

PHILOSOPHICAL LECTURE

THE JUSTIFICATION OF DEDUCTION

By M. A. E. DUMMETT

Fellow of the Academy

Read 10 October 1973

THE standard practice of logicians, in treating of any well-defined fragment of logical theory, is to seek to define two parallel notions of logical consequence, one syntactic and the other semantic, and then attempt to establish a relation between them. The ideal is to establish their extensional equivalence. Proof of such equivalence falls into two parts, a soundness theorem showing that, whenever the syntactic relation obtains, so does the semantic one, and a completeness theorem, showing the converse inclusion. Failure of soundness yields a situation which must be remedied. Failure of completeness cannot always be remedied; a remedy is, however, mandatory wherever it is possible.

When either soundness or completeness fails, the remedy (when there is one) must of course be sought in a modification of either the syntactic or the semantic notion of logical consequence, and, on occasion, it might be the semantic notion rather than the syntactic one which had to be altered. Nevertheless, the semantic notion always has a certain priority: the definition of the syntactic relation is required to be responsible to the semantic relation, rather than the other way about. The syntactic relation is defined by devising a set of primitive rules of inference, and a corresponding notion of a formal deduction. If a semantic notion can be defined with respect to which a soundness proof can be given, we then have a reason for regarding the primitive rules of inference as valid: until then, we have only an intuitive impression of their validity. If the definition of the semantic relation has succeeded in its object of giving the intended meanings of the logical constants, then the fact that one of our rules of inference is not semantically valid shows that our intuition regarding it was unreliable. It is true that, on occasion, the discrepancy may prompt us to repudiate the proposed semantics. But, in such a case, we must always be able to give an independent reason for saying that this semantics did not succeed in

capturing the intended meanings of the constants: we will not abandon it merely because some rule of inference which appeared valid is invalid with respect to it. As far as completeness is concerned, our intuition gives us no assurance. There is no *a priori* reason why there should be any finite set of rules of inference by means of which every semantic consequence of a set of premisses may be derived from them; and, even when such a set exists, we can have no direct assurance, in advance of a completeness proof, that, by writing down all the rules of inference we could think of, we have arrived at such a set.

A soundness or completeness proof thus appears in the light of a justification of the definition of syntactic consequence. By means of a soundness proof, we demonstrate that the primitive rules of inference are in fact valid; by means of a completeness proof, that any valid inference may be effected by the iterated application of these rules. That is the natural way of understanding the standard approach to logical theory, and the one which is encouraged by the usual expositions of that theory.

It is not, however, the attitude most prevalent amongst philosophers. On the contrary, philosophers customarily assume that a justification of deduction is even more evidently impossible than a justification of induction, and for similar, though even more plainly cogent, reasons. There can, of course, be such a thing as a demonstration of the validity of some particular form of argument, namely the kind of demonstration which constitutes a non-trivial proof of syntactic validity: by use of rules of inference taken as primitive, the conclusion of a given form of argument may be shown to be derivable from its premisses. Such a demonstration will, of course, convince anyone who is willing to accept the primitive rules of inference as valid. Obviously, these methods cannot be applied indefinitely. If someone continues to question each rule of inference that is cited, we must eventually reach a point where we have some set of rules no one of which can be reduced to a series of applications of simpler ones, or, at least, of ones which we have not previously justified by appeal to those in the set we now have. At this stage, the only justification that would be possible would be one of a different kind—a semantic rather than a syntactic justification. However, in the view of such philosophers, a soundness proof for our primitive rules would, if offered as a justification of them, incur just the same charge of circularity as if we attempted to justify each of two sets of primitive rules by showing the derivability of each from the other. For, in demonstrating soundness, we should be

bound to employ deductive argument; and, in doing so, we should probably make use either of those very forms of inference which we were supposed to be justifying, or else of ones which we had already justified by reduction to our primitive rules. And, even if we did neither of these things, so that our proof was not strictly speaking circular, we should have used some principles of inference or other, and the question could then be raised what justified them: we should therefore either eventually be involved in circularity, or have embarked upon an infinite regress.

This view, which has been expressly advocated by Nelson Goodman, but probably represents the tacit attitude of most philosophers, has an obvious initial plausibility. What, then, on this view, is the significance or interest of a soundness or completeness proof? Many philosophers would evade the question by saying that such proofs have a 'merely technical' interest: but such a reply is devoid of any immediately comprehensible meaning. Logic is a technical subject in the sense that it employs techniques which need to be learned: but it is not a technical subject in the sense of engineering or agriculture, viz. one whose ends are unproblematic and are given from outside. Relative to a given proof or set of theorems, one can say that a certain notion or a certain formulation has a purely technical interest, meaning that it serves merely to facilitate the execution of that proof or the statement of those theorems: but a soundness or completeness theorem is not a lemma on the way to the proof of a more general theorem, but something that has an interest in its own right. (A completeness theorem may be used to derive a purely model-theoretic result, such as the compactness theorem: but, if this were its whole point, it would be a very uneconomical way of reaching that result.)

A more plausible account would be this: the syntactic notion of logical consequence is required for proving positive results, to the effect that such-and-such a form of argument is valid (because reducible to a number of simpler and intuitively valid steps); the semantic notion is required for negative results, to the effect that such-and-such a form of argument is invalid (because a counter-example can be found in which the premisses are true but the conclusion is not). In order to guarantee that a demonstration of semantic invalidity really does show that the argument in question cannot be reduced to the rules of inference we have taken as primitive, we require a soundness proof. In order to satisfy ourselves that the semantics we have adopted is

adequate in the sense that any form of argument not reducible to the primitive rules is semantically invalid in our sense, we need a completeness proof.

On this view, the syntactic notion would not be responsible to the semantic notion: rather, the converse relationship would obtain. We might regard the set of primitive rules of inference, together with the definition of a derivation employing them, as constitutive of our notion of logical consequence (within this particular area of logic). The interest of the semantic notion would then lie entirely in its use to demonstrate failure of logical consequence. The soundness proof would serve to show that semantic invalidity really did imply invalidity as defined by our set of primitive rules; the completeness proof to show that the device we were employing was adequate to its task.

This would be to reduce the semantic notion of logical consequence to a purely algebraic tool. We have examples of purely algebraic completeness. For instance, the topological interpretations of intuitionistic logic were developed before any connection was made between them and the intended meanings of the intuitionistic logical constants. Thus, intuitionistic sentential and predicate logic is complete with respect to the usual topology on the real line, under a suitable interpretation relative to that topology of the sentential operators and the quantifiers. No one would think of this as in any sense giving the meanings of the intuitionistic logical constants, because we have no idea what it would mean to assign to an actual statement, framed within first-order logic, a 'value' consisting of an open subset of the real line. Here it would be wholly in order to say that the interest of such a completeness proof, which I am calling algebraic as opposed to semantic, was purely technical. We have a mathematical characterization of the set of valid formulas of intuitionistic sentential or predicate logic, which may serve to establish certain general results about that set (for instance, that it contains $A \vee B$ if and only if it contains either A or B): as such, it has an advantage over the purely syntactic characterization in terms of an axiomatic formalization, though it is at a disadvantage as compared to a syntactic characterization in terms of a calculus of sequents.

We now have a position which would exactly correspond to the thesis that soundness and completeness proofs are of 'purely technical' interest, since it obliterates the distinction between a semantic notion of logical consequence, properly so called, and a merely algebraic one. Semantic notions are framed in terms of

concepts which are taken to have a direct relation to the use which is made of the sentences of a language; to take the most obvious example, the concepts of truth and falsity. It is for this reason that the semantic definition of the valuation of a formula under a given interpretation of its schematic letters is thought of as giving the meanings of the logical constants. Corresponding algebraic notions define a valuation as a purely mathematical object—an open set, or a natural number—which has no intrinsic connection with the uses of sentences. On the present view, the distinguishing feature of a semantic as opposed to an algebraic definition of logical consequence is a purely rhetorical flourish, not to be taken seriously. Thus nothing is lost, on this view, if, in the standard semantic treatment of classical sentential logic we replace the truth-values *true* and *false* by the numbers 0 and 1. The whole interest of the soundness and completeness proofs for classical sentential logic lies in the effective method they provide for determining whether or not a formula is derivable from some finite set of formulas: in so far as the words 'true' and 'false' are taken as being connected with the manner in which we effect communication by the use of sentences containing the classical sentential operators, the employment of these words in defining the semantic notion of logical consequence for formulas of classical sentential logic is quite unwarranted; all that we are concerned with is an algebraic device involving functions defined over a two-element set.

Such a position is coherent enough: what is wrong with it is that it simply lacks credibility. It is, indeed, open to argument, not merely whether, for example, the two-valued truth-tables give a correct account of the meanings of certain sentential operators of natural language, but whether they constitute a legitimate form for the explanations of the meanings of any possible sentential operators whatever; from the standpoint of intuitionistic mathematics, for instance, they do not, unless severely restricted as to the contexts in which the operators are permitted to occur. But what is not open to argument is that they purport to constitute such explanations. Specifically, the dispute over their legitimacy must concern the question whether we do or do not possess, for the sentences of our language, notions of truth and falsity such that to each particular utterance of any complete sentence one or other truth-value determinately attaches. This is a large and controversial question. Two things, however, are not controversial. First, if this question is to be answered negatively, then the truth-tables cannot be claimed to

provide at best more than a partial explanation of the meanings of the corresponding operators. And, secondly, if the question is to be answered affirmatively, then the truth-tables provide at least one legitimate way of explaining the meanings of certain possible sentential operators. On the assumption that all our sentences possess determinate truth-values, there is simply nothing that one can think of that a truth-table would leave unexplained concerning the meaning of the sentential operator for which it was correct. I do not propose to argue this here—it would take us too far into the intricate question of the relation between the notions of truth and falsity and that of meaning. Indeed, it might be objected to as not quite accurate: there is, after all, the well-known example, cited by Frege and many others after him, of ‘and’ and ‘but’, which share the same truth-table but differ in meaning. Frege distinguished two ingredients in meaning, *sense* and what in English we might call *tone*; the truth-table determined the sense of the connectives, which they therefore shared, and the residual difference was merely one of tone, the less important of the two ingredients of meaning. In order to make my remark accurate, it would be necessary to appeal to some similar differentiation between types of difference in meaning. Frege drew his distinction in terms of the notions of truth and falsity—a difference in tone could not affect the truth-value of a sentence, a difference in sense would, in general, do so. To make out that a distinction so drawn was genuinely a distinction in kinds of *meaning* would, again, require that we make clear the connection between truth-values and meaning, or part of it, which I have said I do not propose here to embark upon. I simply state it as intuitively obvious (*a*) that there is an important ingredient in meaning in respect of which ‘and’ and ‘but’ are equivalent, and (*b*) that, in respect of this ingredient, a truth-table must constitute a complete explanation of a sentential operator, provided that all sentences to which it is attached determinately possess one or other of the two values, true and false.

It is thus quite impossible that it should be an utter illusion that semantic accounts of the logical constants supply an explanation of their meanings, and that such accounts have no more significance than a purely algebraic characterization of a logical system which no one ever claimed as connected with the meanings of the constants. There is plenty of room for error: particular semantic accounts may be faulty in all sorts of ways. What is not conceivable is that we can rule out in advance the very possi-

bility of a semantics giving a model of meaning in just the way it is ordinarily supposed to do.

The situation is thus the reverse of what seems to be the case with induction. In the case of induction, we appear to have a quite unconvincing argument that there could not in principle be a justification, but we lack any candidate for a justification. In that of deduction, we have excellent candidates, in the soundness and completeness proofs, for arguments justifying particular logical systems; in the face of an apparently convincing argument that no such justification can exist.

The circularity that is alleged against any attempt to justify deduction, viz. to justify a whole system of deductive inference, is not of the usual kind. The validity of a particular form of inference is not a premiss for the semantic proof of its soundness; at worst, that form of inference is employed in the course of the proof. Now, clearly, a circularity of this form would be fatal if our task were to convince someone, who hesitates to accept inferences of this form, that it is in order to do so. But to conceive the problem of justification in this way is to misrepresent the position that we are in. Our problem is not to persuade anyone, not even ourselves, to employ deductive arguments: it is to find a satisfactory explanation of the role of such arguments in our use of language. An explanation often takes the form of constructing a deductive argument, the conclusion of which is a statement of the fact needing explanation: but, unlike what happens in a suasive argument, in an explanatory argument the epistemic direction may run counter to the direction of logical consequence. In a suasive argument, the epistemic direction must coincide with the consequential one: it is necessary that the premisses of the argument be propositions already regarded as true by the person whom we wish to persuade of the truth of the conclusion. Characteristically, in an explanation, the conclusion of the argument is given in advance; and it may well be that our only reason for believing the premisses of the explanatory argument is that they provide the most plausible explanation for the truth of the conclusion. Hence the charge of circularity or of begging the question is not applicable to an explanatory argument in the way that it is to a suasive argument. A philosopher who asks for a justification of the process of deductive reasoning is not seeking to be persuaded of its justifiability, but to be given an explanation of it. Admittedly, the situation is not as straightforward as that in which we have a proposition which we accept as true but want to know how it comes to be true: it is not plain

in advance just what is meant by saying that deductive reasoning is justifiable. We seek, simultaneously, an elucidation of that proposition and an explanatory argument showing what makes it true. Such an argument will, of course, be deductive in character, but that will not rob it of its explanatory power: we already engage in deductive reasoning, and therefore will be ready to admit that the conclusion of a deductive argument which strikes us as valid follows from its premisses; hence, in a suitable case, we shall also be ready to admit that the premisses of such an argument provide an explanation for the truth of the conclusion, even when the conclusion is to the effect that deductive reasoning is justified.

The charge of circularity thus fails to provide a short way with any attempt to justify deduction: but its failure does not show that any justification either is needed or can be provided. The phrase 'the justification of induction' has been scoffed at on the ground that it would be self-defeating to provide a justification of all inductive arguments, including unsound ones: the most that could be asked for is a justification of certain particular forms of inductive reasoning which we conceive to be sound. This is a very bad objection. A philosophical inquiry can begin with the query how there can be such an activity as mathematics, or as philosophy itself, or of what use or value such an activity is: it is no reply to dismiss the query by saying that bad mathematics or bad philosophy is of no use and no value. Obviously it is of no value: but that does not deprive of content the question what value mathematics, or philosophy, has *in general*, or force us to replace it by questions concerning particular mathematical theories or philosophical doctrines. And the case is similar with deduction. The question of justification arises at three levels. The first level is the unproblematic one: the case in which an argument may be validated by constructing a proof, in several steps, from its premisses to its conclusion by the use of simpler forms of inference which are admitted as valid. The second level is that which we also considered, where the correctness of a single basic form of inference, or of a whole systematization of a certain area of logic, is in question: and it is at this level that a proof of semantic soundness or completeness at least purports to provide a justification. But there is yet a third, deeper, level: that at which we require an explanation, not of why we should accept certain forms of argument or canons for judging forms of argument, but of how deductive argument is possible at all.

The existence of deductive inference is problematic because of

the tension between what seems necessary to account for its legitimacy and what seems necessary to account for its usefulness. For it to be legitimate, the process of recognizing the premisses as true must already have accomplished whatever is needed for the recognition of the truth of the conclusion; for it to be useful, a recognition of its truth need not actually have been accorded to the conclusion when it was accorded to the premisses. Of course, no definite contradiction stands in the way of satisfying these two requirements: recognizing the premisses as true involves a possibility of recognizing the conclusion as true, a possibility which will not in all cases be actualized. Yet it is a delicate matter so to describe the connection between premisses and conclusion as to display clearly the way in which both requirements are fulfilled. When we contemplate the simplest basic forms of inference, the gap between recognizing the truth of the premisses and recognizing that of the conclusion seems infinitesimal; but, when we contemplate the wealth and complexity of number-theoretic theorems which, by chains of such inferences, can be proved from the apparently simple set of Peano axioms, we are struck by the difficulty of establishing them and the surprises that they yield. We know, of course, that a man may walk from Paris to Rome, and yet that a single pace will not take him appreciably closer: but epistemic distance is more puzzling to us than spatial distance.

Another way of expressing the perplexity to which the existence of deductive inference gives rise is by asking how it can come about that we have an *indirect* means for recognizing the truth of a statement. Presumably the meaning that we assign to a statement (i.e. to the expressions of which it is composed) determines by what means the statement can be recognized as true. In some cases, the meaning of the statement may be such that an inferential process is necessarily involved in the recognition of it as true. Indeed, it is this insight which is one of the great contributions to the philosophy of language of Quine's celebrated essay *Two Dogmas of Empiricism*, and is there expressed by means of the image of language as an articulated structure of interconnected sentences, upon which experience impinges only at the periphery. The impact of experience may have the eventual effect of inducing us to assign (new) truth-values to sentences in the interior of the structure: but this impact will be mediated by truth-value assignments to other sentences which lie upon a path from the periphery, where the impact is initially felt, to the more centrally located sentences. This metaphor

presumably represents the entirely correct conception that, save for the peripheral sentences, the process of establishing a statement as true does not consist in a sequence of bare sense-perceptions, as on the logical-positivist model of the process of verification, but on the drawing of inferences (which need not, of course, all be strictly deductive) whose ultimate premisses will be based on observation. It is inherent in the meaning of such a sentence as 'The earth goes round the sun' or 'Plague is transmitted by rats' that it cannot be used as a direct report of observation (and thus is not, in Quine's image, located at the periphery of the linguistic structure), but can be established only on the basis of reasoning which takes its departure from what can be directly observed. In extreme cases, for instance, a numerical equation or the statement of the validity of a schema of first-order predicate logic, it is intrinsic to the meaning of the statement that it is to be established by purely linguistic operations, without appeal to observation at all (save the minimum necessary for the manipulation of the symbols themselves).

It is not in cases such as these that there is anything philosophically perplexing. Once we have freed ourselves from the positivist conception of the verification process as consisting in the mere occurrence of some sequence of sense-perceptions, there is no difficulty in acknowledging that it may be inherent in the meanings of certain sentences that some inferential process must enter into anything that will count as conclusively establishing their truth, or even that, in extreme cases, their verification will be exhausted by the production of such a chain of inference. These are the cases in which the most direct means of establishing the given statement as true will involve an inferential process; in which, in terms of Quine's image, the direction of transmission of the sequence of adjustments in truth-value assignments is from the periphery towards the interior. If it is implicit in the meaning of some statement that it can be established as true only by a process involving inference, then there is nothing philosophically puzzling about a chain of inference of the kind needed to establish it; in terms of Quine's image, the meaning of such a statement is determined by the links between it and other sentences adjacent to it in the direction of the periphery, and their meanings in turn by the links that connect them with further sentences yet closer to the periphery, and so on until we reach the observation statements which lie at the periphery itself. Equally devoid of any puzzling character are those extreme cases in which the whole meaning of the sentence

is given by reference to some procedure of proof or computation: if, for example, we consider numerical equations involving addition as given meaning solely in terms of the computation rules which decide their correctness or incorrectness (prescinding completely from their connection with the determination of the cardinality of sets of objects), then there can be nothing philosophically perplexing about the process of computation. But deductive inference does not proceed only in the direction from periphery to interior. It at least appears that chains of deductive reasoning occur which involve, either as premisses or as steps in the proof, statements which lie deeper in the interior than does the conclusion of the argument; even that the conclusion may, on occasion, be a peripheral sentence in the sense of one capable of being used to give a report of observation. In any such case, the conclusion of the deductive argument is being established indirectly, that is, by means of a process our understanding of which is not immediately involved in our grasp of the meaning of the statement. And, in the fact that this is possible, lies another facet of the philosophically puzzling character of deductive inference: how is it that, by means of such inferences, we can establish as true a statement that has not been directly so established, that is, which has not been so established by the means for which our method of conferring meaning on it expressly provides?

The problem so posed is not really distinct from the more general one enunciated previously: it is only the same tension between two features of deductive inference—that in virtue of which we want to say that it yields nothing new and that in virtue of which we want to say the opposite—that we earlier considered. We might put the problem in this way: when a statement is established, conclusively but indirectly, by the use of a deductive argument, in just what sense would it be right to say that, in accepting it as so established, we have remained faithful to the meaning we originally gave it? Or: in what sense, if any, can we say that, when it is established indirectly, we had already implicitly established it directly?

Philosophers have principally stressed, by dicta such as that the premisses contain the conclusion, that inference yields no new knowledge, that logic holds no surprises, and the like, the brevity of the gap between premisses and conclusion in a single inference: Frege, with his emphasis on the fruitfulness of deduction, and his refusal to treat analytically true statements as devoid of cognitive content, was exceptional in stressing the

contrasting feature of deductive inference. As Mill complains, however, few philosophers have made any serious attempt to relieve the tension between the two features: to resort to metaphor, as Frege did, and say that the conclusion is contained in the premisses 'as plants are contained in their seeds, not as beams are contained in a house' (*Grundlagen der Arithmetik*, § 88), is of no great help; we need to know how the metaphor is to be applied.

One of the very few to have attended to the problem of reconciling these two contrary features of deductive inference was Mill. Mill is frequently described as having contributed to the topic by advancing the thesis that every deductive inference is a *petitio principii*. Mill did indeed hold that thesis, but he did not regard it as his contribution to the subject: on the contrary, he cites several other writers as propounding the same thesis, and complains that, in doing so, they succeed in explaining the validity of deductive inference only at the cost of making it appear quite useless. What he took his contribution to be was his attempted explanation of how, while deductive inference really did involve a *petitio principii*, it was nevertheless useful.

Those of whom Mill complained relieved the tension between the two features of deductive inference by in effect repudiating that one of them which renders it fruitful. Wittgenstein, on the other hand, comes close in his *Remarks on the Foundations of Mathematics* to repudiating the other. In that book he held that a proof induces us to accept a new criterion for the truth of the conclusion. There is one sense in which this contention is both indisputable and banal. When the proof is given that a cylinder intersects a plane in an ellipse, we acquire a new criterion for a plane figure's being an ellipse: but, in the sense in which this claim is uncontentious, it does nothing to illuminate the nature or role of proof. For all his professed adherence to the maxim that a philosopher should only draw attention to what everybody knows but has overlooked, Wittgenstein did not intend his thesis to be merely trite: he meant to assert that, in accepting the proof, we have modified the meaning of the statement of the theorem, so that, in our example, the adoption of the new criterion for its application modifies the meaning that we attach to the predicate 'ellipse'. To speak of our accepting something new as a ground for applying a predicate as a modification of its meaning would not be, in itself, to go beyond what is banal, save in the use of the word 'meaning': to give substance to the thesis, we have to construe the modification as consisting, not

merely in our acceptance of the new criterion, but in the possibility of its yielding a different extension for the predicate from that yielded by the old criteria. The standard view of the effect of the proof is that the new criterion which it provides enables us to recognize as ellipses only figures which could already have been so recognized by the criteria we already had: what the proof establishes is precisely that the new criterion, where applicable, must always agree extensionally with that given by the original definition of 'ellipse'; that is why the proof persuades us to adopt the new test as a criterion. If Wittgenstein's thesis is to be more than a statement of the obvious, it must contradict this standard view: it must be understood as involving that there are, or may be, plane figures formed by the intersection of a cylinder with a plane which could not have been recognized as ellipses before the proof was given.

Such a position, whether it be the correct exegesis of Wittgenstein or not, is the reverse of that which Mill ascribes to his predecessors: it makes proof fruitful at the cost of robbing it of that feature which we take as making it compelling. It leaves unexplained the power of proof to induce us to change the meanings of expressions of our language in the way that it represents a proof as doing. On the ordinary view of proof, it is compelling just because, presented with it, we cannot resist the passage from premisses to conclusion without being unfaithful to the meanings we have already given to the expressions employed in it; whereas, on the view I have ascribed to Wittgenstein, its function is precisely to seduce us into such unfaithfulness. Given the view of Mill's predecessors, the puzzle becomes: What possible use is deductive inference? Given the Wittgensteinian view, it is: How does a proof achieve its effect? On any view which does not go to either extreme of denying altogether either that feature which gives deduction its value or that feature which renders it legitimate, the puzzle is, rather: How can any process possess both these characteristics at once?

One of the most obvious objections to the Wittgensteinian view, as I have stated it, is that a proof normally proceeds according to already accepted principles of inference, and that, therefore, by giving a proof we cannot be effecting any alteration of meaning, since the possibility of such a proof was, as it were, already provided for by our linguistic practice, namely by our acceptance of those principles of inference employed in the course of it. Such an objection, thus baldly stated, is based on a holistic view of language: the meaning of an individual

sentence is characterized by the totality of all possible ways that exist within the language for establishing its truth, including ones which involve deductive inference; we therefore cannot fully explain the meaning of an individual sentence without giving an account of the entire language of which it forms part, and, in particular, of all types of inference which might lead to it as conclusion. Even if such a holistic view be adopted, we may still ask whether the introduction of a new rule of inference would modify the meanings of existing sentences of the language: it will do so just in case it allows such sentences to be inferred from premisses from which they could not previously be inferred.

The introduction of such a new rule of inference might be held in itself to involve a modification in the meanings of sentences which are now open to being established as true in circumstances in which they could not previously have been so established. The alteration in meaning would, on such a view, be immediately consequent upon the introduction of the new rule, because a new possibility of establishing certain sentences as true had been introduced: the alteration in meaning would not wait upon the actualization of that possibility. Now, if that is held in fact to be the case with the system of deductive inferences which we now accept, then we arrive at a modification of the Wittgensteinian view, which has a great deal more plausibility. The objection which I just cited to the radical Wittgensteinian position was launched from a holistic position: just because the possibility of the proof was implicit in our existing practice, namely in our accepting the general principles of inference employed in it, the giving of an individual proof cannot be described as effecting a modification in the meaning of the conclusion. But the modified Wittgensteinian doctrine simply is a form of holism. It repudiates the molecular conception of language under which each sentence possesses an individual content which may be grasped without a knowledge of the entire language. Such a conception requires that we can imagine each sentence as retaining its content, as being used in exactly the same way as we now use it, even when belonging to some extremely fragmentary language, containing only the expressions which occur in it and others, of the same or lower levels, whose understanding is necessary to the understanding of these expressions: in such a fragmentary language, sentences of greater logical complexity than the given one would not occur. Our actual language would then be a conservative extension of the fragmentary language: we could not establish, by its use, any sentence of the fragmentary language which

could not already be established in that fragmentary language. The rules of inference which are applied in our language are, on such a molecular view, justified precisely by this fact, the fact, namely, that they remain faithful to the individual contents of the sentences which occur in any deduction carried out in accordance with them.

The modified Wittgensteinian view, which is tantamount to a holistic view of language, rejects this conception. According to it, we could find no way of ascribing an individual content to each sentence of the language which would do justice to the variety of possible ways in which the truth of a sentence might be established, i.e. which would not, in effect, destroy the validity of forms of inference which we are prepared to accept. In particular, we could not take those sentences of our language which fell below a certain level of logical complexity, and exhibit our whole language as a conservative extension of the fragment consisting of those sentences; in particular, we could not do this for those sentences not containing, explicitly or implicitly, any logical constants, that is, the atomic sentences. If we try to imagine our language as it might be if we had no expressions of generality or sentential operators, we could not so describe it that the introduction of these logical constants, and of the principles of inference governing them, would leave undisturbed the use of those atomic sentences: we should inevitably obtain cases in which an atomic sentence could be established in the full language, by means of inference, but could not have been established in the language without the logical constants.

This holistic conception of language no longer appears, like the radical Wittgensteinian view, to account for only one aspect of deductive inference. Whatever its defects as a philosophy of language in general, it has the great advantage of allaying the tension between the two features of deduction. On the holistic view, deduction is useful, because by means of it we can arrive at conclusions, even conclusions of the simplest logical form, which we could not arrive at otherwise. It is justified, simply because it is part of our overall linguistic practice. From a holistic standpoint, no specific ingredient of our general practice in the use of our language needs individual justification: it is justified simply by being part of that general practice. Thus, on such a view, a semantic proof of soundness or completeness can have, at best, a 'merely technical' interest, whatever that may be. This is in line with Wittgenstein's attitude to generally accepted forms of reasoning, namely that they are unassailable,

being, as they are, features of that use we have the right to choose to make of our sentences; it is likewise in line with Quine's preference for a syntactic to a semantic approach to logic, as when he says (in *Philosophy of Logic*) that the intuitionists would do better to rely upon a formalization of their logic rather than on any attempt to explain the meanings which they assign to the logical constants.

I said earlier that the situation in respect of deduction appeared to be the reverse of that in respect of induction: we have no cogent demonstration that there can be no justification of induction, but we lack any plausible candidate for such a justification; we have, on the other hand, plausible candidates for justifications, if not of the procedure of deductive inference in general, then at least for specific systematizations of logical deduction, in face of an apparently cogent argument that there can be no such justification. But, in one respect, the two cases are alike. For in neither case does the conviction that a justification is impossible by itself dispel the impression that a justification is called for. We wanted to know what entitled us to use one or other procedure for arriving at judgements as to the truth of statements. If we are persuaded that any attempt to give an answer to this question will involve us in vicious circularity, then we shall give up the inquiry; but our feeling that a ground of entitlement was needed will remain unassuaged. Holism, however, removes all desire to ask for a justification. We speak as we choose to speak, and our practice, in respect of the whole of our language, determines the meaning of each sentence belonging to it. Forms of deductive inference do not need to be faithful to the individual contents of the sentences which figure in the inference, because there is no individual content other than that determined by the language as a whole, of which those forms of inference are a feature. It is not, therefore, that there is something which must hold good of deductive inference, if it is to be justified, but which, because we should thereby be trapped in a vicious circle, we are unable to demonstrate, but must simply assume: rather, there is no condition whatever which a form of inference can be required to satisfy, and therefore nothing to be shown.

It is when holism is rejected, that is, when we suppose that each sentence may be represented as having a content of its own depending only upon its internal structure, and independent of the language in which it is embedded, that a justification of a system of deductive inference appears to be required. Of course,

even on a molecular view of this kind, no sentence can have a meaning which is independent of *all* the rest of the language. Its meaning depends upon the meanings of the constituent words, and these in turn depend upon the use of other sentences in which they may occur, and also of expressions of a lower level in terms of which they may be explained, or of the same level to which they are logically related: a grasp of the meaning of any sentence must, even on a molecular view of language, depend upon a mastery of some fragment of the language, a fragment which may, in some cases, be quite extensive. Nevertheless, it is essential to such a molecular view that there must be, for each sentence, a representation of its individual content which is independent of a description of the entire language to which the sentence belongs, and that we may distinguish among sentences according to their degree of complexity, where the representation of the meaning of any sentence never involves the representation of that of a sentence of greater complexity. A semantics for a logical theory always makes use of some general form of representation of the meanings of sentences. Since it is concerned only with the logical constants, it does not go beyond the *form* of such a representation: its application to specific sentences or types of sentence becomes the work of the theory of meaning. Thus the standard two-valued semantics for classical logic involves a conception under which to grasp the meaning of a sentence is to apprehend the conditions under which it is, or is not, true. If this is a correct general model for the meaning of any sentence of our language, then the sentential operators and the quantifiers can also be explained in accordance with this model, and the rules of inference governing them which are embodied in classical logic can be justified by reference to that representation of the meanings of the logical constants. The significance of a soundness or completeness proof, in terms of the two-valued semantics, for some systematization of logic depends, therefore, upon a thesis which does not belong to logic, and cannot be tested by it, but belongs, instead, to the theory of meaning: the thesis that the correct representation of meaning for expressions of our language is one given in terms of the truth-conditions of sentences. A Beth tree, on the other hand, considered as providing a semantics for intuitionistic logic, appeals to an alternative form of representation for the meaning of a sentence, namely in terms of the conditions under which it is recognized to have been established as true: once more, it is not part of logic to judge whether this is a correct model of meaning,

but for the theory of meaning if intuitionistic logic is to be considered as generally applicable, and for the philosophy of mathematics if its application is to be restricted to mathematical statements.

Thus it is only in the context of a molecular, of a non-holistic, philosophy of language that a proof of semantic soundness or completeness may be viewed as a justification of a logical theory; and it lies outside the proof itself, or the discipline to which it belongs, to judge whether the semantics in terms of which the proof is given is an acceptable one. But, in this case, we are still faced with the problems which confronted us before, namely: (1) Just what does the proof establish, in view of the fact that, construed as a suasive argument, it would be circular? And (2) how can the validity of deductive argument be reconciled with its utility?

Mill's brave attempt to resolve the second difficulty is a total failure, which is rendered the more difficult of assessment by his faulty analysis of his own chosen example. He holds, first, that, in any case in which someone knows, in the strict sense, the truth of the premisses of a valid deductive argument, he must already know the truth of the conclusion. An illustration, not Mill's, might be that of *modus tollendo ponens*, on the assumption, which must be incorrect from the standpoint of classical logic, that a strict knowledge of the truth of a disjunction must rest on a knowledge of the truth of one of the constituents. He holds, therefore, that an inference may represent a genuine epistemic advance only in a case in which at least one of the premisses of the deductive argument is believed but not strictly known to be true. Here, of course, we may readily agree that there is nothing problematic about a case in which one of the premisses is accepted on the testimony of another: if I am told that the disjunctive premiss of a *modus tollendo ponens* is true, and know the negative premiss by my own observation, then there is nothing puzzling about the fact that I can draw a conclusion which goes beyond both what I observed for myself and what I was told. We can extend this to the case when acceptance of one of the premisses rests on memory, since, in this respect, my memory is merely the testimony of my past self. But it is not this kind of case of which Mill is thinking, but, rather, that in which the ground for accepting one or more of the premisses, though short of conclusive, is, so to speak, the original one, and not derivative from that of another person or of one's former self. In such a case, Mill argues, it cannot be by means of the deduc-

tive inference that the epistemic step is taken, since, in asserting all the premisses, we have thereby already asserted the conclusion. Rather, the step is taken when we pass directly from the minor premiss or premisses to the conclusion, and this is not a deductive but an inductive inference. The deductive inference is of value, not as effecting the epistemic step, but, rather, as analysing the principle in accordance with which that step was taken, or as recording that to which we must be willing to assent if we are to regard that step as justifiable. For the major premiss enunciates the principle in accordance with which the inductive step was taken: we are justified in making the step from the minor premiss or premisses to the conclusion only in so far as we are justified in believing the major premiss to be true.

This account of the matter limps at every step. In the first place, it is impossible to see what can be meant by the contention that an assertion is effected by the making of two or more other assertions, when it would not be effected by the making of any one of them separately; and, if that contention were intelligible, it would apply as much to the unproblematic case in which one of the premisses rests upon testimony as to any other. In the second place, Mill's explanation works only for the case in which the judgement as to the truth of what is designated as the major premiss is subsequent to, or, at best, simultaneous with, the taking of the inductive step. In the case of *modus ponens*, for example, the cases favourable to Mill's account will be those in which someone finds himself disposed, upon learning the truth of *A*, to conclude to that of *B*, and reflects that this disposition can be warranted only in case 'If *A*, then *B*' is true, and that acquiescence in the disposition to which he feels himself inclined requires him to be prepared to assert the conditional statement. There may, indeed, be such cases; but, equally, there are cases in which the commitment to the truth of the conditional was made long in advance of the recognition of the truth of its antecedent. We could, of course, say that such a person had, by his assertion of the conditional, committed himself in advance to concluding to the truth of the consequent should he ever come to assert the antecedent: but that would be merely to say that, in asserting the conditional, he apprehended its deductive force; it would no longer be possible to deny, as Mill wants to do, that, in the epistemic advance which will later occur if the consequent is judged to be true in view of the truth of the antecedent, the judgement as to the truth of the conditional played a real role in that advance. The inability of Mill's account

to handle this case shows plainly in the lameness of his description of mathematical proof. Mill holds, of course, that the axioms of mathematical theories represent inductive generalizations; but he concedes that progress in mathematics consists principally, not in propounding new axioms, but in eliciting by deduction new consequences from those already accepted. But, in that case, the epistemic advance is effected deductively. Mill can in no way evade this conclusion by emphasizing the allegedly inductive basis for our acceptance of the axioms: for, when a new theorem is proved within an axiomatized theory, the axioms were already given, and supply the basis on which the epistemic step is being taken, rather than being arrived at by an analysis of that step.

In the third place, the assumption from which all else proceeds, that a knowledge, strictly so called, of the truth of the premisses of a deductive inference must involve a knowledge of that of the conclusion, is itself fallacious. To think thus is to overlook the complexity of the statements which we can express, and of the processes whereby we establish them as true, to think, in other words, as if each statement required, for its verification, no more than our mere exposure to the relevant sense-impressions. On the contrary, the verification of a statement will frequently demand the recognition of a pattern in what is observed, a pattern which, moreover, may not be accessible to direct inspection, but must be extracted by means of operations, of which counting and measuring are prototypical examples. Consider, as a representative case, Euler's famous solution of the problem of the bridges at Königsberg. Someone who knows Euler's proof can at once infer, from the information that a given person has, on a given day, crossed every bridge, that he has crossed at least one of the bridges at least twice. Mere vacant observation of the person in question, in the course of his peregrinations, would, of itself, assure the observer of neither proposition; and a procedure that established the one would not, by itself, of necessity immediately establish the other. To recognize either proposition as true would be to discern one or another pattern in the complex of perceptions which would make up observation of the entire walk. In a case such as this, either pattern might be noticed, with or without the other, by an observer whose attention was turned in the relevant direction, or might, just as easily, be overlooked; in a more complicated case, when the number of bridges was large, neither pattern would be detected without some operations other than those of mere observation and attention. The proof is convincing because

it displays a means whereby we can effectively transform any representation of the route by means of which the one pattern might be displayed into one by means of which the other could be displayed. Someone who has grasped the general procedure on which the proof depends could infer immediately from an observation sufficient to guarantee the truth of the statement that the person in question had crossed every bridge that he must have crossed some bridge twice, without a direct observation that this was so. Here, then, is a simple model of an indirect means, via deductive inference, to the recognition of a statement as true. The prior acceptance of the conditional statement, based upon the proof given by Euler, here plays an indispensable role in the actual process by which, in such a case, the conclusion is arrived at. The conditional was accepted, not because the proof showed that any verification of the antecedent would, of itself, already constitute a verification of the consequent, but because, rather, it provided a means whereby any sufficiently detailed observations which served to verify the antecedent could be rearranged so as to provide a verification of the consequent. The proof having been accepted, we are willing to proceed from an assertion of the antecedent, however based, to an assertion of the consequent, without necessarily carrying out that operation which the proof supplies which will lead to a direct verification of the statement we are inferring.

Let us at this stage return to the question which concerned us earlier, namely in what degree a semantic proof of soundness or completeness for a systematization of logical inference could be viewed as a justification of it. I distinguished three levels at which the problem of justification, for deductive inference, could arise. The first was the completely unproblematic one at which it is shown that a given form of argument can or cannot be reduced to a deduction within a given formal system by specified primitive rules of inference. The third and deepest level was that at which we ask how deductive inference is possible at all, the question, namely, to which Mill addressed himself, and to which I have returned a preliminary answer. The middle level remains: that at which we ask for a justification of a given set of canons for deductive inference, as embodied, say, within a formalization of some area of logic. We know perfectly well how such a justification may be provided; namely, by a demonstration, in semantic terms, of the soundness or completeness of the formalization. Our problem is whether, and, if so, with what title, such a demonstration may be said to supply a justification.

A sentence is a representation of some facet of reality. Our language—the matrix from which we form our sentences—has two roles: as a medium of communication, and as a vehicle of thought. What can be learned directly can also be communicated to us by others: the statements of others provide me with a vast extension of my own observational powers. But it is also by means of language that we are enabled to impose an order on reality as it is presented to us, to employ concepts whereby we can apprehend aspects of reality not apparent to gross observation. The theory of meaning, which lies at the foundation of the whole of philosophy, attempts to explain the way in which we contrive to represent reality by means of language. It does so by giving a model for the content of a sentence, its representative power. Holism is not, in this sense, a theory of meaning: it is the denial that a theory of meaning is possible. On a holistic view, no model for the individual content of a sentence can be given: we cannot grasp the representative power of any one sentence save by a complete grasp of the linguistic propensities underlying our use of the entire language; and, when we have such a grasp of the whole, there is no way in which this can be systematized so as to give us a clear view of the contribution of any particular part of the apparatus. No sentence can be considered as saying anything on its own: the smallest unit which can be taken as saying something is the totality of sentences believed, at any given time, to be true; and of what this complex totality says no representation is possible—we are part of the mechanism, and cannot view it from outside.

Difficult as is the task of attaining a satisfactory theory of meaning, such pessimism seems to me unwarranted at the present stage of inquiry and a discouragement to further progress. But, if a theory of meaning is possible, if it is possible, that is, to find a satisfactory model for the content of a sentence, and thereby to give an account of the means whereby we use language to represent reality, then we ought not to rest content with saying, of any feature of our linguistic practice, 'That is simply what we do'. Obviously, language must have many arbitrary features: things that are done one way could just as well have been done in some quite different way. But every functional feature of our linguistic practice must be capable of being explained in one of two ways: as contributing either to determining the content of our sentences or to effecting some operation with that content.

Prominent among the practices which make up our use of

language are those of deductive inference and deductive argument. Any satisfactory theory of meaning must, therefore, be able to relate these practices to the model of meaning which it employs: just this is what is done by a semantics for a logical theory.

So regarded, a proof of soundness or of completeness is a test, not so much of the logical theory to which it applies, but of the theory of meaning which underlies the semantics; naturally, this is only one test out of many which a theory of meaning must pass to be acceptable. In so far as the logical theory embodies our actual practice, that is, has primitive rules of inference which we in practice treat as valid, a theory of meaning, if it is to provide a model for our practice, must bring out those rules of inference as semantically valid, and should not bring out as semantically valid any rules which we cannot be brought to accept. There is here a complex interplay between semantic theory and intuitive practice. A semantics which can be shown not to justify a form of inference which is in standard use in ordinary discourse, or to justify one which we should unhesitatingly reject, is, by that fact, subject to criticism; although, even in such a case, we may be quite willing to accept such a semantics as providing a simplified version of some logical constant of everyday speech (as, e.g., many regard the logician's treatment of the conditional). On the other hand, with inferences of any complexity involving modal operators, tenses or higher-order quantification, our intuitions rapidly fail us, which is to say that no standard practice exists in respect of them; and, in such cases, we may be willing to accept as valid certain forms of inference, and to reject as invalid others, concerning neither of which we have any strong intuitions, simply because they are so determined by a semantics which works well for the simpler cases.

From these last remarks, it is apparent that semantic justifications of a logical theory do sometimes operate in a suasive manner, as inducing us to accept or reject certain forms of inference; and this role is often of genuine importance. Nevertheless, I remarked early on that the suasive function of a soundness or completeness proof is not that on which its deep significance depends, and certainly not primarily that in virtue of which we may refer to it as a justification of a logical theory. Rather, its importance lies in its providing for deductive inference what a theory of meaning must provide for every component of our practice in the use of our language, an understanding of the way it works: we seek, not merely a description of our

practice, but a grasp of how it functions. A semantics in terms of which a given fragment of logical theory can be proved to be sound, and, if that is possible at all, complete, supplies an answer to the question: How must our language be conceived to work—what model must we have for the meanings of our sentences—if the practice of deductive inference in which we engage is to be justified? It is not, in general, that we are in doubt as to whether that practice *is* justified: but, so long as we are unable to explain what the justification is, we lack an understanding of how our language works, of what it is that we are doing when we reason. Philosophy is an attempt to understand the world, as it is revealed to us both in our ordinary experience and by the discoveries and theories of science: and until we have achieved an understanding of our language, in terms of which we apprehend the world, and without which, therefore, there is for us no world, so long will our understanding of everything else be imperfect.¹

What a semantics for a logical theory has to be able to show is, first, that the rules of inference we ordinarily employ are in fact valid, that is, that they are justified in the sense that truth is preserved as we pass from premisses to conclusion. Just what this requirement involves will depend upon the semantics being employed; specifically, upon the notion of truth appropriate to that semantics. This is, of course, what is accomplished by a proof of soundness. But the other requirement which any successful account of deductive inference must satisfy, namely that it exhibit such inference as being useful as well as legitimate, must also be met by the semantics that is used, even though the demonstration of this is not ordinarily taken as a task for logic.

Now the question of the utility of deductive inference is, as we have seen, one that has many aspects. As we saw, the possibility of an epistemic advance by means of inference can never be problematic in the case when some of the information on which acceptance of the premisses rests is derived from secondary sources, e.g. from testimony. Nor can it be problematic when the ground of acceptance, though primary, is short of conclusive; here, rather, the doubt should relate, not to the utility of the inference, if justified, but to its justification. The fact that a form

¹ I do not say: without which there *would be* for us no world; dogs, sharks, etc., certainly inhabit *a* world. But our understanding of general features of our world cannot be separated from the understanding of the way in which we express those features. In this sense, all philosophy is a 'critique of language'.

of inference preserves truth does not guarantee that it preserves the level of probability: if certain premisses are accepted because they have a certain degree of probability, a degeneration of probability may well occur in the course of a chain of reasoning which is entirely valid, that is, certified as preserving truth. This fact supplies a rationale to those who, usually on incoherent grounds, distrust complicated chains of argument; it is a fact which is, to my mind, far too often overlooked, from the standpoint both of theory and of practice: but it is not what concerns us here. It is not enough that a model of meaning should allow us to recognize the utility of deductive inference only for cases when the premisses rest upon secondary or inconclusive evidence: an epistemic advance should be possible even in cases when the grounds are both primary and compelling.

Here, again, as we also saw, this splits into two cases, one of which is unproblematic: that in which the conclusion of the inference cannot be established save by an inference of just that form. On any possible view, it is part of the meaning of 'and' that a conjunction cannot be established save by establishing its two constituents; hence there can be no problem about the essential role of the rule of conjunction introduction in anything serving as conclusive grounds for a conjunctive statement. Of course, it will be a matter of the particular semantics adopted whether or not an introduction rule for some given logical constant really plays this role, e.g. whether the rule of existential generalization represents the only means whereby an existential statement may be conclusively established. But we may soften the requirement to allow for the case in which the general description of the means by which a statement of a given form can be established comprises, but is not exhausted by, the use of the relevant introduction rule. For instance, at least intuitionistically, the means by which a conditional is, in general, to be established includes, as a special case, that in which it is derived by if-introduction from a subordinate proof. These cases are very simple examples of what we noted as a general phenomenon, namely that the sense of many sentences is such that inference will play an indispensable role in anything which will count as a conclusive verification of them: there can therefore be nothing problematic about the employment of such inferences.

The problematic case is that in which a statement is established indirectly, even though conclusively, by means of an inference of a kind which is not provided for by a general characterization of the most direct means of verifying the

statement. The direct means of verifying the statement is that which corresponds, step by step, with the internal structure of the statement, in accordance with that model of meaning for the statement and its constituent expressions which is being employed. The possibility of establishing the statement directly must be envisaged by anyone who grasps the meaning of the statement, construed on this model: the possibility of establishing it by indirect means need not be; the indirect inference will involve elimination as well as introduction rules, and hence will involve also statements which do not belong to that fragment of language an understanding of which is essential to an understanding of the statement itself, statements which may therefore be of greater complexity than it.

The possibility of representing an epistemic advance as capable of being made by indirect means of this kind rests upon having a model of meaning which does not equate the truth of a statement with our explicit knowledge of its truth. For consider any case in which an epistemic advance is made by means of an inference of which the premisses may be considered to have been conclusively verified, but in which the conclusion might have, but has not, been established—directly—without the use of such an inference: the Königsberg bridge example will do as well as any other. For there to have been an epistemic advance, it is essential that the recognition of the truth of the premisses did not involve an explicit recognition of that of the conclusion—otherwise we shall be in Mill's difficulty. For the demonstration to be cogent, on the other hand, it is necessary that the passage from step to step involve a recognition of *truth* at each line. For the semantic proof of validity to have any force, that is, really to be a justification of the forms of inference used, this recognition of truth, in following out the demonstration, cannot *constitute* the truth of the statements so recognized: it must be a recognition of a property which is in accordance with the content of the statements, as given by the preferred model of meaning. It is quite different with a direct demonstration. The truth of a conjunction, for instance, simply consists in the truth of the premisses from which it is inferred by means of and-introduction, and so the recognition that it is true is not the recognition of a property which it had independently of the possibility of inferring it in that way.

It may be possible coherently to adopt a strongly idealist view, and equate the truth of a statement with its actual recognition as true, at least by indirect means. But, if epistemic advance by

indirect deductive inference is to be possible, truth must go beyond recognition of truth by *direct* means alone; while, if we are not to fall into holism, it must have some definite relation to the direct means whereby the truth of the statement can be established, since that direct means reflects the content of the statement according to the model of meaning we have adopted. In the case of mathematical statements, the relationship can, if we are disposed to do so, be taken to be as close as this: that a statement is to be recognized as true only if we possess an effective means in principle of establishing its truth by direct means. But, in the general case, we cannot demand a relationship as close as this: we should have, rather, to say that we possess an effective method for arriving at a direct verification of the statement, provided that we are given a sufficiently detailed set of observations. For instance, Euler's proof gives us an effective general means for finding, from any observation of the complete route which leads to a verification of the premiss, a verification of the conclusion: but, in a given case, we may have verified the premiss without having noticed or recorded the whole route in detail.

The relation of truth to the recognition of truth is the fundamental problem of the theory of meaning, or, what is the same thing, of metaphysics: for the question as to the nature of reality is also the question what is the appropriate notion of truth for the sentences of our language, or, again, how we represent reality by means of sentences. What I am affirming here is that the justifiability of deductive inference—the possibility of displaying it as both valid and useful—requires *some* gap between truth and its recognition; that is, it requires us to travel some distance, however small, along the path to realism, by allowing that a statement may be true when things are such as to make it possible for us to recognize it as true, even though we have not accorded it such recognition. Of course, from a realist standpoint, the gap is much wider: the most that can be said, from that standpoint, is that the truth of a statement involves the possibility in principle that it should be, or should have been, recognized as true by a being—not necessarily a human being—appropriately situated and with sufficient perceptual and intellectual powers.

On any molecular theory of meaning, the individual content of a sentence is determined by its internal structure, and relates, in the first place, to whatever constitutes the most direct means of recognizing it as true; on a realist theory, this direct means of

recognition of truth will often be inaccessible to us. Theories of meaning—rival types of semantics—thus differ, in the first instance, in what they represent as being the canonical means whereby the truth of sentences of various forms is to be established. Content, understood in this way, embodies the individual meaning of the sentence, and may be equated with Frege's sense. It is, in effect, cognitive content (when this is not taken relative to the existing knowledge of an individual); such content is *not* required to remain unamplified in the course of a valid chain of deductive inference, but, on the contrary, represents the respect in which such inference can lead to new knowledge. But, in view of the present thesis, that the utility of deduction requires a gap between truth and the recognition of truth by direct means, there is a further respect in which theories of meaning may differ: the notion of truth which they employ. For, in view of that thesis, when we know what constitutes the direct means of establishing a statement to be true, we do not yet know just what picture we need to have of what it is for it to *be* true, even though not established as true in this way. We have, nevertheless, to operate with *some* notion, however attenuated, of things being such as to make a given statement true, whether or not it has been recognized as true, at least by the most direct means. The distinction between these two aspects of a theory of meaning is hard to perceive in the case of a realist theory, just because such a theory operates with the notions of truth and falsity as its basic concepts: but it is in fact to this part of the theory that the assumption of bivalence, which gives it its realist character, belongs. That is: we could imagine a view under which each given statement of our language could be decided as true or as false, by direct means, by a being with sufficient powers and suitably situated, but under which it was nevertheless not held that each such statement is in fact determinately true or false, independently of actually being so decided. The notion of truth employed by a theory of meaning will yield a distinct notion of content—not that employed in this lecture—namely that in which the content of a sentence is a matter of how things have to be for the sentence to be true: in classical realist semantics, this becomes the set of possible worlds in which the sentence is true. This second type of content is that of which it may correctly be said that the condition for the validity of an inference is that content be not added to, or that analytically equivalent sentences have the same content; and it is this second feature of a theory of meaning which determines which forms of inference are validated by it.

These considerations place a restriction on the extent to which it is legitimate to demand that the language as a whole must be a conservative extension of a fragment of it formed by omitting certain expressions—for instance, some set of logical constants—together with the rules of inference governing them. It will be recalled that this condition was earlier stated as one that was necessary for the viability of a molecular view of language, necessary, that is, if holism is to be resisted. In the context of formalized languages, the notion of a conservative extension has a sharp sense, since we are there concerned with a single well-defined property of provability. But, when we consider natural language, there are several different epistemic degrees to which we might take the notion of a conservative extension as relative: and we have already seen that it would be illegitimate to demand that the language as a whole be a conservative extension of each significant fragment relative to any but the strongest of these degrees. Indeed, it cannot be taken as legitimate even relative to conclusive knowledge: if epistemic advance by means of indirect deductive inference is to be possible, then such inference will lead us to conclusions at which, in the actual circumstances, we could not have arrived without the employment of those modes of reasoning. The most that can be demanded is that the extension be conservative relative to the possibility of establishing a statement as true given a sufficiently detailed set of observations.

Indeed, it is not clear that appeal to the notion of a conservative extension is licit at all. The notion is, after all, originally a proof-theoretic one, and I am here stretching it into an epistemic one, whereas we are concerned in this discussion with semantic justifications of principles of inference. If we have a satisfactory semantic notion of truth, then whether or not the introduction of new vocabulary, subject to rules of inference, is a conservative extension of the language is something to which we can be indifferent: if, e.g., we have a language without any logical constants, but determinate truth-conditions for the atomic sentences, then it does not matter if the introduction of the logical constants, with rules governing them, allows us to infer the truth even of some of the atomic sentences in cases in which we could not have established them directly, so long as, in such cases, their truth-conditions are genuinely satisfied. The semantic notion becomes the standard, but our means of establishing truth something to be judged by that standard, not a standard in itself.

Or, rather, this is a misleading way of putting the matter. In

discussing the gap that must exist between truth and its recognition by direct means, I considered how small the gap might be made, and so looked at the topic from the standpoint of a strongly idealist or constructivist model of meaning. From the only standpoint which validates classical logic, the realist model of meaning in terms of truth-conditions, the gap is much wider. Understood on such a model, the condition for the truth of a sentence cannot, in general, be equated with even the possibility in principle of our knowing it to be true, however many observations we were able to make. Given such a model of meaning, there is no justice whatever in the idea that the language as a whole need be a conservative extension, relative to our recognition of truth, of any fragment of it. But what this means is that the model of meaning in terms of truth-conditions can be vindicated only by reference to the whole language. If we consider a fragment of natural language lacking the sentential operators, including negation, but containing sentences not effectively decidable by observation, it would be impossible for that fragment to display features embodying our recognition of the undecidable sentences as determinately true or false. The assumption of bivalence for such sentences shows itself only in the acceptance of certain forms of inference, classically but not intuitionistically valid. Hence it would be unsurprising if the introduction into the language of logical constants, treated as subject to the classical laws, rendered it possible for us, on occasion, to derive the truth of an atomic statement which could not have been recognized without the use of argument: and thus the extended language would not be a conservative extension of the original one relative to our recognition of truth.

What this means in turn is that, even if the two-valued semantics, the realist model of meaning in terms of truth-conditions, is required for the extended language, it was not required for the original fragment. So far as our use of the original, logic-free, language was concerned, there was no need to invoke a notion of truth going beyond the recognition of truth. The model in terms of truth-conditions indeed supplies a representation of the content of the atomic sentences, to which the classical logical laws are faithful; but it is a representation which was not called for by the linguistic practices which existed before the logical constants were introduced. A very clear case would be that of the past tense in a language in which there were no compound tenses, and in which the past tense, considered as an operator, could not be subjected to any of the ordinary logical constants:

in such a language nothing could reveal the assumption that each statement about the past was determinately either true or false.

It thus becomes conceivable that a certain model of meaning is required *only* in order to validate certain forms of inference the employment of which is part of our standard practice. That is, that model would be unnecessary in order to account for the use of that fragment of the language which contained only sentences of a low degree of logical complexity. Earlier, it was suggested that a molecular view of language required us to regard the meaning of each sentence as depending only upon the use of sentences in some quite restricted fragment of language, containing no sentences of greater complexity. But now we see that we may have to qualify this by saying that, although a theory of meaning of the kind aimed at on a molecular conception of language does ascribe an individual content to each sentence, it may be that the ascription could not be justified by reference to the use of the relevant fragment of language on its own, but only by reference to the behaviour of the sentence either as a constituent of more complex sentences or as figuring in inferences involving more complex sentences. And this would mean, therefore, that the meaning which, on such a model, we were taken as assigning to certain sentences, a meaning given in terms of their truth-conditions, was displayed only by our acceptance of certain forms of inference which could not otherwise be validated, rather than by anything involved in the use of those sentences as we learned it when, so to speak, they were on the frontier of the language we were acquiring (the frontier of complexity, that is).

It is just this which an opponent of a realist model of meaning finds incredible: he cannot believe that a grasp of a notion of truth transcending our capacities for its recognition can be acquired, and displayed, only by the acceptance of certain forms of reasoning. He concludes, instead, that these forms of reasoning, though generally accepted, are fallacious. I said earlier that a model of meaning is subject to criticism if it fails to provide a justification for forms of inference which it is part of our general linguistic practice to employ. The idealist (or constructivist) refuses to acknowledge this criticism as devastating: from his point of view, he has uncovered a defect in our language which ought to be corrected, the use of modes of inference which cannot be justified in terms of that model of meaning which fits our progressive acquisition of our language. For the realist, actual practice has to be explained, not

corrected, and his explanation is the only one which will fit the facts of that practice. He is left, of course, with a problem how to account for our acquisition of that grasp of conditions for a transcendent truth-value which he ascribes to us, and to make plausible that ascription. In the resolution of the conflict between these two views lies, as I see it, one of the most fundamental and intractable problems in the theory of meaning; indeed, in all philosophy.