

IF TRUTH IS DETHRONED, WHAT ROLE IS LEFT FOR IT?

BY JOHN CAMPBELL

Michael Dummett's investigation of a speaker's grasp of meaning, the role that this plays in grounding the validity of logical laws, and the relation of that role to realism, is one of the great works in the canon of twentieth-century philosophy. In this essay I want to sketch out some of his main lines of argument, express some reservations and very briefly suggest some alternatives.

Dummett has recently protested against the idea that a philosopher should be taken to be 'writing one long book' all his life. Nevertheless, in setting out what I take to be the main lines of his central arguments here, I have gone back and forth across his work. Dummett is an unusually systematic thinker, and it seems to me that the argument of one piece is often illuminated by the remarks in something else he has written. Nevertheless, I leave it to the reader to judge how far I have put together quite disparate ideas in setting out Dummett's main lines of thought.

I begin by endorsing Dummett's argument that a truth-conditional theory of meaning, thought of as giving the content of a speaker's knowledge of language, should take the form of a properly semantic theory rather than a truth theory. Then I look at the alternative, use-based account of meaning that Dummett recommends instead. This, I argue, ultimately provides us with an unrecognizable picture of our understanding of language. So I go back to look at the prospects for a properly semantic theory as giving the content of speaker knowledge. The prospects, I suggest, are better than Dummett allows.

1. Truth-Conditions and the Epistemology of Inference

In *The Logical Basis of Metaphysics*, Dummett distinguishes between two types of truth-conditional theories of meaning. First, a theory may state the conditions for the truth of a complex sentence in terms of the truth of its constituents. A Davidsonian truth theory does exactly this. In contrast, a properly semantic theory shows how the truth, or lack of truth, of the complex statement is determined by the truth or otherwise of its constituents. Simple examples of such a semantic theory are the ordinary truth-table explanations of the propositional constants. The general strategy of the book is to force the proponent of a truth-conditional theory to accept the properly semantic account, rather than, as many have attempted to do, clinging to the less ambitious description of meaning by truth theory. Once the truth-conditional theorist is forced into this more exposed position, the approach is subjected to a withering attack, and an alternative approach developed, one that describes meaning by describing the use to be made of statements.

At first the distinction between two types of truth-conditional theory can seem difficult to grasp. Suppose we characterize the meaning of a propositional constant by means of a truth theory. Then we will have, “‘A and B’ is true if and only if A is true and B is true’. What is the difference between this characterization of meaning and that given by a truth table? I think the key point here is that on Dummett’s view, the rules of inference governing the sign ‘and’ in the object language cannot be explained by the truth-theory. If the rules of inference for the language in which the theory is stated are classical, for instance, then the rules of inference for the object language will be classical too. The truth-theory simply projects the logic of the meta-

language onto the object language. Therefore, the description of meaning provided by the truth-theory cannot give any insight into why the rules of inference of the object language are right or wrong. In contrast, a truth-table describes the meaning of the constant as given by a function from the truth-values of the constituents to the truth-values of the whole. Whether an elementary introduction rule in the object language is correct will depend on whether it is the least demanding introduction rule that preserves truth from input to output. Whether an elementary elimination rule is correct will depend on whether it is the strongest possible rule that preserves truth from input to output. This kind of approach to the validation of inferences in the object language does not depend on projecting inferences from the meta-language into the object language. In fact, on this approach, the logic of the object language might be quite different to the logic of the meta-language.

The distinction between the two types of truth-conditional theory seems vivid when we consider modals. We can in a truth theory state the truth condition of ‘Necessarily, A’, by saying that it is true if and only if necessarily, A is true. That states the condition for the truth of ‘Necessarily, A’ in terms of the truth of A. But evidently it does not show that the truth or otherwise of ‘Necessarily, A’ is determined by the truth or otherwise of A. To show how A’s being true or otherwise depends on the semantic values of its components we need a semantic theory which assigns to A a semantic value, such as a set of possible worlds at which it is true, and the stipulation for ‘Necessarily, A’ that it is true if and only if A is true at all worlds. From the truth-theoretic stipulation that ‘Necessarily, A’ is true if and only if necessarily, A is true, we gain no understanding of the inferential link between ‘Necessarily, A’ and ‘Possibly, A’, considered as governed by the stipulation that it is true if and only if possibly, A is true. If a grasp of the meta-language already

provides a grasp of the inferential link, then it will indeed be recognized by the truth-theorist, but only the semantic theory will explain why the link exists.

The key point here again is that on a truth-theoretic approach the logic of the object language is maximally sensitive to the logic of the meta-language. The logical links that exist in the object language will be exactly those that exist in the meta-language. If the link between 'Necessarily' and 'Possibly' exists in the meta-language, it will be inherited by the object language; otherwise not. A semantic theory does not have this sensitivity. A semantic theory for a language with modal operators need not be stated in a language containing modal operators. A semantic theory stated in a classical meta-language need not ascribe a classical logic to the object language. A Kripke-tree semantics for intuitionistic logic may be stated in a classical meta-language.

This contrast between a truth theory and a semantic theory is forcefully drawn, but it does not yet explain why, if I am trying to characterize what a speaker knows when he knows a language, I should not settle for what a truth theory has to offer, while acknowledging that it falls short of being a semantic theory. Dummett says that if the notion of truth is to play an important role in a meaning-theory, then it seems reasonably clear that the meaning-theory must incorporate a semantic theory, then he draws the above contrast between a truth theory and a semantic theory. But it is not yet clear why the contrast matters. Perhaps it is true that a comprehensive account of 'how a language functions' will at some point have to include a semantic theory, in order to explain why certain inferential patterns obtain in the language, but it is not immediately obvious that anyone who knows that language must have that degree of insight into the functioning of the language. Perhaps this is simply one

point at which an account of the functioning of language comes apart from a theory of understanding.

One way to fill out the argument is this. One reason why grasping a language is described as 'knowledge' is that it increases one's cognitive powers in a number of ways. For example, having learnt a language, it becomes possible to find out things from the testimony of others, and this greatly increases the range of one's knowledge. But learning a language also makes deductive reasoning possible. It makes it possible to get from premises which one knows already to new knowledge, knowledge of conclusions which one did not know already. An account of a speaker's knowledge of language has to explain how it confers this epistemic ability on the speaker. How does the speaker know which sentences are and which are not logical consequences of others? If we describe the speaker's understanding by using a truth theory, we get no understanding of how the speaker has this knowledge. The speaker is ascribed knowledge of the theorems of the truth-theory, using a meta-language in which various inferential connections are assumed to hold, so the knowledge of the theorems of the truth-theory may be supposed to be sufficient for grasp of certain inferential connections, in that the speaker would not be said to grasp the theorems of the truth-theory unless he recognized those inferential connections. But this tells us only what the speaker knows, and gives us no insight whatever into how he knows it. In contrast, if we assume that the speaker knows the semantics of his language, we can see how it is that the speaker knows which inferential relations obtain. The knowledge depends on the speaker having some understanding of why those inferential relations obtain, and that just requires that the speaker should know a semantic theory for his language, rather than having his understanding exhaustively characterized by the truth theory.

I think that this line of thought must, indeed, be in the background in Dummett's discussion of the informativeness of deduction, the fact that deduction can yield new knowledge (pp. 195-199). Dummett's point here is that recognition of the validity of an inference often involves an ability to spot new patterns in a sentence one has already grasped as structured in one way. For example, one might grasp, 'If you do government business with the chairman, then the chairman will pay your bill' as having the structure 'If A then B'. So one will be able to infer 'the chairman will pay your bill' from 'you do government business with the chairman'. But one might as yet not realise that it is the same people being referred to in the antecedent as in the consequent. So one will not be in a position yet to grasp the inferential connection between 'If you do government business with the chairman, then the chairman will pay your bill' and 'Anyone who has his bill paid by someone with whom he is doing government business will be severely reprimanded.' These two imply that if you do government business with the chairman then you will be severely reprimanded. To grasp this inferential relation, you have to spot that it is the same people who are being referred to in the antecedent as in the consequent of the conditional. So inferences of the first type are trivial, but inferences of the second type, involving a capacity to spot new patterns in a structure already grasped, can be genuinely informative. This makes perfect sense if we suppose that the speaker grasps a semantic theory for his language. But on a truth-theoretic characterization of understanding, it is very hard to see what this distinction comes to. All that we know about the inferential patterns which the speaker ought to recognize in the object-language is that they are the very patterns which obtain in the meta-language, and since we are given no account of how the speaker knows these inferential patterns to obtain, an ability to spot new patterns in familiar sentences is being given no work to

do. So this account of how deduction can be informative, which actually does seem compelling, demands that the subject be thought of as grasping a semantic theory for his language.

If we accept that a truth-conditional meaning-theory ought to take a semantic theory rather than a truth-theory as its base, is there any objection to the resulting position? The fundamental question here is whether the semantic theory is correct. There is a broad sense in which, given any syntactically characterized set of patterns of inference, we can find some model-theoretic characterization or other of logical consequence that validates all and only those inferences. This is what Dummett calls a ‘purely algebraic’ definition of logical consequence, and the key point here is that such a model theory has no tendency to show that the patterns of inference that it validates are actually correct. The condition on the correctness of a semantic theory is that it should be possible to extend it to a meaning-theory for the language. And Dummett’s central point is that a classical semantic theory, which takes the semantic value of a sentence to consist in its being true or not, and so is committed to bivalence, cannot be extended to an adequate meaning-theory.

Dummett’s fundamental point here is that we cannot say what it is to have the conception of a statement being, determinately, true or not true. It cannot, in general, be explained as a matter of our being able to give a verbal explanation of meaning of the statement. But nor could having such a conception be a matter merely of being able to use the statement appropriately in response to the evidence of observation, even when reasoning from observation is needed to confirm or falsify the statement (pp. 301-321). I will return to this argument in §4 below. For the moment, having indicated what I take to be the main destructive thrust of Dummett’s argument, I want now to sketch his positive picture.

2. Harmony

Dummett proposes an alternative approach to the justification of logical laws, one that will again, though Dummett does not highlight the point, be rooted in the speaker's knowledge of language. The guiding idea here is that our use of language has ends which it may attain more or less successfully, and that we can criticize a way of using language on the grounds that it does not attain those ends (p. 210). And the primary reason why a use of language can be defective is that our total linguistic practice involves a multiplicity of principles for getting from one state to another, a multiplicity of transition rules. There are principles that we use in verifying statements, and principles that we use in drawing consequences from statements. And for the linguistic practice to be in order, all the various principles we use must be in harmony with one another.

It is important to emphasize that what gives the demand for harmony its normative status is its relation to our objectives in using language. A truth-conditional semanticist might have made a similar claim about his way of validating the use of particular rules of inference. The 'mythology of truth-conditions' was attempting to describe the objectives we have in our use of language. Our knowledge of what it is for a proposition to be true was thought to define the objective at which our ways of establishing the proposition aim. On the face of it, though, when Dummett describes the approach in terms of harmony, he describes it as one that tries to determine whether our ways of using language do indeed achieve our objectives in using language, without appealing to semantic notions at all. We can ask and answer

the question whether our ways of using language achieve our objectives, without using the notion of truth.

The question of harmony comes up in many ways. For example, we can ask whether there is harmony between what we might call different evidential perspectives. On the face of it, if you say ‘I see a tree’, you do so on a quite different basis from that on which I say of you, ‘You see a tree’. And the question arises whether these different ways of finding out are correctly related to one another. Or again, if on one day I say ‘Today it’s raining’, I do so on a quite different basis than that on which I say, the following week, ‘A week ago it rained’, and again we can ask about the harmony between these different evidential perspectives. And even if we stay with a single evidential perspective, we can ask whether there is harmony between the perceptual basis on which one makes a judgment, and the consequences for action that one takes the judgment to have.

In *The Logical Basis of Metaphysics*, Dummett’s concern is with a still further aspect of harmony: the demand for harmony between the introduction and elimination rules for a logical constant. Given a particular set of introduction rules, we do not want the elimination rules to allow us to derive unwarrantedly strong conclusions, but we do want them to allow us to derive all the conclusions we are entitled to. What is striking about Dummett’s approach is that he attempts to give a proof-theoretic approach to harmony. The central question this raises is whether we can explicate harmony in syntactic terms; again, as I said, without appealing to the notion of truth at all. On the face of it, indeed, there is a straightforward way to get this effect. Suppose we have an introduction rule for each logical constant. This tells us when a statement with that constant as its principal operator is justified, in terms of the conditions under which the sub-sentences are justified. Call an argument

canonical if it begins with atomic sentences as premises and uses only these introduction rules. Then we can say an arbitrary argument, perhaps with complex premises, is valid if, given canonical arguments for its premises, we could construct a canonical argument for its conclusion. Using this procedure, we could justify the elimination rules for the logical constants. For example, given the introduction rule for 'A&B', any canonical argument for 'A&B' will let us derive a canonical argument for A and a canonical argument for B, so the elimination rules are justified. We can also have a procedure that takes the elimination rules as basic. An arbitrary argument will then be valid if any atomic consequence that could be drawn from its conclusion, using only the elimination rules, could already be drawn from its premises, using only the elimination rules. We can use this idea to evaluate proposed introduction rules. For example, suppose we have the usual or-elimination rule, taken as given. Consider now the usual or-introduction rule. The upshot of any application of the introduction rule will be a statement with 'or' as its principal operator, to which we can apply the elimination rule. But any conclusion that we get by applying elimination rules here could also have been obtained by applying elimination rules to the premises of that application of the introduction rule; and that validates the introduction rule.

Consider now the case of classical negation. The elimination rule is, from not-not-A to infer A. That will validate the introduction rule, that from A we can infer not-not-A, and that is all; it does not show how to validate any other introduction rule. The introduction rule we want, for classical negation, is that if, from A as hypothesis, we can derive not-A, then we may discharge the hypothesis that A and conclude not-A. Can we take that introduction rule as given, and use it to justify the elimination rule? Evidently not: to obtain the premise of an application of

the elimination rule, not-not-A, we need to derive not-not-A from the hypothesis that not-A. But that does not of itself give us a way of deriving A by use of the introduction rules alone. So it does seem that by this criterion, the classical rules for negation are not in harmony with one another.

All that we have supplied so far, though, is the point that we cannot see from the rules themselves that classical negation is harmonious. So a speaker who simply grasps those rules does not thereby have the right to use them. The conclusion Dummett draws is only that these rules cannot be regarded as giving a meaning to classical negation. It remains a possibility that there are some empirical facts, relating to the meanings of the particular sentences to which negation is applied, in virtue of which it could be seen to be harmonious. In that case, a speaker who knew the rules and knew those empirical facts would indeed have the right to use classical negation, though it would have lost its status as a logical constant. But this is only an abstract possibility. It is not easy to see, even in outline, what the relevant empirical findings might be.

What exactly does this point establish? Does it show that classical negation should be abandoned? Conservativeness really does matter if we plan on giving only a syntactic explanation of a logical constant, if we are only going to give introduction and elimination rules for it in order to explain it, with no semantic foundation. A logical constant so understood is really just a trick to facilitate inference, and what we want to be sure of is that its use will do no harm; at the same time, we would like to be sure that we are getting as much as we can out of the device. The proof-theoretic account explains how to secure that. But this applies only to the case in which we are trying to rest with a purely syntactic account of the functioning of a constant.

One way of trying to rescue classical negation would be to find a purely proof-theoretic definition of harmony on which something like the classical rules can be seen to be after all harmonious. But the more fundamental question is whether we really can define our objectives in using negation in purely syntactic terms. What makes it seem compelling that classical negation is intelligible is not our shaky grasp of the possibilities for various types of proof-theoretic approaches, but the apparent intelligibility of the classical truth-table. The basic question here is whether it can really be right to think that our objectives in using language are properly described as merely a demand for harmony, without any appeal being made to the notion of truth. If we can dispense with the notion of truth in describing our objectives in using language, then the demand for harmony really does have the global significance that Dummett ascribes to it. If classical negation is not harmonious, then it cannot be explained in purely syntactic terms; everyone should agree to that. It is a more far-reaching claim to say that if classical negation is not harmonious, then it should not be used at all. This claim depends on the idea that our objectives in using language can indeed be explained in purely syntactic terms.

It is difficult to see how Dummett can avoid making this claim, that the normative demands of language can be characterized in purely syntactic terms. It seems to follow on from the abandonment of truth-conditional theories and the consequent insistence on an absolute demand for harmony in the rules of inference for logical constants. It is also quite difficult to see how Dummett can be right about this, for reasons that he more than anyone else made vivid. In a famous early discussion of the notion of truth, Dummett said, 'It is part of the concept of truth that we aim at making true assertions' (*TOE*, p. 2). The question here is how much weight that remark will bear. Does the appeal to the notion of truth do any work

whatever in describing the objectives that we have in using language? If so, then it cannot be right to give a syntactic conception of harmony the overpowering force that Dummett gives it in determining whether our uses of language serve their objectives, and we cannot dismiss the use of a logical constant simply on the grounds that it is not harmonious. On the other hand, suppose that the appeal to the notion of truth does no work in characterizing our objectives in using language. That is, suppose that we can give a complete normative description of the whole enterprise without appealing to the notion of truth. In that case the demand for harmony, syntactically characterized, can indeed be sustained. But we have lost what seemed to be one of Dummett's central insights on the notion of truth. It does not seem too strong to say that we have deprived language of its point.

3. Aiming at Truth

How does Dummett see the role of the notion of truth in an account of the functioning of language? He raises the question whether a statement, for a given language, of the conditions under which each of the sentences of that language is true, constitutes an explanation of the notion of truth for that language. He remarks that one might know all the clauses of a truth-definition of this type for a language, and still have no idea how to go about speaking the language. So it must be possible to explain how the truth of a sentence is related to the use we ordinarily make of it. A first response is to say that in assertion we aim at making true statements. But it must be possible, Dummett says, to explain why we go in for the practice of assertion at all; it must be possible to say what we are about in making assertions. And a theory of the

functioning of language, of this most general type, may be taken to explain what truth is.

This problem was first raised in Dummett's early article, 'Truth'. When we make a statement, we are characteristically trying to tell the truth. This means that an explanation of truth has to do more than simply assign every statement to one of two classes, the truths and the rest. It is not enough for the explanation to get it right about the extension of the concept of truth. It also has to make comprehensible the point of the concept, the way in which the classification relates to our interests. The beginning of the explanation is that we aim at truth in assertion. But we can further ask why that is what we do, why we engage in that activity of assertion.

There is a use of the notion of truth, in remarks such as 'Everything Bill said was true', which does seem capable of being explained by a Tarskian truth-definition, which simply specifies, for each sentence of the language, what it is for that sentence to be true. This use of the notion of truth does not involve any picture of the functioning of language; it is when we go beyond that use, and ask in general how the truth of a statement relates to the use we make of it, that we need the richer account of truth envisaged by Dummett.

In *Frege: Philosophy of Language*, the response Dummett gave to his question, 'What is the point of the practice of assertion?', emphasizes the role of assertion in communication. He holds that there are features of the assertions in any given language which are, by convention, the features that make them assertions, and the convention in question is 'to utter such sentences with the intention of uttering only true ones' (1973, p. 354). However, he continues, this description of the convention:

did not make clear the point of the convention, which can be stated in a general way as follows: we learn to react to the statements of others in the same way that we react to various observed features of our environment It is thus essential to the activity of making an assertion that the making of an assertion will in general modify the behaviour of those to whom it is made.

(1973, p. 355).

So far as I know, this is the only place where Dummett gives a direct reply to his famous question about the point of engaging in assertion. The problem with a disquotational theory was said to be that it did not characterize the point of our practice of assertion; that it missed the importance of the fact that in making assertions we aim at truth. This comment, that ‘we learn to react to the statements of others in the same way that we react to various observed features of our environment’, is what we are to use to fill in what the disquotational theory missed.

The striking thing about this comment is how little the appeal to truth now seems to contribute to the normativity of language-use. Suppose that we have two speakers of individual idiolects. We can describe the way each of them uses language as a pattern of use of sentences in response to observation. We can suppose that this pattern of use is described in entirely syntactic terms – at any rate, without having to appeal to semantic notions. And we can suppose that each speaker individually meets the normative conditions of harmony in use of the idiolect. So far, there has been no appeal to the notion of truth in characterizing the point of their practice. Suppose that now we introduce a disquotational notion of truth for each idiolect, and add the condition that each speaker is to aim at making only true assertions. This will not actually affect what either speaker does, since the pattern of use for sentences not

involving the word 'true' will be exactly as it was before. All that has been added is that each idiolect now has a communicative function, and presumably that the speakers should individually be aware that they will each be reacting to the statements of the other in the same way that they each react to 'various observed features' of the environment. If before the introduction of the notion of truth each speaker was operating simply as a manipulator of syntax, it is difficult to see that there has now been a radical change in the situation.

The pressure Dummett puts on the question, 'In what does knowledge of language consist?' inevitably leads to the view of the individual speaker as the determinant of the meanings of the signs. Certainly the syntactic account of harmony does not acknowledge any essential place for the social character of language. Suppose we ask just what is being conserved when we conservatively extend our practice. Well, what more is there to the practice than the rules governing the logical constants? A further factor is the connections that the sentences have to perception and action. Ultimately, what the proof-theoretic justification procedures are checking is that the use of a particular set of introduction and elimination rules will not disturb those connections that our sentences have to perception and action: that we will not be able to derive excessively strong consequences for action from a given range of perceptions, and that we will be able to derive all the appropriate consequences for action from those perceptions, rather than some unduly weak set. The presupposition here is that the connections between perception and action in our practice are all in order as they stand, and that the task of deduction is merely to respect them, while letting us see them more readily in particular cases. On this view of harmony, one's own perceptions and actions provide a basic perspective, which has to be conservatively extended by any new modes of thought.

Once we have dethroned the notions of truth and falsity from their fundamental place in the theory of meaning, and argued that the normative constraints on the use of sentences are provided by the syntactic criteria of harmony, we seem to be driven to this kind of picture of the functioning of language. Dummett is more alive than most to just how alien this picture of language use is to our ordinary understanding of the situation; I do not think that he would want to endorse this account. The trouble is that we face a dilemma at this point. One possibility is that we should return to the picture on which our understanding of the sentences of our language consists in knowledge of what it is for them to be true or not true, knowledge derived from our knowledge of the references of their parts. This would mean that we do not face the threat of having to suppose that our understanding of language is a matter entirely of syntactic manipulation. Alternatively, we can persist with the dethronement of truth and falsity, and describe merely the use that we make of sentences. But what role shall we find for the dethroned? The trouble is to see how there can now be any significant role for the notions of truth and falsity in setting the norms of language use. The best we can do is to consider the attenuated role envisaged by Dummett, on which they have to do merely with management of the social aspect of the use of language. There is no reason why a single speaker of an individual idiolect could not meet the criteria of harmony in use of that idiolect; talk about the notion of truth merely signals that we are envisaging the possibility of communication and the responsibilities it brings with it.

This does not seem to do justice to how fundamental the practice is of aiming at truth in assertion. I can bring this out by returning to Dummett's original talk of explaining how the practice is connected to our interests. One way of trying to give the explanation is to insist that aiming at truth is instrumental for other concerns that

we have; another approach is to say that it is something we do ‘for its own sake’, something which we can understand by locating it in a broader perspective of interdependent objectives we have, of which aiming at truth is one. At first sight these approaches seem to exhaust the alternatives, and the problem with them is that they leave it seeming that aiming at truth is simply one among many interests that we have. On these approaches, someone who did not go in for aiming at truth at all might be a bit deviant, like someone with no interest in music or in playing games, and we might try various ways of getting him or her to take up this life-enhancing practice. But if we failed, our subject would be, as it were, otherwise normal. The problem is that this does not really seem to be an intelligible scenario: there would be nothing left in respect of which our subject could be said to be ‘otherwise normal’. The problem is really to articulate that datum: to explain why aiming at truth is so fundamental to human life that without it, no distinctively human interests would be left.

In §1 I set out Dummett’s case for a truth-conditional theory of meaning to be a properly semantic theory, rather than merely a truth-theory. It seems to me that this case is compelling. I think therefore that we have to look again at the prospects for a properly semantic truth-conditional theory of meaning. The syntactic conception is not recognizable as an account of our ordinary understanding of language, and it does not seem to be possible to patch it up by finding supplementary supporting roles for the notions of truth and falsity.

4. The Objection to a Truth-Conditional Meaning Theory

Couldn't grasp of a classical conception of truth be displayed by our use of classical reasoning? Could it not be that the reason we ascribe grasp of a classical conception of truth to people is exactly that we want to explain their use of classical reasoning? So far as I have been able to see, Dummett does not explicitly address this natural response in *The Logical Basis of Metaphysics*. I think, though, that an answer is implicit. The fullest and most illuminating presentation of the line of thought here comes in an earlier article, 'The Justification of Deduction':

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.

(1978, p. 316)

The key point to notice here is that the conditional claim is persuasive. Suppose we did have a logic-free fragment of language. What would that fragment look like? Would it contain structured sentences? Typically when we discern structure in language we are discerning inferential structure; but here there are, on the face of it, no inferences to demand structure. It is hard to see how we could even discern singular reference in such a fragment: no complex predicates, no negation. We seem to be describing a 'feature-placing' fragment of language, containing only various unstructured ways of responding to the environment, various cries with which to greet the various features that occur, 'Hot!', 'Foggy!', and so on. Consider someone who

has only this unstructured language, and suppose for the moment, for the sake of argument, that some of these feature-placing sentences are not effectively decidable by observation. Could there be anything about such a speaker that would persuade you that the speaker understands those sentences in such a way that they are, determinately, either true or not true? It seems immediately evident that there could not be. Once you try to spell out why not, though, you are liable to be charged immediately with reductionism or behaviorism about meaning. The natural question to ask is: ‘What could there be in the use that someone made of an unstructured sentence such as ‘Frosty!’ that could make it so that he was understanding the sentence as not effectively decidable, and yet bound to be determinately, either true or not true?’. The idea that the speaker is interpreting the sentence as subject to bivalence seems evidently a fantasy, even though it is hard to spell out why without seeming to appeal to a behaviorist conception of understanding. Suppose, for the moment, we give Dummett the conditional claim, that the speaker grasping a logic-free fragment of language cannot understand these sentences as subject to bivalence. What follows from this? Evidently, we now have to concede that:

The assumption of bivalence for such sentences shows itself only in the acceptance of certain forms of inference, classically but not intuitionistically valid.

(1978, p. 316)

The classical truth-conditional semantics cannot be demanded for the logic-free fragment of language. It is not as if we can suppose that bivalence is already

established for the atomic sentences of a language and then look at what constraints this puts on the introduction of logical constants:

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 fragment 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 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.

(1978, p. 317)

Again, the conditional claim here seems compelling. Suppose we do consider a someone using a fragment of language in which they have only simple past-tense statements, such as 'Rained', and no complex tenses, no logical constants. Would there be anything in this person's understanding of such a fragment that could constitute grasp of bivalence for those simple statements? It again seems evident that this is a fantasy; any argument on this simply laboring the obvious. So why should

we ever appeal to the truth-conditional semantic theory, in which it is taken that every statement is determinately either true or not true?

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.

(1978, p. 317)

And now Dummett claims his conclusion:

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 those forms of reasoning, though generally accepted, are fallacious.

(1978, p. 318)

The idea is that the notions of semantic theory to which we appeal in explaining the validity of inferences should not be allowed to become posits, invoked to explain inference without having any independent grounding. It is not hard to sympathize with this concern. If we resist the idea of a purely algebraic approach to validity, on which any inference we like can be regarded as valid because there will be some notion of model available to underwrite it, do we not have to explain in what sense we can be said to know the semantic theory that validates correct inferences, otherwise than merely by making those inferences?

Nonetheless, it seems to me that the focus on a logic-free fragment of language pushes the point too far. The following pair of claims seems consistent:

(a) There is such a thing as grasp of a classical semantic theory for one's language, and it is one's grasp of a classical semantic theory that explains the correctness of the inferences in which one engages.

(b) Grasp of a classical semantic theory can be exercised only in connection with reasoning and inference. A subject who possessed only a logic-free fragment of language would have no use for semantic notions, and could not put them to work.

This way of putting the point does immediately invite the three questions on which Dummett focuses in *The Logical Basis of Metaphysics*. These are:

1. What is it to grasp a semantic theory for your language? This can't be identified with the capacity to engage in suitable reasoning; it is rather to be what makes it right for you to engage in those patterns of reasoning. Dummett's claim is that there is nothing we could find to constitute grasp of the references of the expressions of your language; all we find are ways of using sentences in response to experience.

2. Once we have characterized what it is to grasp the truth-conditions of sentences, we face the problem of explaining how the capacity to use those sentences can be derived from that grasp of truth-conditions. And, Dummett

says, there evidently is no way in which the ability to use a sentence could be derived from a grasp of truth-conditions.

3. Even if we get this far, we shall find that there is no way in which we could, ordinarily, know that you and I are associating the same references with the same signs. All we can observe is the use that we make of the signs. The alleged grasp of truth-conditions, lying behind the use that is made of the signs, cannot itself be observed; it is ultimately unknowable.

(1991, pp. 301-321)

We do have to answer those questions if we are to sustain the roles of truth and falsity as central to the theory of meaning.

5. Experience as Acquaintance

I will not here try to give a comprehensive response. I will try to indicate what I think the promising lines are. For the moment, though, I hope it is enough if I have indicated why I think we should persevere with truth-conditional semantic theories, and that Dummett's objections should be viewed as helpful impositions of constraints rather than as indicating fatal errors in the approach.

I think that a key point is one that generally goes unremarked: the role of experience in our understanding of language. Dummett, like almost everyone except Russell who writes on meaning, implicitly takes it that the only role experience can play in our understanding of language is to allow us to verify sentences about our

surroundings. But the role of experience may be to acquaint us with the references of our terms, just as Russell thought. There seems to be no obvious reason why the fundamental role of experience in understanding should not be to provide us with knowledge of reference, and ultimately knowledge of truth-conditions.

As I said, moreover, it does press Dummett's point too hard to insist that it must be possible to display grasp of the semantic theory in relation to grasp of a logic-free fragment of language. It may be that grasp of the references of the parts of a sentence can be exercised only in connection with reasoning and inference; that does not of itself mean that grasp of references has to be exhausted by engagement in a particular style of reasoning.

The same point applies when Dummett develops his second argument against the truth-conditional conception. Assuming that we associate classical truth-conditions with the sentences we use means, he says, that we cannot explain how the ability to recognize the truth of statements is derived from knowledge of their truth-conditions (pp. 306-311). However, an opponent may respond by putting pressure on the kind of 'derivation' in question here. What is required is not that the speaker should, without using any principles of reasoning or evidence, somehow extract the right principles to use in confirming a statement, from a grasp of the truth-condition of that statement. That would indeed be an impossible task, analogous to the task of providing what Dummett calls a 'suasive' justification of deduction (p. 202). Rather, all that can be asked is that the speaker's grasp of truth-conditions should provide what Dummett calls an 'explanatory' justification of those principles of confirmation. A truth-conditional theorist need not work with a picture on which the speaker is simply given knowledge of the truth-conditions of statements, with no idea how to verify them or draw implications from them, and then left to get on with it. Rather,

we should think of the speaker as being provided with a mass of techniques for reasoning, confirming, and acting on the basis of statements, together with a grasp of the truth-conditions of those statements. So long as the grasp of truth-conditions provides the speaker with an explanation of the point of those techniques, that is all the ‘derivation’ that is required.

Is the situation here that I am postulating various patterns of reasoning as being absolutely constitutive of conceptual thought, and that the notions of truth and reference have to be regarded as being indeed mere posits that we invoke to try to justify our use of these patterns of inference? Certainly not. It is part of the general picture I am recommending that there is a role for experience in our understanding of basic concepts. And the role of experience is to acquaint us with the semantic values of those concepts; truth and reference are not here mere posits. Our encounters with the references of our terms are as direct as any encounter could be.

This is particularly vivid in the case of perceptual demonstratives. There are some basic patterns of inference in which we use demonstratives, such as those involving identity and complex predicates, or existential generalization:

This thing is F

This thing is G

Hence, this thing is both F and G.

This thing is F

Hence, something is F.

The reference of the term ‘this’ is not a mere posit invoked to explain the validity of the inference; it is an object with which you are directly confronted. You understand the term ‘this’ only because you have experience of the object. And this provides you with an explanatory justification of the validity of the arguments above.

If we take the full force of the argument of Dummett’s that I outlined in §4 above, though, we would have to regard these patterns of inference as being justified by the fact that they provide conservative extensions of a more primitive, feature-replacing level of language. Dummett does indeed begin to sketch how such a program might go (1973, pp. 570-583). But from our present perspective this aim seems misplaced. We begin in the middle. There are certain patterns of reasoning we regard as fundamental, and we are acquainted with the semantic values of various basic concepts. This acquaintance with semantic values supplies an explanatory justification of the patterns of reasoning. There is no need to attempt to reconstruct the whole enterprise by asking whether it conservatively extends some more basic fragment of language.

The second line of response I want to suggest to Dummett’s argument against a truth-conditional semantic theory begins by remarking a certain solipsism implicit in the approach in terms of harmony. One way to bring this out is to consider what I earlier called the demand for harmony between different evidential perspectives. This comes up most sharply when we consider statements involving indexicals. For example, we want there to be harmony between the basis on which Jack says ‘I see a tree’, and the basis on which someone else says of Jack, ‘He sees a tree’. And we want there to be harmony between the basis on which one says on one day, ‘It is raining’, and the basis on which one says, a week later, ‘A week ago it rained’. To understand how there can be harmony here, we need the conception of a single world

on which there are many evidential perspectives: a single world on which one's current evidential perspective is simply one among many. And it is not just that this conception has to be something that the theorist can formulate: it must be available to the subject, if the subject is to have any grasp of what is going on in the co-ordination of his current evidential perspective with his own past evidential perspectives and other people's current evidential perspectives.

What allows us to transcend this solipsism is that we each of us have the conception of our own causal relation to the world, to each other and to our own past and future selves. We are not simply agents who intervene in the world; we know what we are doing and what is being done to us. This reflective grasp of agency provides you with an understanding of what you are about in your actions. It means you have an understanding of why you act as you do, and why in particular you observe the things that you do. You grasp that what you observe depends in part on what is there, waiting to be observed, and in part on your own position in the world, your own relation to it. Dummett has often stressed that having this idea of the world as there, waiting to be investigated, is an ingredient in our ordinary notion of truth. But this is not something to which we have to win our way by means of increasingly tenuous extensions from the primitive fragment of language in which there is no distinction between truth and verification. Rather, this reflective understanding of our own causal relation to the world is part of our most primitive grasp of concepts. Grasp of your causal relation to the world is constitutive of the capacity for conceptual thought.

Let me relate this line of thought to the project of explaining the point of our interest in truth; explaining why we aim at truth. Dummett's challenge was to explain the way in which the concept of truth relates to our interests. How should we do that?

It seems to be a precondition of our having any desires whatever that we also have some beliefs or other; desiring something would have no impact on your behavior unless you had some beliefs about how to bring it about. And surely the formation of beliefs involves aiming at truth. You might suggest that this gives a kind of transcendental account of the point of aiming at truth. It is not that it is connected with some interests that we happen to have, but rather that it is a precondition of our having any interests whatever. But if aiming at truth is given a role only as the necessary complement to desire in the mediation of perception and action, the point of aiming at truth is still being explained in terms only of the mediation of perception and action.

It seems to me that this notion of truth is actually all that we need for the judgments and beliefs of creatures have only first-order desires and no way of criticizing or correcting their desires. But this is not the ordinary human case. We do tell ourselves stories about what we are up to in our actions, stories which articulate a picture of our causal relations to the environment and each other, and which we use to explain, justify and correct our motivations. This capacity for causal narrative construction is really basic to ordinary human motivation. I want to suggest finally that we see the point of the ordinary notion of truth as its being something aiming at which is a precondition of this kind of narrative construction. So the point of the notion of truth is not that it meets some particular desires we have, nor that aiming at truth is a precondition of desire in general, but that it is a precondition of the kind of motivational structure that separates us from many other animals; the kind of motivational structure which makes us free agents.

At one point in *The Logical Basis of Metaphysics* Dummett speaks of himself as leading a mountaineering expedition which he plans to take no further than the foothills. I have been arguing that it is possible to complain about the direction in which we are being taken, that we are missing the best routes, and so on. It is consistent with that to acknowledge that the whole topography with which we are grappling – the mountains and the lakes, the glaciers and the swamps – all sprang into being in response to his probing.

REFERENCES

- Dummett, Michael. 1978. 'The Justification of Deduction'. In *Truth and Other Enigmas*, 290-318. London: Duckworth.
- Dummett, Michael. 1973. *Frege: Philosophy of Language*. London: Duckworth.
- Dummett, Michael. 1991. *The Logical Basis of Metaphysics*. London: Duckworth.