern in this paper was not in the least with the exegesis of actual intuitionist writings: however little it may jibe with the view of the intuitionists themselves, the considerations that I have sketched appear to me to form one possible type of argument in favour of adopting an intuitionistic version of mathematics in place of a classical one (at least as far as the logic employed is concerned), and, moreover, an argument of considerable power. I shall not take the time here to attempt an evaluation of the argument, which would necessitate enquiring how the platonist might reply to it, and how the debate between them would then proceed: my interest lies, rather, in asking whether this is the only legitimate route to the adoption of an intuitionistic logic for mathematics.

Now the first thing that ought to strike us about the form of argument which I have sketched is that it is virtually independent of any considerations relating specifically to the *mathematical* character of the statements under discussion. The argument involved only certain considerations within the theory of meaning of a high degree of generality, and could, therefore, just as well have been applied to any statements whatever, in whatever area of language. The argument told in favour of replacing, as the central notion for the theory of meaning, the condition under which a statement is true, whether we know or can know when that condition obtains, by the condition under which we acknowledge the statement as conclusively established, a

condition which we must, by the nature of the case, be capable of effectively recognising whenever it obtains. Since we were concerned with mathematical statements, which we recognise as true by means of a proof (or, in simple cases, a computation), this meant replacing the notion of truth by that of proof: evidently, the appropriate generalisation of this, for statements of an arbitrary kind, would be the replacement of the notion of truth, as the central notion of the theory of meaning, by that of verification; to know the meaning of a statement is, on such a view, to be capable of recognising whatever counts as verifying the statement, i.e. as conclusively establishing it as true. Here, of course, the verification would not ordinarily consist in the bare occurrence of some sequence of sense-experiences, as on the positivist conception of the verification of a statement. In the mathematical case, that which establishes a statement as true is the production of a deductive argument terminating in that statement as conclusion; in the general case, a statement will, in general, also be established as true by a process of reasoning, though here the reasoning will not usually be purely deductive in character, and the premisses of the argument will be based on observation; only for a restricted class of statements—the observation statements—will their verification be of a purely observational kind, without the mediation of any chain of reasoning or any other mental, linguistic or symbolic process.

It follows that, in so far as an intuitionist position in the philosophy of mathematics (or, at least, the acceptance of an intuitionistic logic for mathematics) is supported by an argument of this first type, similar, though not necessarily identical, revisions must be made in the logic accepted for statements of other kinds. What is involved is a thesis in the theory of meaning of the highest possible level of generality. Such a thesis is vulnerable in many places: if it should prove that it cannot be coherently applied to any one region of discourse, to any one class of statements, then the thesis cannot be generally true, and the general argument in favour of it must be fallacious. Construed in this way, therefore, a position in the philosophy of mathematics will be capable of being undermined by considerations which have nothing directly to do with mathematics at all.

Is there, then, any alternative defence of the rejection, for mathematics, of classical in favour of intuitionistic logic? Is there any such defence which turns on the fact that we are dealing with *mathematical* statements in particular, and leaves it entirely open whether or not we wish to extend the argument to statements of any other general class?

Such a defence must start from some thesis about mathematical statements the analogue of which we are free to reject for statements of other kinds. It is plain what this thesis must be: namely, the celebrated thesis that mathematical statements do not relate to an objective mathematical reality existing independently of us. The adoption of such a view apparently leaves us free either to reject or to adopt an analogous view for statements of any other kind. For instance, if we are realists about the physical universe, then we may contrast mathematical statements with statements ascribing physical properties to material objects: on this combination of views, material-object statements do relate to an objective reality existing independently of ourselves, and are rendered true or false, independently of our knowledge of their truth-values or of our ability to attain such knowledge or the particular means, if any, by which we do so, by that independently existing reality; the assertion that mathematical statements relate to no such external reality gains its substance by contrast with the physical case. Unlike material objects, mathematical objects are, on this thesis, creations of the human mind: they are objects of thought, not merely in the sense that they can be thought about, but in the sense that their being is to be thought of; for them, *esse est concipi*.

On such a view, a conception of meaning as determined by truth-conditions is available for any statements which do relate to an independently existing reality, for then we may legitimately assume, of each such statement, that it possesses a determinate truth-value, true or false, independently of our knowledge, according as it does or does not agree with the constitution of that external reality which it is about. But, when the statements of some class do not relate to such an external reality, the supposition that each of them possesses such a determinate truth-value is empty, and we therefore cannot regard them as being given meanings by associating truth-conditions with them; we have, in such a case, *faute de mieux*, to take them as having been given meaning in a different way, namely by associating with them conditions of a different kind—conditions that we are capable of recognising when they obtain—namely, those conditions under which we take their assertion or their denial as being conclusively justified.

The first type of justification of intuitionistic logic which we considered conformed to Kreisel's dictum, 'The point is not the existence of mathematical objects, but the objectivity of mathematical truth': it bore directly upon the claim that mathematical statements possess objective truth-values, without raising the question of the ontological status of mathematical objects or the metaphysical character of mathematical reality. But a justification of the second type violates the dictum: it makes the question whether mathematical statements possess objective truth-values depend upon a prior decision as to the being of mathematical objects. And the difficulty about it lies in knowing on what we are to base the premiss that mathematical objects are the creations of human thought in advance of deciding what is the correct model for the meanings of mathematical statements or what is the correct conception of truth as relating to them. It appears that, on this view, before deciding whether a grasp of the meaning of a mathematical statement is to be considered as consisting in a knowledge of what has to be the case for it to be true or in a capacity to recognise a proof of it when one is presented, we have first to resolve the metaphysical question whether mathematical objects—natural numbers, for example—are, as on the constructivist view, creations of the human mind, or, as on the platonist view, independently existing abstract objects. And the puzzle is to know on what basis we could possibly resolve this metaphysical question, at a stage at which we do not even know what model to use for our understanding of mathematical statements. We are, after all, being asked to choose between two metaphors, two pictures. The platonist metaphor assimilates mathematical enquiry to the investigations of the astronomer: mathematical structures, like galaxies, exist, independently of us, in a realm of reality which we do not inhabit but which those of us who have the skill are capable of observing and reporting on. The constructivist metaphor assimilates mathematical activity to that of the artificer fashioning objects in accordance with the creative power of his imagination. Neither metaphor seems, at first sight, especially apt, nor one more apt than the other: the activities of the mathematician seem strikingly unlike those either of the astronomer or of the artist. What basis can exist for deciding which metaphor is to be preferred? How are we to know in which respects the metaphors are to be taken seriously, how the pictures are to be used?

Preliminary reflection suggests that the metaphysical question ought not to be answered first: we cannot, as the second type of approach would have us do, *first* decide the ontological status of mathematical objects, and then, with that as premiss, deduce the character of mathematical truth or the correct model of meaning for mathematical statements. Rather, we have first to decide on the correct model of meaning—either an intuitionistic one, on the basis of an argument of the first type, or a platonistic one, on the basis of some rebuttal of it; and then one or other picture of the metaphysical character of mathematical reality will force itself on us. If we have decided upon a model of the meanings of mathematical statements according to which we have to repudiate a notion of truth considered as determinately attaching, or failing to attach, to such statements independently of whether we can now, or ever will be able to, prove or disprove them, then we shall be unable to use the picture of mathematical reality as external to us and waiting to be discovered. Instead, we shall inevitably adopt the picture of that reality as being the product of our thought, or, at least, as coming into existence only as it is thought. Conversely, if we admit a notion of truth as attaching objectively to our mathematical statements independently of our knowledge, then, likewise, the picture of mathematical reality as existing, like the galaxies, independently of our observation of it will force itself on us in an equally irresistible manner. But, when we approach the matter in this way, there is no puzzle over the interpretation of these metaphors: psychologically inescapable as they may be, their non-metaphorical content will consist entirely in the two contrasting models of the meanings of mathematical statements, and the issue between them will become simply the issue as to which of these two models is correct. If, however, a view as to the ontological status of mathematical objects is to be treated as a *premiss* for deciding between the two models of meaning, then the metaphors cannot without circularity be explained solely by reference to those models; and it is obscure how else they are to be explained.

These considerations appear, at first sight, to be reinforced by reflection upon Frege's dictum, 'Only in the context of a sentence does a name stand for anything'. We cannot refer to an object save in the course of saying something about it. Hence, any thesis concerning the ontological status of objects of a given kind must be, at the same time, a thesis about what makes a statement involving reference to such objects true, in other words, a thesis about what properties an object of that kind can have. Thus, to say that fictional characters are the creations of the imagination is to say that a statement about a fictional character can be true only if it is imagined as being true, that a fictional character can have only those properties which it is part of the story that he has; to say that something is an object of sense—that for it *esse est percipi*—is to say that it has only those properties it is perceived as having: in both cases, the ontological thesis is a ground for rejecting the law of excluded middle as applied to statements about those objects. Thus we cannot separate the question of the ontological status of a class of objects from the question of the correct notion of truth for statements about those objects, i.e. of the kind of thing in virtue of which such statements are true, when they are true. This conclusion corroborates the idea that an answer to the former question cannot serve as a premiss for an answer to the latter one.

Nevertheless, the position is not so straightforward as all this would make it appear. From the possibility of an argument of the first type for the use of intuitionistic logic in mathematics, it is evident that a model of the meanings of mathematical statements in terms of proof rather than of truth need not rest upon any particular view about the ontological character of

mathematical objects. There is no substantial disagreement between the two models of meaning so long as we are dealing only with decidable statements: the crucial divergence occurs when we consider ones which are not effectively decidable, and the linguistic operation which first enables us to frame effectively undecidable mathematical statements is that of quantification over infinite totalities, in the first place over the totality of natural numbers. Now suppose someone who has, on whatever grounds, been convinced by the platonist claim that we do not create the natural numbers, and yet that reference to natural numbers is not a mere *façon de parler*, but is a genuine instance of reference to objects: he believes, with the platonist, that natural numbers are abstract objects, existing timelessly and independently of our knowledge of them. Such a person may, nevertheless, when he comes to consider the meaning of existential and universal quantification over the natural numbers, be convinced by a line of reasoning such as that which I sketched as constituting the first type of justification for replacing classical by intuitionistic logic. He may come to the conclusion that quantification over a denumerable totality cannot be construed in terms of our grasp of the conditions under which a quantified statement is true, but must, rather, be understood in terms of our ability to recognise a proof or disproof of such a statement. He will therefore reject a classical logic for number-theoretic statements in general, admitting only intuitionistically valid arguments involving them. Such a person would be accepting a platonistic view of the existence of mathematical objects (at least the objects of number theory), but rejecting a platonistic view of the objectivity of mathematical statements.

Our question is, rather, whether the opposite combination of views is possible: whether one may consistently hold that natural numbers are the creations of human thought, but yet believe that there is a notion of truth under which each number-theoretic statement is determinately either true or false, and that it is in terms of our grasp of their truth-conditions that our understanding of number-theoretic statements is to be explained. If such a combination is possible, then, it appears, there can be no route from the ontological thesis that mathematical objects are the creations of our thought to the model of the meanings of mathematical statements which underlies the adoption of an intuitionistic logic.

This is not the only question before us: for, even if these two views cannot be consistently combined, it would not follow that the ontological thesis could serve as a premiss for the constructivist view of the meanings of mathematical statements; our difficulty was to understand how the ontological thesis could have any substance if it were not merely a picture encapsulating that conception of meaning. The answer is surely this: that,

10

while it is surely correct that a thesis about the ontological status of objects of a given kind, e.g. natural numbers, must be understood as a thesis about that in which the truth of certain statements about those objects consists, it need not be taken as, in the first place, a thesis about the entire class of such statements; it may, instead, be understood as a thesis only about some restricted subclass of such statements, those which are basic to the very possibility of making reference to those objects. Thus, for example, the thesis that natural numbers are creations of human thought may be taken as a thesis about the sort of thing which makes a numerical equation or inequality true, or, more generally, a statement formed from such equations by the sentential operators and bounded quantification. To say that the only notion of truth we can have for number-theoretic statements generally is that which equates truth with our capacity to prove a statement is to prejudge the issue about the correct model of meaning for such statements, and therefore cannot serve as a premiss for the constructivist view of meaning. But to say that, for decidable number-theoretic statements, truth consists in provability, is not in itself to prejudge the question in what the truth of undecidable statements, involving unbounded quantification, consists: and hence the possibility is open that a view about the one might serve as a premiss for a view about the other. Our problem is to discover whether it can do so in fact: whether there is any legitimate route from the thesis that natural numbers are creations of human thought, construed as a thesis about the sort of thing which makes decidable number-theoretic statements true, to a view of the meanings of number-theoretic statements generally which would require the adoption for them of an intuitionistic rather than a classical logic.

In order to resolve this question, it is necessary for us to take a rather closer look at the notion of truth for mathematical statements, as understood intuitionistically. The most obvious suggestion that comes to mind in this connection is that the intuitionistic notion of truth conforms, just as does the classical notion, to Tarski's schema:

(T) S is true iff A,

where an instance of the schema is to be formed by replacing 'A' by some number-theoretic statement and 'S' by a canonical name of that sentence, as, e.g., in:

'There are infinitely many twin primes' is true iff there are infinitely many twin primes.

It is necessary to admit counter-examples to the schema (T) in any case in which we wish to hold that there exist sentences which are neither true nor false: for if we replace 'A' by such a sentence, the left-hand side of the biconditional becomes false (on the assumption that, if the negation of a sentence is true, that sentence is false), although, by hypothesis, the right-hand side is not false. But, in intuitionistic logic, that semantic principle holds good which stands to the double negation of the law of excluded middle as the law of bivalence stands to the law of excluded middle itself: it is inconsistent to assert of any statement that it is neither true nor false; and hence there seems no obstacle to admitting the correctness of the schema (T). Of course, in doing so, we must construe the statement which appears on the right-hand side of any instance of the schema in an intuitionistic manner. Provided we do this, a truth-definition for the sentences of an intuitionistic language, say that of Heyting arithmetic, may be constructed precisely on Tarski's lines, and will yield, as a consequence, each instance of the schema (T).

However, notoriously, such an approach leaves many philosophical problems unresolved. The truth-definition tells us, for example, that

'$598017 + 246532 = 844549$' is true

just in the case in which $598017 + 246532 = 844549$. We may perform the computation, and discover that $598017 + 246532$ does indeed equal 844549: but does that mean that the equation was already true before the computation was performed, or that it would have been true even if the computation had never been performed? The truth-definition leaves such questions quite unanswered, because it does not provide for inflections of tense or mood of the predicate 'is true': it has been introduced only as a predicate as devoid of tense as are all ordinary mathematical predicates; but its rôle in our language does not reveal why such inflections of tense or even of mood should be forbidden.

These difficulties raise their heads as soon as we make the attempt to introduce tense into mathematics, as intuitionism provides us with some inclination to do; this can be seen from the problems surrounding the theory of the creative subject. These problems are well brought out in Troelstra's discussion of the topic. It is evident that we ought to admit as an axiom

(α) $(\vdash_n A) \rightarrow A$;

if we know that, at any stage, A has been (or will be) proved, then we are certainly entitled to assert A. But ought we to admit the converse in the form (β) $$A \to \exists n\,(\vdash_n A)\,?$$

Its double negation

(γ) $$A \to \neg\neg\exists n\,(\vdash_n A)$$

is certainly acceptable: if we know that A is true, then we shall certainly never be able to assert, at least on purely mathematical grounds, that it will never be proved. But can we equate truth with the obtaining of a proof at some stage, in the past or in the future, as the equivalence:

(δ) $$A \leftrightarrow \exists n\,(\vdash_n A)$$

requires us to do? (To speak of 'truth' here seems legitimate, since, while Tarski's truth-predicate is a predicate of sentences, the sentential operator to which it corresponds is a redundant one, which can be inserted before or deleted from in front of any clause without change of truth-value.)

If we accept the axiom (β), and hence the equivalence (δ), we run into certain difficulties, on which Troelstra comments. The operator '$\exists n\,(\vdash_n \ldots)$' becomes a redundant truth-operator, and hence may be distributed across any logical constant, as in

(ε) $$(\vdash_k \forall m\, A(m)) \to \forall m\, \exists n\,(\vdash_n A(m)).$$

As Troelstra observes, this appears to have the consequence that, if we have once proved a universally quantified statement, we are in some way committed to producing, at some time in the future, individual proofs of all its instances, whereas, palpably, we are under no such constraint. The solution to which he inclines is that proposed by Kreisel, namely that the operator '$\vdash_n$' must be so construed that a proof, at stage n, of a universally quantified statement counts as being, at the same time, a proof of each instance, so that we could assert the stronger thesis

(ζ) $$(\vdash_k \forall m\, A(m)) \to \forall m\,(\vdash_k A(m)).$$

(Troelstra in fact recommends this interpretation on separate grounds, as enabling us to escape a paradox about constructive functions; he himself points out, however, that this paradox can alternatively be avoided by introducing distinctions of level which seem intrinsically plausible.) The difficulty about this solution is that it must be extended to every recognised logical consequence. From

(η) $$(m \leqslant n \,\&\, (\vdash_m A)) \to (\vdash_n A)$$

we have

(θ) $$(n = \max(m, k) \,\&\, (\vdash_m A) \,\&\, (\vdash_k C)) \to ((\vdash_n A) \,\&\, (\vdash_n C)),$$

while from (δ) we obtain

(ι) $$(\vdash_m A) \,\&\, (\vdash_k (A \to B)) \to \exists n\,(\vdash_n B).$$

We could in the same way complain that this committed us, whenever we had proved a statement A and had recognised some other statement B as being a consequence of A, to actually drawing that consequence some time in the future; and, if our interpretation of the operator '$\vdash_n$' is to be capable of dealing with this difficulty in the same way as with the special case of instances of a universally quantified statement, we should have to allow that a proof that a theorem had a certain consequence was, at the same time, a proof of that consequence, and, likewise, that a proof of a statement already known to have a certain consequence was, at the same time, a proof of that consequence; we should, that is, have to accept the law

(κ) $$(n = \max(m, k) \,\&\, (\vdash_m A) \,\&\, (\vdash_k (A \to B))) \to (\vdash_n B).$$

We should thus have so to construe the notion of proof that a proof of a statement is taken as simultaneously constituting a proof of anything that has already been recognised as a consequence of that statement. We can, no doubt, escape having to say that it is simultaneously a proof of whatever, in a platonistic sense, is as a matter of fact an intuitionistic consequence of the statement: but when are we to be said to have recognised that one statement is a consequence of another? If a proof of a universally quantified statement is simultaneously a proof of all its instances, it is difficult to see how we can avoid conceding that a demonstration of the validity of a schema of first-order predicate logic is simultaneously a demonstration of the truth of all its instances, or an acceptance of the induction schema simultaneously an acceptance of all cases of induction. The resulting notion of proof would be far removed indeed from actual mathematical experience, and could not be explained as no more than an idealisation of it.

The trouble with all this is that, as a representation of actual mathematical experience, we are operating with too simplified a notion of proof. The axiom (η) is acceptable in the sense that, prescinding from the occasional accident, once a theorem has been proved, it always remains *available* to be subsequently appealed to: but the idea that, having acknowledged the two premisses of a modus ponens, we have *thereby* recognised the truth of the conclusion, is plausible only in a case in which we are simultaneously bearing in mind the truth of the two premisses. To have once proved a statement is not thereafter to be continuously aware of its truth: if it were, then we should indeed always know the logical consequences of everything which we know, and should have no need of proof.

Acceptance of axiom (β) leads to the conclusion that we shall eventually prove every logical consequence of everything we prove. This, as a representation of the intuitionist notion of proof, is an improvement upon Beth trees, as normally presented: for these are set up in such a way that, at any stage (node), every logical consequence of statements true at that stage is already true; the Beth trees are adapted only to situations, such as those involving free choice sequences, where new information is coming in that is not derived from the information we have at earlier stages. But the idea that we shall eventually establish every logical consequence of everything we know is implausible and arbitrary: and it cannot be rescued by construing each proof as, implicitly, a proof also of the consequences of the statement proved, save at the cost of perverting the whole conception. If we wish to do so, there seems no reason why we should not take the stages represented by the numerical subscripts as punctuated by proofs, however short the stages thereby become, and the notion of proof as relating only to what is quite explicitly proved, so that, at each stage, one and only one new statement is proved, and consider what axioms hold under the resulting interpretation of the symbol '$\vdash_n$'. It thus appears that, under this interpretation, the axiom (β) must be rejected in favour of the weaker axiom (γ).

Looked at in another way, however, the stronger axiom (β) seems entirely acceptable. If, that is, we interpret the implication sign in its intuitionistic sense, the axiom merely says that, given a proof of A, we can effectively find a proof that A was proved at some stage; and this seems totally innocuous and banal. But, if axiom (β) is innocuous, how did we arrive at our earlier difficulties? The only possibility seems to be that our logical laws are themselves at fault. For instance, the law

(λ) $$\forall x\, A(x) \rightarrow A(m)$$

leads, via axiom (β), to the conclusion

(μ) $$\forall x\, A(x) \rightarrow \exists n\, (\vdash_n A(m)),$$

which appears, on the present interpretation of '$\vdash_n$', to say that we shall explicitly prove every instance of every universally quantified statement which we prove; so perhaps the error lies in the law (λ) itself. A law such as (λ) is ordinarily justified by saying that, given a proof of $\forall x\, A(x)$, we can, for each m, effectively find a proof of $A(m)$. If this is to remain a sufficient justification of (μ), then (μ) must be construed as saying that, given a proof of $\forall x\, A(x)$, we can effectively find a proof that $A(m)$ will be proved at

some stage. How can we do this, for given m? Obviously, by proving $A(m)$ and noting the stage at which we do so. This means, then, that the existentially quantified statement

(ν) $$\exists n\, (\vdash_n A(m))$$

is to be so understood that its assertion does not amount to a claim that we shall, as a matter of fact, prove $A(m)$ at some stage n, but only that we are capable of bringing it about that $A(m)$ is proved at some stage. Our difficulties thus appear to have arisen from understanding the existential quantifier in (β) in an excessively classical or realistic manner, namely as meaning that there will in fact be a stage n at which the statement is proved, rather than as meaning that we have an effective means, if we choose to apply it, of making it the case that there is such a stage. The point here is that it is not merely a question of interpreting the existential quantifier intuitionistically rather than classically in the sense that we can assert that there is a stage n at which a statement will be proved only if we have an effective means for identifying a particular such stage. Rather, if quantification over temporal stages is to be introduced into mathematical statements, then it must be treated like quantification over mathematical objects and mathematical constructions: the assertion that there is a stage n at which such-and-such will hold is justified provided that we possess an enduring capability of bringing about such a stage, regardless of whether we ever exercise this capability or not.

The confusions concerning the theory of the creative subject which we have been engaged in disentangling arose in part from a perfectly legitimate desire, to relate the intuitionistic truth of a mathematical statement with a use of the logical constants which is alien to intuitionistic mathematics. Troelstra's difficulties sprang from his desire to construe the expression '$\exists n\, (\vdash_n A)$' as meaning that A would in fact be proved at some stage: but, whether we interpret the existential quantifier classically or constructively, such a way of construing it fails to jibe with the way it and the other logical constants are construed within ordinary mathematical statements, and hence, however we try to modify our notion of a statement's being proved, we shall not obtain anything equivalent to the mathematical statement A itself. Nevertheless, the desire to express the condition for the intuitionistic truth of a mathematical statement in terms which do not presuppose an understanding of the intuitionistic logical constants as used within mathematical statements is entirely licit. Indeed, if it were impossible to do so, intuitionists would have no way of conveying to platonist mathematicians what it was that they were about: we should have a situation quite different from that which in fact obtains, namely one in which some people found it natural to extend basic computational mathematics in a classical direction, and others found it natural to extend it in an intuitionistic direction, and neither could gain a glimmering of what the other was at. That we are not in this situation is because intuitionists and platonists can find a common ground, namely statements, both mathematical and non-mathematical, which are, in the view of both, decidable, and about whose meaning there is therefore no serious dispute and which both sides agree obey a classical logic. Each party can, accordingly, by use of and reference to these unproblematic statements, explain to the other what his conception of meaning is for those mathematical statements which are in dispute. Such an explanation may not be accepted as legitimate by the other side (the whole point of the intuitionist position is that undecidable mathematical statements cannot legitimately be given a meaning by laying down truth-conditions for them in the platonistic manner): but at least the conception of meaning held by each party is not wholly opaque to the other.

This dispute between platonists and intuitionists is a dispute over whether or not a realist interpretation is legitimate for mathematical statements: and the situation I have just indicated is quite characteristic for disputes concerning the legitimacy of a realist interpretation of some class of statements, and is what allows a *dispute* to take place at all. Typically, in such a dispute there is some auxiliary class of statements about which both sides agree that a realist interpretation is possible (depending upon the grounds offered by the anti-realists for rejecting a realist interpretation for statements of the disputed class, this auxiliary class may or may not consist of statements agreed to be effectively decidable); and, typically, it is in terms of the truth-conditions of statements of this auxiliary class that the anti-realist frames his conception of meaning, his non-classical notion of truth, for statements of the disputed class, while the realist very often appeals to statements of the auxiliary class as providing an analogy for his conception of meaning for statements of the disputed class. Thus, when the dispute concerns statements about the future, statements about the present will form the auxiliary class; when it concerns statements about material objects, the auxiliary class will consist of sense-data statements; when the dispute concerns statements about character-traits, the auxiliary class will consist of statements about actual or hypothetical behaviour; and so on.

If the intuitionistic notion of truth for mathematical statements can be explained only by a Tarski-type truth-definition which takes for granted the meanings of the intuitionistic logical constants, then the intuitionist notion of truth, and hence of meaning, cannot be so much as conveyed to anyone who does not accept it already, and no debate between intuitionists and platonists is possible, because they cannot communicate with one another. It is therefore wholly legitimate, and, indeed, essential, to frame the condition for the intuitionistic truth of a mathematical statement in terms which are intelligible to a platonist and do not beg any questions, because they employ only notions which are not in dispute.

The obvious way to do this is to say that a mathematical statement is intuitionistically true if there exists an (intuitionistic) proof of it, where the existence of a proof does not consist in its platonic existence in a realm outside space and time, but in our actual possession of it. Such a notion of truth, obvious as it is, already departs at once from that supplied by the analogue of the Tarski-type truth-definition, since the predicate 'is true', thus explained, is significantly tensed: a statement not now true may later become true. For this reason, when 'true' is so construed, the schema (T) is incorrect: for the negation of the right-hand side of any instance will be a mathematical statement, while the negation of the left-hand side will be a non-mathematical statement, to the effect that we do not as yet possess a proof of a certain mathematical statement, and hence the two sides cannot be equivalent. We might, indeed, seek to restore the equivalence by replacing 'is true' on the left-hand side by 'is or will be true': but this would lead us back into the difficulties we encountered with the theory of the creative subject, and I shall not further explore it.

What does require exploration is the notion of proof being appealed to, and that also of the existence of a proof. It has often, and, I think, correctly, been held that the notion of proof needs to be specialised if it is to supply a non-circular account of the meanings of the intuitionistic logical constants. It is possible to see this by considering disjunction and existential quantification. The standard explanation of disjunction is that a construction is a proof of $A \vee B$ just in case it is a proof either of A or of B. Despite this, it is not normally considered legitimate to assert a disjunction, say in the course of a proof, only when we actually have a proof of one or other disjunct. For instance, it would be quite in order to assert that

$$10^{10^{10}} + 1 \text{ is either prime or composite}$$

without being able to say which alternative held good, and to derive some theorem by means of an argument by cases. What makes this legitimate, on the standard intuitionist view, is that we have a method which is in principle effective for deciding which of the two alternatives is correct: if we were to take the trouble to apply this method, the appeal to an argument by cases could be dispensed with. Generally speaking, therefore, if we take a statement as being true only when we actually possess a proof of it, an assertion of a disjunctive statement will not amount to a claim that it is true, but only to a claim that we have a means, effective in principle, for obtaining a proof of it. This means, however, that we have to distinguish between a proof proper, a proof in the sense of 'proof' used in the explanations of the logical constants, and a cogent argument. In the course of a cogent argument for the assertibility of a mathematical statement, a disjunction of which we do not possess an actual proof may be asserted, and an argument by cases based upon this disjunction. This argument will not itself be a proof, since any initial segment of a proof must again be a proof: it merely indicates an effective method by which we might obtain a proof of the theorem if we cared to apply it. We thus appear to require a distinction between a proof proper—a canonical proof—and the sort of argument which will normally appear in a mathematical article or textbook, an argument which we may call a 'demonstration'. A demonstration is just as cogent a ground for the assertion of its conclusion as is a canonical proof, and is related to it in this way: that a demonstration of a proposition provides an effective means for finding a canonical proof. But it is in terms of the notion of a canonical proof that the meanings of the logical constants are given. Exactly similar remarks apply to the existential quantifier.

There is some awkwardness about this way of looking at disjunction and existential quantification, namely in the divorce between the notions of truth and of assertibility. It might be replied that the significance of the act of assertion is not, in general, uniquely determined by the notion of truth: for instance, even when we take the notion of truth for mathematical statements as given, it still needs to be stipulated whether the assertion of a mathematical statement amounts to a claim to have a proof of it, or whether it may legitimately be based on what Polya calls a 'plausible argument' of a non-apodictic kind. (We can imagine people whose mathematics wholly resembles ours, save that they do not construe an assertion as embodying a claim to have more than a plausible argument.) It nevertheless remains that, if the truth of a mathematical statement consists in our possession of a canonical proof of it, while its assertion need be based on possession of no more than a demonstration, we are forced to embrace the awkward conclusion that it may be legitimate to assert a statement even though it is *known* not to be true. Still, if the sign of disjunction and the existential quantifier were the only logical constants whose explanation appeared to call for a distinction between canonical proofs and demonstrations, the distinction might be avoided altogether by modifying their explanations, to allow that a proof of a disjunction consisted in any construction of which we could

recognise that it would effectively yield a proof of one or other disjunct, and similarly for existential quantification: we should then be able to say that a statement could be asserted only when it was (known to be) true.

However, the distinction is unavoidable if the explanations of universal quantification, implication and negation are to escape circularity. The standard explanation of implication is that a proof of $A \rightarrow B$ is a construction of which we can recognise that, applied to any proof of A, it would yield a proof of B. It is plain that the notion of proof being used here cannot be one which admits unrestricted use of modus ponens: for, if it did, the explanation would be quite empty. We could admit anything we liked as constituting a proof of $A \rightarrow B$, and it would remain the case that, given such a proof, we had an effective method of converting any proof of A into a proof of B, namely by adding the proof of $A \rightarrow B$ and performing a single inference by modus ponens. Obviously, this is not what is intended: what is intended is that the proof of $A \rightarrow B$ should supply a means of converting a proof of A into a proof of B without appeal to modus ponens, at least, without appeal to any modus ponens containing $A \rightarrow B$ as a premiss. The kind of proof in terms of which the explanation of implication is being given is, therefore, one of a restricted kind. On the assumption that we have, or can effectively obtain, a proof of $A \rightarrow B$ of this restricted kind, an inference from $A \rightarrow B$ by modus ponens is justified, because it is in principle unnecessary. The same must, by parity of reasoning, hold good for any other application of modus ponens in the main (though not in any subordinate) deduction of any proof. Thus, if the intuitionistic explanation of implication is to escape, not merely circularity, but total vacuousness, there must be a restricted type of proof—canonical proof—in terms of which the explanation is given, and which does not admit modus ponens save in subordinate deductions. Arguments employing modus ponens will be perfectly valid and compelling, but they will, again, not be proofs in this restricted sense: they will be demonstrations, related to canonical proofs as supplying a means effective in principle for finding canonical proofs. Exactly similar remarks apply to universal quantification *vis-à-vis* universal instantiation and to negation *vis-à-vis* the rule ex falso quodlibet: the explanations of these operators presuppose a restricted type of proof in which the corresponding elimination rules do not occur within the main deduction.

What exactly the notion of a canonical proof amounts to is obscure. The deletion of elimination rules from a canonical proof suggests a comparison with the notion of a normalised deduction. On the other hand, Brouwer's celebrated remarks about fully analysed proofs in connection with the bar theorem do not suggest that such a proof is one from which unnecessary detours have been cut out—the proof of the bar theorem consists in great part in cutting out such detours from a proof taken already to be in 'fully analysed' form. Rather, Brouwer's idea appears to be that, in a fully analysed proof, all operations on which the proof depends will actually have been carried out. That is why such a proof may be an infinite structure: a proof of a universally quantified statement will be an operation which, applied to each natural number, will yield a proof of the corresponding instance; and, if this operation is carried out for each natural number, we shall have proofs of denumerably many statements. The conception of the mental construction which is the fully analysed proof as being an infinite structure must, of course, be interpreted in the light of the intuitionist view that all infinity is potential infinity: the mental construction consists of a grasp of general principles according to which any finite segment of the proof could be explicitly constructed. The direction of analysis runs counter to the direction of deduction; while one could not be convinced by an actually infinite proof-structure (because one would never reach the conclusion), one may be convinced by a potentially infinite one, because its infinity consists in our grasp of the principles governing its analysis. Indeed, it might reasonably be said that the standard intuitionistic meanings of the universal and conditional quantifiers involve that a proof is such a potentially infinite structure. Nevertheless, the notion of a fully analysed proof, that is, of the result of applying every operation involved in the proof, is far from clear, because it is obscure what the effect of the analysis would be on conditionals and negative statements. We can systematically display the results of applying the operation which constitutes a proof of a statement involving universal quantification over the natural numbers, because we can generate each natural number in sequence. But the corresponding application of the operation which constitutes the proof of a statement of the form $A \rightarrow B$ would consist in running through all putative canonical proofs of A and either showing, in each case, that it was not a proof of A, or transforming it into a proof of B; and, at least without a firm grasp upon the notion of a canonical proof, we have no idea how to generate all the possible candidates for being a proof of A.

The notion of canonical proof thus lies in some obscurity; and this state of affairs is not indefinitely tolerable, because, unless it is possible to find a coherent and relatively sharp explanation of the notion, the viability of the intuitionist explanations of the logical constants must remain in doubt. But, for present purposes, it does not matter just how the notion of canonical proof is to be explained; all that matters is that we require some distinction between canonical proofs and demonstrations, related to one another in the way that has been stated. Granted that such a distinction is necessary, there is no motivation for refusing to apply it to the case of disjunctions and existential statements.

Let us now ask whether we want the intuitionistic truth of a mathematical statement to consist in the existence of a canonical proof or of a demonstration. If by the 'existence' of a proof or demonstration we mean that we have actually explicitly carried one out, then either choice leaves us with certain counter-intuitive consequences. On either view, naturally, a valid rule of inference will not always lead from true premisses to a true conclusion, namely if we have not explicitly drawn the inference: this will always be so on any view which equates truth with our actual possession of some kind of proof. If we take the stricter line, and hold a statement to be true only when we possess a canonical proof of it, then, as we have seen, we shall have to allow that a statement may be asserted even though it is known not to be true. If, on the other hand, we allow that a statement is true when we possess merely a demonstration of it, then truth will not distribute over disjunction: we may possess a demonstration of $A \vee B$ without having a demonstration either of A or of B. Now, admittedly, once we have admitted a significant tense for the predicate 'is true', then, as we have noted, the schema (T) cannot be maintained as in all cases correct: but our instinct is to permit as little divergence from it as possible, and it is for this reason that we are uneasy about a notion of truth which is not distributive over disjunction or existnetial quantification.

A natural emendation is to relax slightly the requirement that a proof or demonstration should have been explicitly given. The question is how far we may consistently go along this path. If we say merely that a mathematical statement is true just in case we are aware that we have an effective means of obtaining a canonical proof of it, this will not be significantly different from equating truth with our actual possession of a demonstration. It might be allowed that there would be some cases when we had demonstrated the premisses of, say, an inference by modus ponens in which we were aware that we could draw the conclusion, though we had not quite explicitly done so; but there will naturally be others in which we were not aware of this, i.e. had not noticed it; if it were not so, we could never discover new demonstrations. It is therefore tempting to go one step further, and say that a statement is true provided that we are in fact in possession of a means of obtaining a canonical proof of it, whether or not we are aware of the fact. Would such a step be a betrayal of intuitionist principles?

In which cases would it be correct to say that we possess an effective means of finding a canonical proof of a statement, although we do not know that we have such a means? Unless we are to suppose that we can attain so sharp a notion of a canonical proof that it would be possible to enumerate effectively all putative such proofs of a given statement (the supposition whose implausibility causes our difficulty over the notion of a fully analysed proof), there is only one such case: that in which we possess a demonstration of a disjunctive or existential statement. Such a demonstration provides us with what we recognise as an effective means (in principle) for finding a canonical proof of the disjunctive or existential statement demonstrated. Such a canonical proof, when found, will be a proof of one or other disjunct, or of one instance of the existentially quantified statement: but we cannot, in general, tell which. For example, when $A(x)$ is a decidable predicate, the decision procedure constitutes a demonstration of the disjunction '$A(\bar{n}) \vee \neg A(\bar{n})$', for specific n; but, until we apply the procedure, we do not know which of the two disjuncts we can prove. It is very difficult for us to resist the temptation to suppose that there is already, unknown to us, a determinate answer to the question which of the two disjuncts we should obtain a proof of, were we to apply the decision procedure; that, for example, that it is already the case either that, if we were to test it out, we should find that $10^{10^{10}}+1$ is prime, or that, if we were to test it out, we should find that it was composite. What is involved here is the passage from a subjunctive conditional of the form:

$$A \rightarrow (B \vee C)$$

to a disjunction of subjunctive conditionals of the form

$$(A \rightarrow B) \vee (A \rightarrow C).$$

Where the conditional is interpreted intuitionistically, this transition is, of course, invalid: but the subjunctive conditional of natural language does not coincide with the conditional of intuitionistic mathematics. It is, indeed, the case that the transition is not in general valid for the subjunctive conditional of natural language either: but, when we reflect on the cases in which the inference fails, it is difficult to avoid thinking that the present case is not one of them.

There are two obvious kinds of counter-example to this form of inference for ordinary subjunctive conditionals: perhaps they are really two subvarieties of a single type. One is the case in which the antecedent A requires supplementation before it will yield a determinate one of the disjuncts B and C. For instance, we may safely agree that, if Fidel Castro were to meet President Carter, he would either insult him or speak politely to him; but it might not be determinately true, of either of those things, that he would do it, since it might depend upon some so far unspecified further condition, such as whether the meeting took place in Cuba or outside. Schematically, this kind of case is one in which we can assert:

$$A \rightarrow (B \vee C),$$
$$(A \,\&\, Q) \rightarrow B,$$
$$(A \,\&\, \neg Q) \rightarrow C,$$

but in which the subjunctive antecedent A neither implies nor presupposes either Q or its negation; in such a case, we cannot assert either $A \rightarrow B$ or $A \rightarrow C$. The other kind of counter-example is that in which we do not consider the disjuncts to be determined by anything at all: no supplementation of the antecedent would be sufficient to decide between them in advance. If that light-beam were to fall upon an atom, either it would assume a higher energy level, or it would remain in its ground state; but nothing can determine for certain in advance which would happen. Similar cases will arise, for those who believe in free will in the traditional sense, in respect of human actions.

If we were to carry out the decision procedure for determining the primality or otherwise of some specific large number N, we should either obtain the result that N is prime or obtain the result that N is composite. Is this, or is it not, a case in which we may conclude that it either holds good that, if we were to carry out the procedure, we should find that N is prime, or that, if we were to carry out the procedure, we should find that N is composite? The difficulty of resisting the conclusion that it is such a case stems from the fact that it does not display either of the characteristics found in the two readily admitted types of counter-example to the form of inference we are considering. No further circumstance could be relevant to the result of the procedure—this is part of what is meant by calling it a computation; and, since at each step the outcome of the procedure is determined, how can we deny that the overall outcome is determinate also?

If we yield to this line of thought, then we must hold that every statement formed by applying a decidable predicate to a specific natural number already has a definite truth-value, true or false, although we may not know it. And, if we hold this, it makes no difference whether we chose at the outset to say that natural numbers are creations of the human mind or that they are eternally existing abstract objects. Whichever we say, our decision how to interprét undecidable statements of number theory, and, in the first place, statements of the forms $\forall x A(x)$ and $\exists x A(x)$, where $A(x)$ is decidable, will be independent of our view about the ontological status of natural numbers. For, on this view of the truth of mathematical statements, each decidable number-theoretic statement will already be determinately true or false, independently of our knowledge, just as it is on a platonistic view; any thesis about the ontological character of natural numbers will then be quite irrelevant to the interpretation of the quantifiers. As we noted, it would be possible for someone to be prepared to regard natural numbers as timeless abstract objects, and to regard decidable predicates as being determinately true or false of them, and yet to be convinced by an argument of the first type, based on quite general considerations concerning meaning, that unbounded quantification over natural numbers was not an operation which in all cases preserved the property of possessing a determinate truth-value, and therefore to fall back upon a constructivist interpretation of it. Conversely, if someone who thought of the natural numbers as creations of human thought also believed, for the reasons just indicated, that each decidable predicate was determinately true or false of each of them, he might accept a classical interpretation of the quantifiers. He would do so if he was unconvinced by the general considerations about meaning which we reviewed, i.e., by the first type of argument for the adoption of an intuitionistic logic for mathematics: the fact that he was prepared to concede that the natural numbers come into existence only in virtue of our thinking about them would play no part in his reflections on the meanings of the quantifiers. Dedekind, who declared that mathematical structures are free creations of the human mind, but nevertheless appears to have construed statements about them in a wholly platonistic manner, may perhaps be an instance of just such a combination of ideas.

One who rejects the idea that there is already a determinate outcome for the application, to any specific case, of an effective procedure is, however, in a completely different position. If someone holds that the only acceptable sense in which a mathematical statement, even one that is effectively decidable, can be said to be true is that in which this means that we presently possess an actual proof or demonstration of it, then a classical interpretation of unbounded quantification over the natural numbers is simply unavailable to him. As is frequently remarked, the classical or platonistic conception is that such quantification represents an infinite conjunction or disjunction: the truth-value of the quantified statement is determined as the infinite sum or product of the truth-values of the denumerably many instances. Whether nor not this be regarded as an acceptable means of determining the meaning of these operators, the explanation presupposes that all the instances of the quantified statement themselves already possess determinate truth-values: if they do not, it is impossible to take the infinite sum or product of these. But if, for example, we do not hold that such a predicate as 'x is odd $\rightarrow$ x is not perfect' already has a determinate application to each natural number, though we do not know it, then it is just not open to us to think that, by attaching a quantifier to this predicate, we obtain a statement that is determinately true or false.

One question which we asked earlier was this: Can the thesis that natural numbers are creations of human thought be taken as a premiss for the adoption of an intuitionistic logic for number-theoretic statements? And another question was: What content can be given to the thesis that natural numbers are creations of human thought that does not prejudge the question what is the correct notion of truth for number-theoretic statements in general? The tentative answer which we gave to this latter question was that the thesis might be taken as relating to the appropriate notion of truth for a restricted class of number-theoretic statements, say numerical equations, or, more generally, decidable statements. From what we have said about the intuitionistic notion of truth for mathematical statements, it has now become apparent that there is one way in which the thesis that natural numbers are creations of the human mind might be taken, namely as relating precisely to the appropriate notion of truth for decidable statements of arithmetic, which would provide a ground for rejecting a platonistic interpretation of number-theoretic statements generally, without appeal to any general thesis concerning the notion of meaning. This way of taking the thesis would amount to holding that there is no notion of truth applicable even to numerical equations save that in which a statement is true when we have actually performed a computation (or effected a proof) which justifies that statement. Such a claim must rest, as we have seen, on the most resolute scepticism concerning subjunctive conditionals: it must deny that there exists any proposition which is now true about what the result of a computation which has not yet been performed would be if it were to be performed. Anyone who can hang on to a view as hard-headed as this has no temptation at all to accept a platonistic view of number-theoretic statements involving unbounded quantification: he has a rationale for an intuitionistic interpretation of them which rests upon considerations relating solely to mathematics, and demanding no extension to other realms of discourse (save in so far as the subjunctive conditional is involved in explanations of the meanings of statements in these other realms). But, for anyone who is not prepared to be quite as hard-headed as that, the route to a defence of an intuitionistic interpretation of mathematical statements which begins from the ontological status of mathematical objects is closed; the only path that he can take to this goal is that which I sketched at the outset: one turning on the answers given to general questions in the theory of meaning.