

Forthcoming in B. Leiter (ed.), *The Future for Philosophy*, Oxford: Oxford University Press

Past the Linguistic Turn?

Timothy Williamson

The Linguistic Turn is the title of an influential anthology edited by Richard Rorty, published in 1967. In his introduction, Rorty explained:

The purpose of the present volume is to provide materials for reflection on the most recent philosophical revolution, that of linguistic philosophy. I shall mean by “linguistic philosophy” the view that philosophical problems are problems which may be solved (or dissolved) either by reforming language, or by understanding more about the language we presently use. (1967: 3)

‘The linguistic turn’ has subsequently become a standard vague phrase for a diffuse event — some regard it as *the* event — in twentieth century philosophy, one not confined to signed-up linguistic philosophers in Rorty’s sense. For those who took the turn, language was somehow the central theme of philosophy. There is an increasingly widespread sense that the linguistic turn is past. In this essay I ask how far the turn has been, or should be, reversed.

1. A. J. Ayer, my predecessor but two in the Wykeham Chair of Logic at Oxford, was the first

of its holders to take the linguistic turn.¹ In 1936, back from Vienna but not yet in the Chair, he announced an uncompromisingly formal version of linguistic philosophy:

[T]he philosopher, as an analyst, is not directly concerned with the physical properties of things. He is concerned only with the way in which we speak about them. In other words, the propositions of philosophy are not factual, but linguistic in character — that is, they do not describe the behaviour of physical, or even mental, objects; they express definitions, or the formal consequences of definitions. (1936: 61-2).

Ayer traced his views back ultimately to the empiricism of Berkeley and Hume (1936: 11). His contrast between definitions of words and descriptions of objects is, roughly, the linguistic analogue of Hume's contrast between relations of ideas and matters of fact. For an empiricist, the *a priori* methods of philosophy cannot provide us with knowledge of synthetic truths about matters of fact ('the behaviour of physical, or even mental, objects'); they yield only analytic truths about relations of ideas ('definitions, or the formal consequences of definitions'). A rather traditional empiricism later overshadowed the linguistic theme in Ayer's work.

Ayer was the predecessor of Michael Dummett in the Wykeham Chair. Dummett gave a classic articulation of the linguistic turn, attributing it to Frege:

Only with Frege was the proper object of philosophy finally established: namely, first, that the goal of philosophy is the analysis of the structure of *thought*; secondly, that the study of *thought* is to be sharply distinguished from the study of the psychological process of *thinking*; and, finally, that the only proper method for analysing thought

consists in the analysis of *language*. [...] [T]he acceptance of these three tenets is common to the entire analytical school (1978: 458).

On this view, thought is essentially expressible (whether or not actually expressed) in a public language, which filters out the subjective noise, the merely psychological aspects of thinking, from the inter-subjective message, that which one thinks. Dummett's own corpus constitutes one of the most imposing monuments of analytic philosophy as so defined.

Elsewhere, Dummett makes it clear that he takes this concern with language to be what distinguishes 'analytical philosophy' from other schools (1993: 4). His account of its inception varies slightly. At one points, he says:

[A]nalytical philosophy was born when the 'linguistic turn' was taken. This was not, of course, taken uniformly by any group of philosophers at any one time: but the first clear example known to me occurs in Frege's *Die Grundlagen der Arithmetik* of 1884. (1993: 5).

Later, we read:

If we identify the linguistic turn as the starting-point of analytical philosophy proper, there can be no doubt that, to however great an extent Frege, Moore and Russell prepared the ground, the crucial step was taken by Wittgenstein in the *Tractatus Logico-Philosophicus* of 1922. (1993: 127).

Presumably, in Frege the linguistic turn was a fitful insight, in Wittgenstein, a systematic

conception.

That ‘analytical philosophers’ in Dummett’s sense coincide with those usually classified as such is not obvious. Some kind of linguistic turn occurred in much of what is usually called ‘continental [supposedly non-analytic] philosophy’. It is not obvious that Jacques Derrida does not subscribe in his own way to Dummett’s three tenets: if some stretching of terms is required, it is for the later Wittgenstein too. Conversely, Bertrand Russell did not subscribe to the three tenets, although often cited as a paradigm ‘analytical philosopher’. Over the last twenty years, fewer and fewer of those who would accept the label ‘analytic philosophy’ for their work would also claim to take the linguistic turn (I am not one of those few). Even philosophers strongly influenced by Dummett, such as Gareth Evans, Christopher Peacocke and John Campbell, no longer give language the central role that he describes. For Dummett, they belong to a tradition that has grown out of ‘analytical philosophy’ without themselves being ‘analytical philosophers’ (1993: 4-5). In effect, they aim to analyse thought directly without taking a diversion through the analysis of language.

The philosophy of mind has famously displaced the philosophy of language at the centre of much current debate. This is hardly a form of the linguistic turn, even granted the importance of Jerry Fodor’s notion of a language of thought (the brain’s computational code) in the philosophy of mind (Fodor 1975). However, the notion of a *mental representation* is central to the new philosophy of mind. A concept is a mental representation in this sense, whether or not it corresponds to an expression in a language of thought. One might therefore classify both thought and language together under the more general category of representation, and argue that the linguistic turn was just the first phase of the *representational turn*, on which the goal of philosophy is the analysis (in a generous sense) of representation. For the classification to be appropriate, we need a generous sense of

‘representation’ too, correlative with a loose notion of *aboutness*. We think about things and talk about them. We represent how things are when we know or believe or assert that they are some way; we represent how things are to be when we intentionally bring it about or hope or ask someone to bring it about that they are that way. Things are or are not as they are represented to be. In both thought and language, questions of truth and falsity arise: even if there is some thinking or speaking neither truly nor falsely, there would be no thinking or speaking at all if there were no thinking or speaking truly or falsely.² On some views, perception involves non-conceptual representations of one’s environment: they too raise the question of misrepresentation. As non-conceptual, they presumably fall outside Dummett’s category of thought; they threaten even his first tenet, that the goal of philosophy is the analysis of the structure of thought. Nor is it clear how far contemporary philosophy of mind accepts his second tenet, for it does not always sharply distinguish the study of thought from the study of the psychological process of thinking. *Naturalists* hold that everything is part of the natural world, and should be studied as such; how should one study thought as part of the natural world if not by studying the psychological process of thinking?

We could have brought thought and language together as forms of representation in other words by saying that both are forms of *intentionality*. This terminology reminds us how little the representational turn is confined to what would ordinarily be called ‘analytic philosophy’. The phenomenological tradition may constitute another form of the representational turn. In the hermeneutic study of interpretation and various shades of postmodernist discourse about discourse the representational turn takes a more specifically linguistic form.

Have we stretched our terms so far that it is vacuous to say that a piece of philosophy takes the representational turn? No. What language and thought have most obviously in

common is that they are both manifestations of *mind*. If we reject idealism in all its forms, we take mind to constitute only a small fraction of reality. It is no platitude to claim that the goal of philosophy is to analyse manifestations of that small fraction. To put it very schematically, let *idealism about the subject-matter of philosophy* be the view that what philosophy studies is mind, in contrast to *ontological idealism*, the view that what exists is mind. Although idealism about the subject-matter of philosophy does not entail ontological idealism, it is unclear why we should accept idealism about the subject-matter of philosophy if we reject ontological idealism. Of course, one might reject idealism about the subject-matter of philosophy while nevertheless holding that the correct method for philosophy is to study its not purely mental subject-matter by studying mental or linguistic representations of that subject-matter. This methodological claim will be considered later; for present purposes, we merely note how much weaker it is than those formulated by Ayer and Dummett.

The statement that mind constitutes only a small fraction of reality may be accused of violating Dummett's second tenet by confusing thought with the process of thinking. Almost everyone agrees that psychological events constitute only a small fraction of reality, but that is not yet to concede that thought in a non-psychologistic sense is similarly confined. John McDowell, for instance, argues:³

[T]here is no ontological gap between the sort of thing one can mean, or generally the sort of thing one can think, and the sort of thing that can be the case. When one thinks truly, what one thinks *is* what is the case. So since the world is everything that is the case [...] there is no gap between thought, as such, and the world. Of course thought can be distanced from the world by being false, but there is no distance from the world implicit in the very idea of thought. (1994: 27).

For McDowell, the sort of thing one can think is a conceptual content: the conceptual has no outer boundary beyond which lies unconceptualized reality. He denies the accusation of idealism on the grounds that he is not committed to any contentious thesis of mind-dependence.

The sort of thing that can be the case is that a certain object has a certain property. McDowell's claim is not that the object and the property *are* concepts; it is merely that we can in principle form concepts *of* them, with which to think that the object has the property. Indeed, we can in principle form many different concepts of them: we can think of the same object as Hesperus or as Phosphorus. In Fregean terms, different senses determine the same reference. McDowell admits 'an alignment of minds with the realm of sense, not with the realm of reference' (1994: 179). For objects, his claim that the conceptual is unbounded amounts to the claim that any object can be thought of. Likewise for the sort of thing that can be the case: the claim is, for example, that whenever an object has a property, it can be thought, of the object and the property, that the former has the latter. But, on a coherent and natural reading of 'the sort of thing that can be the case', such things are individuated coarsely, by the objects, properties and relations that they involve. Thus, since Hesperus *is* Phosphorus, what is the case if Hesperus is bright *is* what is the case if Phosphorus is bright: the objects are the same, as are the properties. On this reading, McDowell's claim 'When one thinks truly, what one thinks *is* what is the case' is false, because what one thinks is individuated at the level of sense while what is the case is individuated at the level of reference. Although McDowell's claim is true on other readings, it is not clear that they will bear the weight that his argument puts on them.

McDowell's argument seems to require the premise that everything (object, property, relation, state of affairs, ...) is thinkable. That premise is highly contentious. What reason

have we to assume that reality does not contain *elusive objects*, incapable in principle of being individually thought of? Although we can think of them collectively—for example, as elusive objects—that is not to single out any one of them in thought. Can we be sure that ordinary material objects are not constituted of clouds of elusive sub-sub-atomic particles? We might know them by their collective effects while unable to think of any single one of them. Of course, McDowell does not intend the conceptual to be limited by the merely medical limitations of human beings, but the elusiveness might run deeper than that: the nature of the objects might preclude the kind of separable causal interaction with complex beings that isolating them in thought would require. In Fregean terminology again, a sense is a mode of presentation of a referent; a mode of presentation of something is a way of presenting it to a possible thinker, if not an actual one; for all that McDowell has shown, there may be necessary limitations on all possible thinkers.⁴ We do not know whether there are elusive objects. It is unclear what would motivate the claim that there are none, if not some form of idealism. We should adopt no conception of philosophy that on methodological grounds excludes elusive objects.⁵

Suppose, for the sake of argument, that there are no elusive objects. That by itself would still not vindicate a restriction of philosophy to the conceptual, the realm of sense or thought. The practitioners of any discipline have thoughts and communicate them, but they are rarely studying those very thoughts: rather, they are studying what their thoughts are about. Most thoughts are not about thoughts. To make philosophy the study of thought is to insist that philosophers' thoughts should be about thoughts. It is not obvious why philosophers should accept that restriction.

Biology and physics are not studies of thought. In their most theoretical reaches, they merge into the philosophy of biology and the philosophy of physics. Why then should

philosophers of biology and philosophers of physics study only thought? Sometimes they do study biologists' and physicists' thought, but sometimes they study what that thought is about, in an abstract and general manner. Why should such activities not count as philosophy?

There is a more central example. Much contemporary metaphysics is not primarily concerned with thought or language at all. Its goal is to discover what fundamental kinds of things there are and what properties and relations they have, not how we represent them. It studies substances and essences, universals and particulars, space and time, possibility and necessity. Although nominalist or conceptualist reductions of all these matters have been attempted, such theories have no methodological priority and often turn out to do no justice to what they attempt to reduce.

The usual stories about the history of twentieth-century philosophy fail to fit much of the liveliest, exactest and most creative achievements of the final third of that century: the revival of metaphysical theorizing, realist in spirit, often speculative, often commonsensical, associated with Saul Kripke, David Lewis, Kit Fine, Peter van Inwagen, David Armstrong and many others: work that has, to cite just one example, made it anachronistic to dismiss essentialism as anachronistic.⁶ On the traditional grand narrative schemes in the history of philosophy, this activity must be a throwback to pre-Kantian metaphysics. It ought not to be happening; but it is. Many of those who practise it happily acknowledge its continuity with traditional metaphysics; appeals to the authority of Kant, or history, ring hollow, for they are unbacked by any argument that has withstood the test of recent time.

One might try to see in contemporary metaphysics a Quinean breakdown of divisions between philosophy and the natural sciences. But if it is metaphysics naturalized, then so is the metaphysics of Aristotle, Descartes and Leibniz. Broadly *a priori* argument retains a

central role, as do the modal notions of possibility and necessity. Although empirical knowledge constrains the attribution of essential properties, results are more often reached through a subtle interplay of logic and the imagination. The crucial experiments are thought experiments.

Might the contrast between the new-old metaphysics and the representational turn be less stark than it appears to be? These metaphysicians firmly resist attempts to reconstrue their enterprise as the analysis of thought—unlike Sir Peter Strawson, who defined his ‘descriptive metaphysics’ as ‘content to describe the actual structure of our thought about the world’ (1959: 9). But perhaps one cannot reflect on thought or talk about reality without reflecting on reality itself, for the aboutness of thought and talk is intrinsic to it, and its very point. This idea has been emphasized by David Wiggins, Dummett’s successor and my predecessor, author of some of the most distinguished essentialist metaphysics, in which considerations of logic and biology harmoniously combine. Wiggins writes:

Let us forget once and for all the very idea of some knowledge of language or meaning that is not knowledge of the world itself. (Wiggins 2001: 12)

In defining words—for example, natural kind terms—we must point at real specimens.

What there is determines what there is for us to mean. In knowing what we mean, we know something about what there is. That may prompt us to wonder how far the analysis of thought or language could be pursued autonomously, with any kind of methodological priority.

Dummett claimed not that the traditional questions of metaphysics could not be answered but that the way to answer them was by the analysis of thought and language. For example, in order to determine whether there are numbers, one must determine whether

number words such as ‘7’ function semantically like proper names in the context of sentences uttered in mathematical discourse. But what is it so to function? Devil words such as ‘Satan’ appear to function semantically like proper names in the context of sentences uttered in devil-worshipping discourse, but one should not jump to the conclusion that there are devils. However enthusiastically devil-worshippers use ‘Satan’ as though it referred to something, that does not make it refer to something. Although empty names *appear* to function semantically like referring names in the context of sentences uttered by those who believe the names to refer, the appearances are deceptive. ‘Satan’ refers to something if and only if some sentence with ‘Satan’ in subject position (such as ‘Satan is self-identical’) expresses a truth, but the analysis of thought and language may not be the best way to discover whether any such sentence is indeed true.

Historians of philosophy on the grand scale will probably be too Whiggish or Hegelian to regard the linguistic or representational turn as merely a false turning from which philosophy is withdrawing now that it recognizes its mistake. We are supposed to go forward from it, not back. At the very least, we should learn from our mistakes, if only not to repeat them. But if the representational turn was a mistake, it was not a simple blunder; it went too deep for that.

2. It is time to get down to serious work. My aim in what follows is to explore these issues about the nature of philosophy in a microcosm. As a case study, I will consider the problem of vagueness, partly out of habit, but also because it appears to be a paradigm of a philosophical problem about thought and language. For vagueness is generally conceived as a feature of our representations of the world, not of the world itself. Admittedly, some philosophers find tempting the idea of mind-independently vague objects, such as Mount

Everest, vague in their spatiotemporal boundaries and mereological composition, if not in their identity. That kind of vagueness is not my concern here. I will consider an example of a quite standard type, involving a vague predicate.⁷

Suppose that there was once plenty of water on the planet Mars; it was clearly not dry. Ages passed, and very gradually the water evaporated. Now Mars is clearly dry. No moment was clearly the first on which it was dry or the last on which it was not. For a long intermediate period it was neither clearly dry nor clearly not dry. Counting the water molecules would not have enabled us to determine whether it was dry; other measurements would have been equally inconclusive. We have no idea of any investigative procedure that would have resolved the issue. It was a borderline case. No urgent practical purpose compels us to decide whether Mars was dry then, but only a limited proportion of thought and talk in any human society is driven by urgent practical purposes. We should like to know the history of Mars. When necessary, we can always use words other than ‘dry’. Nevertheless, we reflect on the difficulty of classifying Mars as dry or as not dry at those intermediate times, even given exact measurements. We may wonder whether it was either. We ask ourselves:

Was Mars always either dry or not dry?

Henceforth I will refer to that as *the original question*. More precisely, I will use that phrase to designate that interrogative sentence, as used in that context (which is not to deny that the word ‘question’ can also be applied to what interrogative sentences express rather than the sentences themselves).

The original question is at least proto-philosophical in character. It is prompted by an obstacle both hard to identify and hard to avoid that we encounter in applying our conceptual

distinctions. It hints at a serious threat to the validity of our most fundamental forms of logical reasoning. Philosophers disagree about its answer, on philosophical grounds that will be explored below. A philosophical account of vagueness that did not tell us how to answer the original question would thereby be incomplete. Without an agreed definition of ‘philosophy’, we can hardly expect to *prove* that the original question or any other is a philosophical question; but when we discuss its answer, we find ourselves invoking recognisably philosophical considerations. Before we worry about the answer, let us examine the original question itself.

The question queries just the supposition that Mars was always either dry or not dry, which we could formalise as a theorem of classical logic, $\forall t(Dmt \vee \sim Dmt)$.⁸ In words: for every time t , either Mars was dry at t or Mars was not dry at t . The question is composed of expressions that are not distinctively philosophical in character: ‘Mars’, ‘always’, ‘either ... or ...’, ‘not’, ‘was’ and ‘dry’. All of them occur in a recognisably unphilosophical question such as ‘Was Mars always either uninhabited or not dry?’, which someone might ask on judging that Mars is both uninhabited and dry and wondering whether there is a connection. Although philosophical issues can be raised *about* the words in both questions, it does not follow that merely in using those words one is in any way engaging in philosophy. One difference between the two questions is that it is not obviously futile to try to argue on *a priori* grounds that Mars was always either dry or not dry, whereas it is obviously futile to try to argue on *a priori* grounds that Mars was always either uninhabited or not dry.

The original question does not itself *ask* whether it is metaphysically necessary, or knowable *a priori*, or analytic, or logically true that Mars was always either dry or not dry. It simply asks whether Mars always *was* either dry or not dry. Expressions such as ‘metaphysically necessary’, ‘knowable *a priori*’, ‘analytic’ and ‘logically true’ do not occur in

the original question; one can understand it without understanding any such philosophically distinctive terms. This is of course neither to deny nor to assert that it *is* metaphysically necessary, or knowable *a priori*, or analytic, or logically true that Mars was always either dry or not dry. For all that has been said, the proposition may be any combination of those things. But that is not what the original question asks.

In other circumstances, we could have answered the original question on philosophically uninteresting grounds. For instance, if there had never been liquid on Mars, then it would always have been dry, and therefore either dry or not dry. In order to pose a question which could not possibly be answered in that boring way, someone who already grasped one of those philosophically distinctive concepts might ask whether it is metaphysically necessary, or knowable *a priori*, or analytic, or logically true that Mars was always either dry or not dry. The philosophically distinctive expression might then fall under various kinds of suspicion, which would extend to the question in which it occurred. But the original question itself cannot be correctly answered in the boring way with respect to the originally envisaged circumstances. Its philosophical interest, however contingent, is actual.

We could generalize the original question in various ways. We might ask whether *everything* is always either dry or not dry. Then we might notice that discussing that question is quite similar to discussing whether everything is either old or not old, and so on. We might therefore ask whether for every property everything either has that property or lacks it. Such generalizing over properties might itself fall under various kinds of suspicion, which would extend to the question in which it occurred. Someone might even doubt whether there is such a property as dryness. But the original question itself does not attempt such generality. That it has the same kind of philosophical interest as many other questions does not imply that it has no philosophical interest. If that interest is obscured by problematic features of the apparatus

with which we try to generalize it, we can refrain from generalizing it, and stick with the original question.⁹

What is the question about? ‘About’ is not a precise term. On the most straightforward interpretation, a sentence in a context is about whatever its constituents refer to in that context. Thus, taken at face value, the original question is about the planet Mars, the referent of ‘Mars’ in this context; perhaps it is also about dryness, the referent of ‘dry’, and the referents of other constituents too. Since the original question contains no metalinguistic expressions, it is not about the name ‘Mars’ or the adjective ‘dry’. Evidently, the original question is not explicitly about words.

Is the original question implicitly about language? Someone might claim so on the grounds that it is equivalent to questions that are explicitly about language, such as these:

Is the sentence ‘Mars was always either dry or not dry’ true? (Does it express a truth as used in this context?)

Did Mars always belong either to the extension of the word ‘dry’ or to the anti-extension of ‘dry’ (as the word ‘dry’ is used in this context)?

But parallel reasoning would lead to the conclusion that the unphilosophical question ‘Was Mars always either uninhabited or not dry?’ and all everyday questions are also implicitly about language. Since they are not about language in any distinctive sense, the reasoning does not show that the original question is about language in any distinctive sense. Even if the equivalences did show that the original question was in some sense implicitly about language, they could be read in both directions: they would also show that the explicitly metalinguistic

questions were in an equally good sense implicitly not about language.

The equivalences between the questions are in any case only as strong as the corresponding disquotational biconditionals:

(T1) ‘Mars was always either dry or not dry’ is true if and only if Mars was always either dry or not dry.

(T2a) For any time t , Mars belongs to the extension of ‘dry’ at t if and only if Mars is dry at t .

(T2b) For any time t , Mars belongs to the anti-extension of ‘dry’ at t if and only if Mars is not dry at t .

On the face of it, these biconditionals express at best contingent truths. For the word ‘dry’ could have meant *wet*, in which case Mars would have belonged to the extension of ‘dry’ when wet and to the anti-extension of ‘dry’ when dry: for *we* use the word ‘dry’ to mean *dry* even when we are talking *about* circumstances in which it would have meant something else, because we are not talking *in* those circumstances. Thus T2a and T2b do not express necessary truths. Similarly, the sentence ‘Mars was always either dry or not dry’ could have failed to express a truth even though Mars was always either dry or not dry, since ‘always’ could have meant *never*. Thus T1 does not express a necessary truth. We should not assume that a useful notion of aboutness would transfer across merely contingent biconditionals. Perhaps one could interpret T1, T2a and T2b as expressing necessary truths by individuating linguistic expressions in some non-standard manner on which their semantic properties are

essential to them; it is a delicate matter. Moreover, some theorists of vagueness have denied even the actual truth of biconditionals such as T1, T2a and T2b; they might respond to the original question in one way and to the explicitly metalinguistic questions in another. Thus the questions are not pragmatically, dialectically or methodologically equivalent within the context of debates on vagueness. For present purposes, we need not resolve the status of the disquotational biconditionals, because we have already seen that the sense in which they make the original question implicitly about words is too indiscriminate to be useful.

We can argue more directly that the original question is not implicitly about the word ‘dry’ by appeal to a translation test. For consider the translation of the original question into another language, such as Serbian:

Da li je Mars uvek bio suv ili nije bio suv?

The Serbian translation is not implicitly about the English word ‘dry’. But since the questions in the two languages mean the same, what they are implicitly about should also be the same. Therefore, the original question is not implicitly about the word ‘dry’. By similar reasoning, it is not about any word of English or any other language. Of course, given the informality of the notion of implicit aboutness, the argument is not fully rigorous. Nevertheless, the translation test emphasizes how far one would have to water down the notion of reference in order to reach a notion of implicit aboutness on which the original question would be implicitly about a word.

The translation test does not show that the original question is not implicitly about a *concept*, something like the meaning of a word rather than the word itself, for the English word ‘dry’ and its Serbian synonym ‘suv’ both express the concept *dry*. But what basis is

there for the claim that the original question is implicitly about the concept *dry*? We might argue that the original question is in some sense equivalent to a metaconceptual question:

Did Mars always belong either to the extension of the concept *dry* or to the anti-extension of *dry*?

For we might apply the notions of extension and anti-extension to concepts by means of biconditionals similar to T2a and T2b respectively:

(TCa) For any time t , Mars belongs to the extension of *dry* at t if and only if Mars is dry at t .

(TCb) For any time t , Mars belongs to the anti-extension of *dry* at t if and only if Mars is not dry at t .

TC2a and TC2b have a better chance than T2a and T2b of expressing necessary truths, for the contingent relation between words and their meanings has no straightforward analogue for concepts. Concepts are individuated semantically: rather than merely having meanings, they *are* meanings, or something like them.¹⁰ Nevertheless, the argument that the original question is implicitly about the concept *dry* in virtue of being equivalent to the metaconceptual question wildly overgeneralizes, just like the argument that the original question is implicitly about the word ‘dry’ in virtue of being equivalent to the metalinguistic question. For parallel reasoning would lead to the conclusion that the unphilosophical question ‘Was Mars always either uninhabited or not dry?’ is implicitly about the concept *dry*, and likewise for any other unphilosophical question. Since those questions are not about concepts in any distinctive

sense, the original reasoning does not show that the original question is about concepts in any distinctive sense. Even if the equivalences did show that the original question was in some sense implicitly about thought, they can be read in both directions: they would also equally show that the explicitly metaconceptual questions were in an equally good sense implicitly not about thought.

A Fregean might argue: the original question is *explicitly* about the concept *dry*, because it contains the predicate ‘... is dry’ (in the past tense), which refers to the concept *dry*. In that sense, the question ‘Was Mars always either uninhabited or not dry?’ would also be explicitly about the concept *dry*. However, the Fregean is not using the word ‘concept’ with its contemporary meaning, on which concepts are mental or semantic representations, in the realm of sense rather than reference. The Fregean referent of a predicate (a Fregean concept) is simply the function that maps everything to which the predicate applies to the true and everything else to the false: it could be treated as the extension of the predicate, except that it is a function rather than an object. If the predicate refers to the property of dryness or to the set of dry things, then the original question is about the property of dryness or the set of dry things, but that has no tendency to show that it is about thought. Similarly, the Fregean claim has no tendency to show that the question is about thought, for the Fregean concept is in the realm of reference, not in the realm of thought. Like the property and the set, it is no representation but that which is represented.

Representations do not always represent representations. To judge by its overt compositional structure, the original question in particular does not represent any representations. It is not a metalinguistic or metaconceptual question. We have encountered no reason to regard its overt structure as at all misleading in that respect. Our provisional conclusion must therefore be that the original question, although at least proto-philosophical,

is not about thought or language in any distinctive sense. It does not support the representational turn.

3. If the original question, read literally, had too obvious an answer, either positive or negative, that would give us reason to suspect that someone who uttered it had some other meaning in mind, to which the overt compositional structure of the question might be a poor guide. But competent speakers of English may find themselves quite unsure how to answer the question, so we have no such reason for interpreting it non-literally.

The most straightforward reason for answering the original question positively is that ‘Mars was always either dry or not dry’ is a logical truth, a generalization over instances of the law of excluded middle ($P \vee \sim P$, ‘It is either so or not so’) for various times. In my view, that reasoning is sound. However, many think otherwise. They deny the validity of excluded middle for vague terms such as ‘dry’.

The simplest way of opposing the law of excluded middle is to deny outright when Mars is a borderline case that it is either dry or not dry, and therefore to answer the original question in the negative. For instance, someone may hold that Mars was either dry or not dry at time t only if one can know (perhaps later) whether it was dry at t , given optimal conditions for answering the question (and no difference in the history of Mars): since one cannot know, even under such conditions, whether it was dry when the case was borderline, it is not either dry or not dry. One difficulty for this negative response to the original question is that it seems to imply that in a borderline case Mars is neither dry nor not dry: in other words, both not dry and not not dry. That is a contradiction, for ‘not not dry’ is the negation of ‘not dry’.

Intuitionistic logic provides a subtler way to reject the law of excluded middle without denying any one of its instances. Intuitionists ground logic in states of increasing but

incomplete information, rather than a once-for-all dichotomy of truth and falsity; they deny that anything can be both proved and refuted, but they do not assert that everything can be either proved or refuted. For intuitionists, the denial of an instance of excluded middle ($\sim(P \vee \sim P)$, ‘It is not either so or not so’) entails a contradiction ($\sim P \ \& \ \sim\sim P$, ‘It is both not so and not not so’), just as it does in classical logic, and contradictions are as bad for them as for anyone else. Thus they cannot assert that Mars was once not either dry or not dry ($\exists t \sim(Dmt \vee \sim Dmt)$), for that would imply that a contradiction once obtained ($\exists t(\sim Dmt \ \& \ \sim\sim Dmt)$, ‘Mars was once both not dry and not not dry’), which is intuitionistically inconsistent. However, although intuitionists insist that proving an existential claim in principle involves proving at least one instance of it, they allow that disproving a universal claim need not in principle involve disproving at least one instance of it. The claim that something lacks a property is intuitionistically stronger than the claim that not everything has that property. Thus one might assert that Mars was not always either dry or not dry ($\sim\forall t(Dmt \vee \sim Dmt)$), on the general grounds that there is no adequate procedure for sorting all the times into the two categories, without thereby committing oneself to the inconsistent existential assertion that it was once not either dry or not dry. Hilary Putnam once proposed the application of intuitionistic logic to the problem of vagueness for closely related reasons.¹¹ Thus one might use intuitionistic logic to answer the original question in the negative.

On closer inspection, this strategy looks less promising. For a paradigm borderline case is the worst case for the law of excluded middle with respect to a term such as ‘dry’ for which threats to the law other than from vagueness are not relevant, in the sense that both proponents and opponents of the law can agree that it holds in a paradigm borderline case only if it holds universally. In symbols, if τ is a time when Mars was a paradigm borderline case, then we have $(Dm\tau \vee \sim Dm\tau) \rightarrow \forall t(Dmt \vee \sim Dmt)$ (‘If Mars was either dry or not dry at

time τ , then Mars was always either dry or not dry'). But on this approach the law does not hold always hold in this case ($\sim\forall t(Dmt \vee \sim Dmt)$, 'Mars was not always either dry or not dry'), from which intuitionistic logic allows us to deduce that it does not hold in the paradigm borderline case ($\sim(Dm\tau \vee \sim Dm\tau)$, 'Mars was not either dry or not dry at τ '), which is a denial of a particular instance of the law, and therefore intuitionistically inconsistent (it entails $\sim Dm\tau$ & $\sim\sim Dm\tau$, 'Mars was both not dry and not not dry at τ '). Thus the intuitionistic approach forces one to deny that there are such paradigm borderline cases for vague terms. That denial is hard to reconcile with experience: after all, the notion of a borderline case is usually explained by examples.

The problems for the intuitionistic approach do not end there. One can show that the denial of the conjunction of any finite number of instances of the law of excluded middle is intuitionistically inconsistent.¹² A universal generalization of the law over a finite domain is therefore intuitionistically false too. If time is infinitely divisible, the formula $\forall t(Dmt \vee \sim Dmt)$ generalizes the law over an infinite domain of moments of time, but that is not crucial to the phenomena of vagueness. We could just as well have asked the original question about a long finite series of moments at one-second intervals; it would have been equally problematic. The classical sorites paradox depends on just such a finite series: a heap of sand consists of only finitely many grains, but when they are carefully removed one by one, we have no idea how to answer the question 'When did there cease to be a heap?'. It is intuitionistically inconsistent to deny that Mars was dry or not dry at each of those moments. Thus intuitionistic logic provides a poor basis for a negative answer to the original question.

Other theorists of vagueness refuse to answer the original question either positively or negatively. They refuse to assert that Mars was always either dry or not dry; they also refuse to assert that it was not always either dry or not dry.

A simple version of this approach classifies vague sentences (relative to contexts) as definitely true (T), definitely false (F) or indefinite (I); borderline sentences are classified as indefinite. The generalized truth-tables of a three-valued logic are used to calculate which of these values to assign to a complex sentence in terms of the values assigned to its constituent sentences. The negation of P , $\sim P$, is definitely true if P is definitely false, definitely false if P is definitely true and indefinite if P is indefinite:

P	$\sim P$
T	F
I	I
F	T

A conjunction $P \& Q$ (' P and Q ') is definitely true if every conjunct is definitely true; it is definitely false if some conjunct is definitely false; otherwise it is indefinite. A disjunction $P \vee Q$ ('Either P or Q ') is definitely true if some disjunct is definitely true; it is definitely false if every disjunct is definitely false; otherwise it is indefinite:

P	Q	$P \& Q$	$P \vee Q$
T	T	T	T
T	I	I	T
T	F	F	T
I	T	I	T
I	I	I	I
I	F	F	I

F	T	F	T
F	I	F	I
F	F	F	F

A universal generalization is treated like the conjunction of its instances: it is definitely true if every instance is definitely true; it is definitely false if some instance is definitely false; otherwise it is indefinite. An existential generalization is treated like the disjunction of its instances: it is definitely true if some instance is definitely true; it is definitely false if every instance is definitely false; otherwise it is indefinite. The three-valued tables generalize the familiar two-valued ones in the sense that one recovers the latter by reading ‘T’ as ‘true’, ‘F’ as ‘false’, and ignoring lines with ‘I’.

Let us apply this three-valued approach to the original question. If Mars is definitely dry or definitely not dry at t , then Dmt is definitely true or definitely false, so the instance of excluded middle $Dmt \vee \sim Dmt$ is definitely true. But if Mars is neither definitely dry nor definitely not dry at t , then Dmt is indefinite, so $\sim Dmt$ is indefinite too by the table for negation, so $Dmt \vee \sim Dmt$ is classified as indefinite by the table for disjunction. Since Mars was once a borderline case, the universal generalization $\forall t(Dmt \vee \sim Dmt)$ has a mixture of definitely true and indefinite instances; hence it is classified as indefinite. Therefore its negation $\sim \forall t(Dmt \vee \sim Dmt)$ is also indefinite. Thus three-valued theoreticians who wish to assert only definite truths neither assert $\forall t(Dmt \vee \sim Dmt)$ nor assert $\sim \forall t(Dmt \vee \sim Dmt)$. They answer the original question neither positively nor negatively.

Three-valued logic replaces the classical dichotomy of truth and falsity by a three-way classification. Fuzzy logic goes further, replacing it by a continuum of degrees of truth between definite truth and definite falsity. According to proponents of fuzzy logic, vagueness

should be understood in terms of this continuum of degrees of truth. For example, ‘It is dark’ may increase continuously in degree of truth as it gradually becomes dark. On the simplest version of the approach, degrees of truth are identified with real numbers in the interval from 0 to 1, with 1 as definite truth and 0 as definite falsity. The semantics of fuzzy logic provides rules for calculating the degree of truth of a complex sentence in terms of the degrees of truth of its constituent sentences. For example, the degrees of truth of a sentence and of its negation sum to exactly 1; the degree of truth of a disjunction is the maximum of the degrees of truth of its disjuncts; the degree of truth of a conjunction is the minimum of the degrees of truth of its conjuncts. For fuzzy logic, although the three-valued tables above are too coarse-grained to give complete information, they still give correct results if one classifies every sentence with an intermediate degree of truth, less than the maximum and more than the minimum, as indefinite.¹³ Thus the same reasoning as before shows that fuzzy logicians should answer the original question neither positively nor negatively.

Although three-valued and fuzzy logicians reject both the answer ‘Yes’ and the answer ‘No’ to the original question, they do not reject the question itself. What they reject is the restriction of possible answers to ‘Yes’ and ‘No’. They require a third answer, ‘Indefinite’, when the queried sentence takes the value I. More formally, consider the three-valued table for the sentence operator Δ , read as ‘definitely’ or ‘it is definite that’:

P	ΔP
T	T
I	F
F	F

Even for fuzzy logicians this table constitutes a complete semantics for Δ , since the only output values are T and F, which determine unique degrees of truth (1 and 0). A formula of the form $\sim\Delta P \ \& \ \sim\Delta\sim P$ ('It is neither definitely so nor definitely not so') characterizes a borderline case, for it is definitely true if P is indefinite and definitely false otherwise. In response to the question $P?$, answering 'Yes' is tantamount to asserting P , answering 'No' is tantamount to asserting $\sim P$, and answering 'Indefinite' is tantamount to asserting $\sim\Delta P \ \& \ \sim\Delta\sim P$. On the three-valued and fuzzy tables, exactly one of these three answers is correct (definitely true) in any given case; in particular, the correct answer to the original question is 'Indefinite'.

On the three-valued and fuzzy approaches, to answer 'Indefinite' to the question 'Is Mars dry?' is to say something about Mars, just as it is if one answers 'Yes' or 'No'. It is not a metalinguistic response. For Δ is no more a metalinguistic operator than \sim is. They have the same kind of semantics, given by a many-valued truth-table. Just as the negation $\sim P$ is about whatever P is about, so are ΔP and $\sim\Delta P \ \& \ \sim\Delta\sim P$. Thus the answer 'Indefinite' to the original question involves no semantic ascent to a metalinguistic, metaconceptual or metarepresentational level. It remains at the level of discourse about Mars.

The three-valued and fuzzy approaches have many suspect features. For instance, they treat any sentence of the form ΔP as perfectly precise, because it always counts as definitely true or definitely false, never as indefinite, whatever the status of P ; thus $\Delta\Delta P \vee \Delta\sim\Delta P$ ('It is definite whether it is definitely so') is always definitely true. This result does not fit the intended interpretation of Δ . For 'Mars is definitely wet' is not perfectly precise. Just as no moment is clearly the last on which Mars was wet or the first on which it was not, so no moment is clearly the last on which it was definitely wet or the first on which it was not definitely wet. Just as it was sometimes unclear whether Mars is wet, so it was sometimes

unclear whether it was definitely wet. This is one form of the notorious problem of higher-order vagueness: in other words, there are borderline cases of borderline cases, and borderline cases of borderline cases of borderline cases, and so on. The problem has never received an adequate treatment within the framework of three-valued or fuzzy logic, and it is far from obvious that it could.¹⁴

Some philosophers, often under the influence of the later Wittgenstein, deny the relevance of formal semantic theories to vague natural languages. They regard the attempt to give a systematic statement of the truth-conditions of English sentences in terms of the meanings of their constituents as vain. For them, the formalization of ‘Mars was always either dry or not dry’ as $\forall t(Dmt \vee \sim Dmt)$ is already suspect. This attitude suggests a premature and slightly facile pessimism. No doubt formal semantics has not described any natural language with perfect accuracy; what has not been shown is that it provides no deep insights into natural languages. In particular, it has not been shown that the main semantic effects of vagueness are not susceptible to systematic formal analysis. In any case, for present purposes the claim that there can be no systematic theory of vagueness is just one more theory of vagueness, although —unless it is self-refuting —not a systematic one; it does not even answer the original question. Even if that theory were true, the other theories of vagueness, however false, would still exist, and would still have been accepted by some intelligent and linguistically competent speakers.

This is no place to resolve the debate between opposing theories of vagueness. The present point is just that **different theories support contrary answers to the original question.** All these theories have their believers. Any answer to the original question, positive, negative or indefinite, is contentious. Of course, if everyone found their own answer obvious, but different people found different answers obvious, then we might suspect that they were

interpreting the question in different ways, talking past each other. But that is not so: almost everyone who reflects on the original question finds it difficult and puzzling. Even when one has settled on an answer, one can see how intelligent and reasonable people could answer differently while understanding the meaning of the question in the same way. If it has an *obvious* answer, it is the answer ‘Yes’ dictated by classical logic, but those of us who accept that answer can usually imagine or remember the frame of mind in which one is led to doubt it. Thus the original question, read literally, has no obvious answer in any sense that would give us reason to suspect that someone who uttered it had some other reading in mind.

Without recourse to non-literal readings, some theorists postulate ambiguity in the original question. For example, some three-valued logicians claim that ‘not’ in English is ambiguous between the operators \sim (strong negation) and $\sim\Delta$ (weak negation): although $\sim P$ and $\sim\Delta P$ have the same value if P is definitely true or definitely false, $\sim\Delta P$ is true while $\sim P$ is indefinite if P is indefinite. While $P \vee \sim P$ (‘It is so or not so’) can be indefinite, $P \vee \sim\Delta P$ (‘It is so or not definitely so’) is always definitely true. On this view, the original question queries $\forall t(Dmt \vee \sim Dmt)$ on one reading, $\forall t(Dmt \vee \sim\Delta Dmt)$ on another; the latter is definitely true (Mars was always either dry or not definitely dry) while the former is indefinite. Thus the correct answer to the original question depends on the reading of ‘not’. It is ‘Indefinite’ if ‘not’ is read as strong negation, ‘Yes’ if ‘not’ is read as weak negation. The three-valued logician’s reasoning here is undermined by higher-order vagueness, but that is not the present issue.

If ‘not’ were ambiguous in the way indicated, it would still not follow that the dispute over the original question is merely verbal. For even when we agree to consider it under the reading of ‘not’ as strong negation, which does not factorize in the manner of $\sim\Delta$, we yet find theories of vagueness in dispute over the correct answer. We have merely explained our terms

in order to formulate more clearly a difficult question about Mars.

Still, it might be suggested, the dispute between different theories of vagueness is verbal in the sense that their rival semantics characterize different possible languages or conceptual schemes: our choice of which of them to speak or think would be pragmatic, based on considerations of usefulness rather than of truth.

To make sense of the pragmatic view, one must suppose that vague atomic sentences are classifiable according to both the bivalent scheme (as true or false) and the trivalent scheme (as definitely true, indefinite or definitely false), and that the truth-tables of each scheme define intelligible connectives, although the connective defined by a trivalent table should be distinguished from the similar-looking connective defined by the corresponding bivalent table. Definite truth implies truth, and definite falsity implies falsity, but indefiniteness does not discriminate between truth and falsity: although all borderline atomic sentences are indefinite, some are true and others false. As Mars dries, ‘Mars is dry’ is first false and definitely false, then false but indefinite, then true but indefinite, and finally true and definitely true. However, this attempted reconciliation of the contrasting theories does justice to neither side. For trivalent logicians, once we know that a sentence is indefinite, there is no further question of its truth or falsity to which we do not know the answer: the category of the indefinite was introduced in order not to postulate such a mystery. Similarly, for fuzzy logicians, once we know the intermediate degree of truth of a sentence, there is no further question of its truth or falsity to which we do not know the answer: intermediate degrees of truth were introduced in order not to postulate such a mystery. In formal terms, trivalent and fuzzy logics are undoubtedly less convenient than bivalent logic: the justification for introducing them was supposed to be the inapplicability of the bivalent scheme to vague sentences. If a bivalent vague language is even a possible option, then the

trivalent and fuzzy accounts of vagueness are mistaken. Conversely, from a bivalent perspective, the trivalent and fuzzy semantics do not fix possible meanings for the connectives, because they do not determine truth-conditions for the resultant complex sentences: for example, the trivalent table for \sim does not specify when $\sim P$ is true in the bivalent sense. It would therefore be a fundamental misunderstanding of the issue at stake between theories of vagueness to conceive it as one of a pragmatic choice of language.

We already speak the language of the original question; we understand those words and how they are put together; we possess the concepts that they express; we grasp what is being asked. That semantic knowledge may be necessary if we are to know the answer to the original question.¹⁵ It is not sufficient, for it does not by itself put one in a position to arbitrate between conflicting theories of vagueness. For each of those theories has been endorsed by some competent speakers of English who fully grasp the question.

Competent speakers may of course fail to reflect adequately on their competence. Although the proponents of conflicting theories of vagueness presumably have reflected on their competence, their reflections may have contained mistakes. Perhaps reflection of sufficient length and depth on one's competence would lead one to the correct answer to the original question. But the capacity for such more or less philosophical reflection is not a precondition of semantic competence. Philosophers should resist the professional temptation to require all speakers to share their proclivities.

We can distinguish two levels of reflection, the logical and the metalogical. In response to the original question, logical reflection would involve reasoning with terms of the kind in which the question is phrased; the aim is to reach a conclusion that answers the question. For example, one might conclude by classical logic that Mars was always either dry or not dry; one might conclude by fuzzy logic that it is indefinite whether it was always one

or the other. The logical level is not purely mechanical. When the reasoning is complex, one needs skill to select from the many permissible applications of the rules one sequence that leads to an answer to the question. When the reasoning is informal, one needs good judgement to select only moves that really are permissible applications of the rules. But one is still thinking about whatever the question was about. It is at the metalogical level of reflection that one starts to think about the semantics of the logical connectives and other expressions that one employed at the logical level. For example, at the metalogical level one may assert or deny that the sentence 'Mars was always either dry or not dry' is a logical truth. Only at the metalogical level can one be expected to articulate the rules that one was using at the previous level.

It must be possible to think logically without thinking metalogically, for otherwise by the same principle thinking metalogically would involve thinking metametalogically, and so *ad infinitum*: our thinking never goes all the way up such an infinite hierarchy. What can prompt ascent to the metalogical level are hard cases in which one feels unclear about the permissibility of a given move at the logical level. One's mastery of the language and possession of concepts leave one quite uncertain how to go on. In the case of the original question, a salient line of classical reasoning leads to a positive answer: it persuades some competent speakers while leaving others unconvinced. Even to discuss the contentious reasoning we must semantically ascend. We cannot hope to resolve the dispute undogmatically if we never leave the lower level.

4. The argument so far has reached two conclusions that at first sight are hard to reconcile with each other. First, the original question is not about thought or language. Second, to answer it adequately one must assess rival theories of vagueness in thought and language.

How can that way of reaching an answer be appropriate to the original question? We might therefore find ourselves tempted back to the idea that somehow the original question was surreptitiously about thought or language.

On further reflection, the apparent conflict between the two conclusions is less surprising. Many non-philosophical questions that are not about thought or language cannot be resolved without inquiry into thought or language.

Suppose that a court of law must decide whether Smith killed Jones. The question is not who said or thought what. Nevertheless, the crucial arguments may be over whether to trust the witnesses' testimony. How is what they say now related to what they think now or thought then? How is what they think now or thought then related to what actually happened? Are they lying or sincere? Are their memories confused or clear? Those are questions about their thought and speech, their representations. They hold the key to whether Smith killed Jones, even though that question is not about representations. Of course, the questions about representations are not about them in isolation from what they represent: they are relevant because they concern the relation between the representations and what they represent.

The court must decide the issue on the evidence before it. In a criminal case, does the evidence put it beyond reasonable doubt that Smith killed Jones? In a civil case, does the evidence make it more probable than not? If the court is really deciding a question about testimonial evidence, is that not already a question about representations?¹⁶ But the question about the evidence arises in virtue of its bearing on the primary question, whether Smith killed Jones. Indeed, the question about the evidence is exactly a question about its bearing on the primary question. So the point stands.

Historians are often in a similar position. They want to know what happened. The way to achieve that is largely by considering documents, representations of what happened —not

in isolation, but in relation to what they represent. Most obviously, the historian wants to know whether the documents accurately represent what happened, but to answer that question one must in turn ask about their provenance: who produced them, when and why? Thus the history of the events of primary interest requires a history of representations of those events. Those histories typically overlap, for the production of a representation of some part of a complex human event is often another part of the same complex event.

Something analogous occurs in the methodology of the natural sciences. We wish to know the value of some physical quantity. We must devise apparatus to measure it. We may find ourselves in disputes over the functioning of different devices. Although the primary question was not about those measuring devices, we cannot answer it adequately without considering them. We need a theory about the relation between the value of the quantity and the representations of it that we record when we use our instruments. The scientific investigation of the physical quantity widens to include the scientific investigation of its interaction with our experimental equipment. After all, our apparatus is part of the same natural world as the primary topic of our enquiry.

These analogies make it less surprising that when we try to answer the original question, which is not a question about thought or language, our main task is to adjudicate between rival theories of vague thought and language. A theory of vagueness validates some deduction that concludes with an answer to the original question. That deduction uses but does not mention vague thought or language. It is formulated at the logical level, like the original question itself, not at the metalogical level. But to justify trusting that deduction, rather than one that reaches another conclusion by other rules, one must assess the rival theories of vagueness.

That theories of vagueness conflict in their answers to the original question shows that

they are not confined to claims about thought and talk. Theories such as epistemicism and supervaluationism which employ classical logic have ‘Mars was always either dry or not dry’ as a theorem, once they are formulated in a suitably expressive language. To reiterate, that theorem is not about thought or talk.

For three-valued and fuzzy approaches, the matter is only slightly more complicated. Their proponents assert:

(C) It is indefinite whether Mars was always either dry or not dry.

On those approaches, C does not count as about thought or language. Strictly speaking, however, C does not follow from the three-valued or fuzzy theory of vagueness itself; for all the theory implies, there was never any liquid on Mars, in which case it would always have been either dry or not dry, even by three-valued or fuzzy standards, and it would not have been indefinite. The theory implies only a conditional theorem:

(P1) If it was once indefinite whether Mars was dry then it is indefinite whether Mars was always either dry or not dry.

Three-valued or fuzzy theorists can combine P1 with what they regard as an empirical truth about Mars:

(P2) It was once indefinite whether Mars was dry.

From P1 and P2 they use the rule of modus ponens (from ‘If P then Q ’ and ‘ P ’ infer ‘ Q ’) to

infer C, the answer to the original question. Although their theorem P1 does not answer the question by itself, it is no more about thought or language than C is. Their theories are just as committed as classical ones to making claims that are not about thought or language.

In order to give an adequately reflective answer to a simple, non-technical, non-metarepresentational question, we were forced to adjudicate between complex, technical, metarepresentational theories. This phenomenon seems to have been overlooked by those who complain about the ‘arid’ technical minuteness of much philosophy in the analytic tradition. A question may be easy to ask but hard to answer. Even if it is posed in dramatic and accessible terms, that does not entail that the reflections needed to select between rival answers are equally dramatic and accessible. Such contrasts are commonplace in other disciplines; it would have been amazing if they had not occurred in philosophy. Impatience with the long haul of technical reflection is a form of shallowness, often thinly disguised by histrionic advocacy of depth. Serious philosophy is always likely to bore those with short attention-spans.¹⁷

The rise of modern logic from Frege onwards has provided philosophers with conceptual instruments of unprecedented power and precision, enabling them to formulate hypotheses with more clarity and to determine their consequences with more reliability than ever before. Russell’s theory of descriptions showed vividly how differences between the surface form of a sentence and its underlying logical form might mislead us as to its logical relations and thereby create philosophical illusions. The development of formal model-theoretic truth-conditional semantics by Tarski and others has provided a rigorous framework for thinking about the validity of our inferences. These theoretical advances have enormous intellectual interest in their own right. They may have made it tempting to suppose that all philosophical problems are problems of language: but they do not really provide serious

evidence for that conjecture.

Many have been attracted to the idea that all philosophical problems are linguistic or conceptual for deeper philosophical reasons, which I can barely touch on here. For instance, if the method of philosophy is *a priori* reflection, how can it lead to substantive knowledge of the world? Those who find that question compelling may propose that it informs us of relations of ideas rather than matters of fact, or that its truths are analytic rather than synthetic, or that it presents rules of grammar disguised as descriptions, or that its aim is the analysis of thought or language. We have noted how poorly these claims fit the case of vagueness. We may suspect the presence of empiricist presuppositions in the background — or, as with Ayer, in the foreground. Not starting with such presuppositions, **we should be open to the idea that thinking just as much as perceiving is a way of learning how things are.**

Although we do not fully understand *how* thinking can provide new knowledge, the cases of logic and mathematics constitute overwhelming evidence that it does so. The case of the original question, which is philosophical yet queries a theorem of classical logic, shows that we cannot segregate logic from philosophy and claim that *a priori* thinking illuminates the former but not the latter. In particular, conceptions of logic and mathematics as (unlike philosophy) as somehow trivial or purely formal have not been vindicated by any clear account of ‘trivial’ or ‘purely formal’. Whether a given formal system of logic or mathematics is consistent is itself a non-trivial question of logic or mathematics; we know from Gödel’s second incompleteness theorem that the consistency of a system of elementary mathematics can only be decided in less elementary mathematics, unless the original system was inconsistent.

To deny that all philosophical questions are about thought or language is not to deny the obvious, that many are. We have also seen how in practice the attempt to answer a

question which is not about thought or language can largely consist in thinking about thought and language. Some contemporary metaphysicians appear to believe that they can safely ignore formal semantics and the philosophy of language because their interest is in a largely extra-mental reality. They resemble an astronomer who thinks that he can safely ignore the physics of telescopes because his interest is in the extra-terrestrial universe. In delicate matters, his attitude makes him all the more likely to project features of his telescope confusedly onto the stars beyond. Similarly, the metaphysicians who most disdain language are the most likely to be its victims. Again, those who neglect logic in order to deduce their philosophy from natural science make frequent logical errors in their deductions; their philosophical conclusions are not logical consequences of their scientific premises.

Analytic philosophy at its best uses logical rigour and semantic sophistication to achieve a sharpness of philosophical vision unobtainable by other means. To sacrifice those gains would be to choose blurred vision. Fortunately, good vision is not restricted to looking at eyes.¹⁸

Notes

- 1 Ayer's three immediate predecessors were John Cook Wilson, H.H. Joachim and H.H. Price.
- 2 That language and thought are representational in this sense does not imply that sentences or thoughts represent entities such as states of affairs; the representing might be done by constituent words or concepts. When Davidson denies that language is representational, he is denying that sentences as contrasted with singular terms represent objects of some kind (1990: 281, 304).
- 3 Although McDowell is sometimes classified as a 'post-analytic' philosopher, he finds his own way to accept Dummett's 'fundamental tenet of analytical philosophy', that 'philosophical questions about thought are to be approached through language' (1994: 125).
- 4 McDowell's invocation of humility (1994: 40) addresses contingent limitations rather than necessary ones.
- 5 Mark Johnston (1993: 96-97) discusses 'the Enigmas, entities essentially undetectable by us'. He stipulates that they are collectively as well as individually undetectable; thus our elusive objects need not be his Enigmas. If we cannot have good evidence that there are no Enigmas, it may well be a waste of time to worry whether there are

Enigmas. But it would not follow that it is a waste of time to worry whether there *can be* Enigmas. Their definition does not rule out knowledge of the possibility of such things, and knowledge of that possibility may itself be philosophically useful (indeed, Johnston uses it for his philosophical purposes).

- 6 See, for example, Kripke 1980, French, Uehling and Wettstein 1986, Fine 1994 and 1995 and Wiggins 2001.
- 7 On vagueness in general see Keefe and Smith 1997 and Williamson 1994. On vague objects see Williamson 200? and references therein.
- 8 Classical logic is the standard logic of expressions such as ‘every’, ‘either ... or ...’ and ‘not’ on the assumption that there is a mutually exclusive, jointly exhaustive dichotomy of sentences into the true and the false.
- 9 See also Quine 1970: 11.
- 10 Even if a word retains its linguistic meaning, its reference may shift with the context of utterance (‘I’, ‘now’, ‘here’). If ‘dry’ undergoes such contextual shifts, T2a and T2b may fail when interpreted as generalizations about utterances of ‘dry’ in contexts other than the theorist’s own. It might be argued that concepts can also undergo contextual shifts in reference: you use the concept *I* to refer (in thought) to yourself but I use the same concept to refer to myself; at noon we use the concept *now* to think of noon but at midnight we use the same concept to refer to midnight; at the North

Pole we use the concept *here* to refer to the North Pole but at the South Pole we use the same concept to refer to the South Pole. If so, TCa and TCb may also fail when interpreted as generalizations about uses of the concept *dry* in contexts other than the theorist's own.

- 11 For intuitionist logic in general see Dummett 1977. For its application to the problem of vagueness see Putnam 1983 and, for discussion, Williamson 1996, Chambers 1998, and references therein.
- 12 One proves by mathematical induction on n that if A_n is the conjunction of n instances of excluded middle then $\sim A_n$ is intuitionistically inconsistent.
- 13 This point does not generalize to the semantics of conditionals in fuzzy logic, given the standard rule that if the consequent is lower than the antecedent in degree of truth then the degree of truth of the conditional falls short of 1 by the amount by which the consequent falls short of the antecedent in degree of truth; otherwise the degree of truth of the conditional is 1. Thus if P has a higher degree of truth than Q but both are indefinite then $P \rightarrow Q$ is indefinite while $Q \rightarrow P$ is definitely true. Thus the information that the antecedent and consequent are indefinite does not determine whether the conditional is indefinite.
- 14 See Williamson 1999b for an analysis of higher-order vagueness.
- 15 Of course, monolingual speakers of another language might know whether Mars was

always dry or not dry without ever hearing of the original question, which is an interrogative sentence of English; they use a synonymous sentence of their own language. They do not know whether the original English question has a positive answer. Someone might even have that knowledge without understanding the original question, because the knowledge can be passed along a chain of testimony; understanding of the original question is needed only at one end of the chain. These quibbles do not affect the argument.

- 16 Non-testimonial evidence may be taken to include non-representational items such as a bloodied knife; this is what lawyers call ‘real evidence’. For an argument that all evidence in an epistemologically central sense of the term is propositional see Williamson 2000: 194-200. For example, the evidence in this sense might include the proposition that the bloodied knife was found at the scene of the crime, but not the knife itself.
- 17 Popularization has its place, in philosophy as in physics, but should not be confused with the primary activity.
- 18 Thanks to Brian Leiter and discussion groups at Oxford University and the University of Texas at Austin for comments on earlier versions of this paper.

Bibliography

- Ayer, A. J. 1936. *Language, Truth and Logic*. London: Victor Gollancz.
- Chambers, T. 1998. 'On vagueness, *sorites*, and Putnam's "intuitionistic strategy"'. *Monist* 81: 343-348.
- Davidson, D. 1990. 'The structure and content of truth'. *Journal of Philosophy*: 278-328.
- Dummett, M. A. E. 1977. *Elements of Intuitionism*. Oxford: Oxford University Press.
- Dummett, M. A. E. 1978. *Truth and Other Enigmas*. London: Duckworth.
- Dummett, M. A. E. 1993. *Origins of Analytical Philosophy*. London: Duckworth.
- Fodor, J. 1975. *The Language of Thought*. New York: Thomas Y. Crowell.
- Fine, K. 1994. 'Essence and modality', in J. Tomberlin, ed., *Philosophical Perspectives*, 8, *Logic and Language*. Atascadero, CA: Ridgeview.
- Fine, K. 1995. 'Senses of essence', in W. Sinnott-Armstrong, D. Raffman and N. Asher, eds., *Modality, Morality, and Belief: Essays in Honor of Ruth Barcan Marcus*. Cambridge: Cambridge University Press.
- French, P.A., Uehling, T.E., and Wettstein, H.K., eds. 1986. *Midwest Studies in Philosophy XI: Studies in Essentialism*. Minneapolis: University of Minnesota Press.
- Johnston, M. 1993. 'Objectivity refigured: pragmatism without verificationism', in J. Haldane and C. Wright, eds., *Reality, Representation and Projection*. Oxford: Oxford University Press.
- Keefe, R., and Smith, P. (eds.) 1997. *Vagueness: A Reader*. Cambridge, Mass.: MIT Press.
- Kripke, S.A. 1980. *Naming and Necessity*. Oxford: Blackwell.
- McDowell, J. H. 1994. *Mind and World*. Cambridge, Mass.: Harvard University Press.

- Putnam, H. 1983. 'Vagueness and alternative logic'. *Erkenntnis* 19: 297-314.
- Quine, W. V. 1970. *Philosophy of Logic*. Englewood Cliffs, N. J.: Prentice-Hall.
- Rorty, R. (ed.) 1967. *The Linguistic Turn: Recent Essays in Philosophical Method*. Chicago: University of Chicago Press.
- Strawson, P. F. 1959. *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.
- Wiggins, D. R. P. 2001. *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.
- Williamson, T. 1994. *Vagueness*. London: Routledge.
- Williamson, T. 1996. 'Putnam on the sorites paradox'. *Philosophical Papers* 25: 47-56.
- Williamson, T. 1999a. 'Truthmakers and the converse Barcan formula'. *Dialectica* 53: 253-270.
- Williamson, T. 1999b. 'On the structure of higher-order vagueness'. *Mind* 108: 127-143.
- Williamson, T. 200?. 'Vagueness in reality', in M. Loux and D. Zimmerman, eds., *The Oxford Handbook of Metaphysics*. Oxford: Oxford University Press.