

THE SIGNIFICANCE OF PHILOSOPHICAL SCEPTICISM

BARRY STROUD

CLARENDON PRESS · OXFORD

experiencing and thinking persons I can know anything about if Descartes is correct. Other people, as I understand them, are not simply sensory experiences of mine; they too, if they exist, will therefore inhabit the unreachable world beyond my sensory experiences, along with the tables and chairs and other things about which I can know nothing. So at least with respect to what I can know I could not console myself with thoughts of a like-minded community of perceivers all working together and cheerfully making do with what a communal veil of perception provides. I would have no more reason to believe that there are any other people than I have to believe that I am now sitting in a chair writing. The representations or sensory experiences to which Descartes's conclusion would restrict my knowledge could be no other than my own sensory experiences; there could be no communal knowledge even of the veil of perception itself. If my own sensory experiences do not make it possible for me to know things about the world around me they do not make it possible for me to know even whether there are any other sensory experiences or any other perceiving beings at all.

The consequences of accepting Descartes's conclusion as it is meant to be understood are truly disastrous. There is no easy way of accommodating oneself to its profound negative implications. But perhaps by now we have come far enough to feel that the whole idea is simply absurd, that ultimately it is not even intelligible, and that there can be no question of 'accepting' Descartes's conclusion at all. I have no wish to discourage such a reaction. I would only insist that the alleged absurdity or unintelligibility must be identified and made out. I think that is the only way we can hope to learn whatever there is to be learned from Descartes's investigation. In the next chapter I consider a powerful form of criticism along these lines and try to sketch a certain conception of the relation between the philosophical investigation of knowledge and our everyday standards and procedures for assessing knowledge. If that conception can be explained and defended, the sceptical conclusion will remain intact and its scope and negative significance will be undiminished.

II

Philosophical Scepticism and Everyday Life

I think that when we first encounter the sceptical reasoning outlined in the previous chapter we find it immediately gripping. It appeals to something deep in our nature and seems to raise a real problem about the human condition. It is natural to feel that either we must accept the literal truth of the conclusion that we can know nothing about the world around us, or else we must somehow show that it is not true. Accepting it and holding to it consistently seem disastrous, and yet rejecting it seems impossible. But what is the 'literal truth' of that conclusion? Both responses depend on a firm understanding of what it says and means; without that there would be nothing determinate to accept as true or to reject as false. That proper understanding of the sceptical conclusion is what I want to concentrate on. That is why I suggest we look to the source of that conclusion—how it is arrived at and how it becomes so unavoidable—and in particular at just how closely Descartes's requirement that the dream-possibility must always be eliminated corresponds to our ordinary standards or requirements for knowledge in everyday life.

In suggesting that we try to determine exactly what the sceptical reasoning manages to establish I do not mean to deny that it does raise deep problems about the human condition and can reveal something of great significance about human knowledge. It might seem as if that is not so, since it might seem that as soon as we even glance in the direction of the standards and procedures we follow in everyday life we will find that there is nothing at all in Descartes's argument. It is obvious that we do not always insist that people know they are not dreaming before we allow that they know something in everyday life, or even in science or a court of law, where the standards are presumably

stricter. So it can easily look as if Descartes reaches his sceptical conclusion only by violating our ordinary standards and requirements for knowledge, perhaps substituting a new and different set of his own. If that were so his conclusion would not have the consequences it seems to have for our everyday and scientific knowledge and beliefs. So understood, it would not have the significance we originally take it to have.

One example of a diagnosis of scepticism along these lines goes as follows. Suppose someone makes the quite startling announcement that there are no physicians in the city of New York. That certainly seems to go against something we all thought we knew to be true. It would really be astonishing if there were no physicians at all in a city that size. When we ask how the remarkable discovery was made, and how long this deplorable state of affairs has obtained, suppose we find that the bearer of the startling news says it is true because, as he explains, what he means by 'physician' is a person who has a medical degree and can cure any conceivable illness in less than two minutes.¹ We are no longer surprised by his announcement, nor do we find that it contradicts anything we all thought we knew to be true. We find it quite believable that there is no one in the whole city who fulfils all the conditions of that peculiar 're-definition' of 'physician'. Once we understand it as it was meant to be understood, there is nothing startling about the announcement except perhaps the form in which it was expressed. It does not deny what on first sight it might seem to deny, and it poses no threat to our original belief that there are thousands and thousands of physicians in New York.

The suggestion is that the sceptical conclusion is in the same boat. It too is said to rest on a misunderstanding or distortion of the meanings of the words in which it is expressed. It is at first astonishing to be told that no one can ever know anything about the world around us, but once we learn that the 'knowledge' in question is 'knowledge' that requires the fulfilment of a condition which is not in fact required for the everyday or scientific knowledge we are

¹ See P. Edwards, 'Bertrand Russell's Doubts About Induction' in A. Flew (ed.), *Logic and Language*, First Series (Oxford, 1955).

interested in, we will no longer be surprised or disturbed by that announcement. We do not insist that the dream-possibility must always be known not to obtain in order to know things in everyday or scientific life. When we find that Descartes's sceptical reasoning does insist on that requirement, we will find that his sceptical conclusion does not contradict anything we thought we knew at the outset. We might find it quite believable that there is no knowledge of the world fulfilling all the conditions of Descartes's special 're-definition' of knowledge. But properly understood, his conclusion would not deny what its peculiar linguistic form originally led us to suppose it denies, and it would pose no threat to our everyday knowledge and beliefs. Any exhilaration or disquiet we might have felt on first encountering it must therefore have been due to nothing but illusion.

If there were nothing more behind Descartes's sceptical conclusion than there is behind the peculiar announcement about physicians in New York it would indeed be profoundly uninteresting. If Descartes simply imposes on knowledge an unreasonable or outrageous requirement, and then points out (even quite correctly) that it can never be fulfilled, there will be no reason to go along with him, even temporarily. What he says would reveal nothing more about the everyday or scientific knowledge that we want a philosophical theory to illuminate than that crazy announcement manages to reveal about physicians in New York. Someone is no less a physician even though there are many patients he never happens to cure, and if Descartes is simply distorting the requirements for knowledge, what we possess in everyday life and in science will be no less knowledge even though we do not usually fulfil the outrageous condition that we must know we are not dreaming. I think many philosophers find philosophical scepticism uninteresting and the study of it unprofitable on grounds such as these. Descartes's assessment of his own position is thought to deviate so radically and so obviously from our familiar assessments that it cannot be expected to reveal anything of deep or lasting significance about the human knowledge we are interested in.

It is perhaps not so immediately obvious that a change or distortion of meaning has occurred in the philosophical case

as it is in the announcement about physicians in New York. If we are at all taken in by the sceptical reasoning, the misunderstanding must be to that extent hidden from us, just as it is presumably hidden from the sceptical philosopher himself. Giving that special meaning to the word 'physician' is nothing more than a crazy whim. But the philosopher's alleged change in the meaning of the word 'know' might not be unmotivated; certainly it is not just a personal whim. What lies behind the sceptical conclusion might therefore turn out to be more interesting and more worthy of investigation than what lies behind that 're-definition' of 'physician'.

But still, it will be felt, the philosophical case will be interesting only to the extent to which it is interesting to find out how and why philosophers so persistently go wrong—why they continue to insist, as they apparently do, on misunderstanding or distorting the meanings of the familiar words they examine and use. The investigation of philosophical scepticism would then be of pathological interest only. Aside from revealing how easy it is for philosophers to fall into confusion or make mistakes, it could not be expected to reveal anything deep or of lasting significance about human knowledge itself.

J. L. Austin, for example, thought an inquiry into the sources of the sceptical conclusion was 'a matter of unpicking, one by one, a mass of seductive (mainly verbal) fallacies, or exposing a wide variety of concealed motives—an operation which leaves us, in a sense, just where we began'.² In a positive vein, he thought we might at most learn something about the meanings of certain English words that are interesting in their own right.³ Many recent philosophers who care less than Austin did about the meanings of those English words would hold that if we simply keep our wits about us and guard against the errors that have led older philosophers astray we will find no reason to follow them down the garden path to philosophical scepticism. The misguided sceptical conclusion is held to reveal nothing about our everyday or scientific knowledge and beliefs because it is not really about that knowledge or those beliefs at all, any more than

² J. L. Austin, *Sense and Sensibilia* (Oxford, 1962), pp. 4–5.

³ *Sense and Sensibilia*, p. 5.

that crazy announcement is really about the physicians in New York.

I have tried so far to suggest that the best strategy in the face of the sceptical argument is to examine more carefully the requirement that we must know we are not dreaming if we are to know anything about the world around us. I think that is much more promising than accepting Descartes's condition as a genuine condition of knowledge and trying to show that it can be fulfilled. I argued that that strategy cannot succeed. If it is in general a necessary condition of our knowing something about the world around us that we know we are not dreaming, it follows that we can never know that we are not dreaming. That is why I think the only hope lies in avoiding that condition. But I do not share the impression that what Descartes says is a condition of knowledge of the world is obviously no such condition at all. There seems to me to be no question that the meaning of 'physician' has been changed in that trivial example, but in Descartes's reasoning I think much deeper and more complex issues are raised. And what is at stake is more than simply a mass of avoidable mistakes or confusions by traditional philosophers. I think the right kind of investigation into the sources of Descartes's requirement promises to illuminate something about our actual conception of knowledge, or about what we seek when we try to understand it, or perhaps even about human knowledge itself.

Let us suppose for the moment that the critics are right, and that what Descartes says is a requirement for knowledge of the world is really no such requirement at all. How could that be known, if it were true? What shows or would show that Descartes is or must be distorting or misunderstanding what knowledge is when he insists that we must know we are not dreaming if we are to know anything about the world around us? When critics of Descartes's conclusion argue that the meaning of 'know' does not in fact require what Descartes apparently requires of it, that knowledge is not 'closed under logical consequence', or that the word 'know' does not 'penetrate' to all the logical consequences of what is known, or to what are known to be its logical consequences, or even to what are known to be the consequences of knowing

it,⁴ how are such claims about knowledge or about the meaning of the word 'know' themselves to be supported? I will try to bring out some of the difficulties this question raises by looking at one of the most persuasive and most influential versions of that line of criticism.

The sceptical conclusion is reached in the course of an assessment of our knowledge of the world—an investigation of how we know the things we think we know about the world around us. J. L. Austin thinks philosophers in the course of such assessments have not paid sufficient attention to 'what sort of thing does actually happen when ordinary people are asked "How do you know?"',⁵ and in his 'Other Minds' he tries to show how the typical philosophical investigation deviates from our normal practices.

If asked how I know there is a goldfinch in the garden, for example, I might reply by explaining how I have come to know about goldfinches, or about small British birds in general, or I might explain how I came to be in a position to recognize and hence to know about the goldfinch in the garden in this particular case. This second kind of reply to the question 'How do you know?' might be inadequate because my system of classification is wrong—what I think are goldfinches are really something else—or my response might be challenged on the grounds that what I have said about how I know is not enough. If I said I knew it was a goldfinch by its red head, it might be objected 'But that's not enough: plenty of other birds have red heads. What you say doesn't prove it. For all you know, it may be a woodpecker' (OM, 51). This amounts to raising a possibility compatible with everything I have said but which, if actual, would imply that I do not know that there is a goldfinch in the garden. It therefore brings us close to the kind of objection Descartes raises against our ordinary knowledge of the world around us.

⁴ Many recent philosophers have argued against scepticism on some such grounds. One version of the idea is worked out in some detail in F. Dretske, 'Reasons and Consequences', *Analysis*, 1968; 'Epistemic Operators', *The Journal of Philosophy*, 1971. For a more recent version along the same lines see R. Nozick, *Philosophical Explanations* (Cambridge, Mass., 1981), ch. 3. I believe the basic idea is to be found in Austin.

⁵ J. L. Austin, 'Other Minds', in his *Philosophical Papers* (Oxford, 1961), p. 45. (Hereafter cited as OM.)

Austin thinks philosophers in their assessments tend to concentrate on questions about 'reality' and to some extent about 'sureness and certainty', and that although questions of those kinds are of course raised about knowledge in everyday life, the philosopher's special investigation of knowledge distorts or abandons our everyday procedures for answering them. In objecting to a piece of putative knowledge on the grounds that what has been said is not enough or does not prove what is claimed to be known, he says, we all ordinarily accept that:

- (a) If you say 'That's not enough', then you must have in mind some more or less definite lack . . . If there is no definite lack, which you are at least prepared to specify on being pressed, then it's silly (outrageous) just to go on saying 'That's not enough'.
- (b) Enough is enough: it doesn't mean everything. Enough means enough to show that (within reason, and for present intents and purposes) it 'can't' be anything else, there is no room for an alternative, competing description of it. It does *not* mean, for example, enough to show it isn't a *stuffed* goldfinch. (OM, 52.)

When philosophers go on to raise questions about 'reality' ('But do you know it's a *real* goldfinch?') they intend to question the reliability of the 'facts' put forward in support of the original claim to know. That too, of course, is something we do in everyday life. Austin thinks philosophers do not always satisfy the above conditions when they press our ordinary knowledge-claims in this way.

The doubt or question 'But is it a *real* one?' has always (must have) a special basis, there must be some 'reason for suggesting' that it isn't real, in the sense of some specific way, or limited number of specific ways, in which it is suggested that this experience or item may be phoney. Sometimes (usually) the context makes it clear what the suggestion is: . . . If the context doesn't make it clear, then I am entitled to ask 'How do you mean? Do you mean it may be stuffed or what? *What are you suggesting?*' (OM, 55.)

Austin suggests that a philosopher interested in knowledge—or at any rate someone he calls 'the metaphysician'—does not fulfil this condition in his typical challenges. His 'wile' consists in asking 'Is it a real table?' without specifying or limiting the ways he has in mind in which it might not be real. This leaves us at a loss in trying to answer him, just as

we are left baffled and uneasy by the conjurer's invitation 'Will some gentleman kindly satisfy himself that this is a perfectly ordinary hat?' (OM, 55n). It is no wonder we feel the philosophical objection to our ordinary knowledge cannot be met if this is what it trades on.

It should be clear that this unflattering description of the philosopher's or 'metaphysician's' procedure does not apply to Descartes's argument as I have outlined it. In his assessment of his claim to know that he is sitting by the fire with a piece of paper in his hand he does not simply complain in general terms that his grounds are not sufficient to prove that he is really sitting there. He is fully prepared to specify, in fact the whole force of his argument turns on his explicitly specifying, a particular way in which his grounds are inadequate, a particular possibility compatible with the facts he is relying on but incompatible with his knowing that he is really sitting there. His grounds are found inadequate in a perfectly determinate way; he might be dreaming. The 'wile of the metaphysician' as Austin describes it cannot explain why it is difficult or impossible to meet Descartes's objection to our knowledge of the world.

Austin might still be right that Descartes does violate our ordinary standards or procedures in another closely-related way. Once it has been made determinate precisely what question must be answered before one can be said to know in a particular case, Austin says, the question can then be answered 'by means of recognized procedures (more or less roughly recognized, of course) appropriate to the particular type of case' (OM, 55). In fact Austin strongly suggests, without saying so explicitly, that the existence of such 'recognized procedures' is a simple consequence of the determinateness of the original criticism of the knowledge-claim; as soon as it is made clear what doubt or deficiency the critic of knowledge has in mind (e.g., 'How do you know you are not dreaming?'), it will follow that there are recognized procedures for making up the deficiency or allaying the doubt.⁶ There are recognized ways of distinguishing between

⁶ Austin's reason for believing that this connection holds is probably to be found in his conception of what he calls 'the normal procedure of language', which he 'schematizes' as follows: if that complex of features originally taken

dreaming and waking', Austin says, '(how otherwise should we know how to use and to contrast the words?), and of deciding whether a thing is stuffed or live, and so forth' (OM, 55).

Austin does not say much about what he thinks the 'procedures' or 'recognized ways' of telling that one is not dreaming actually are. He seems content with the idea that there must be such procedures or else we would not be able to use and to contrast the words 'dreaming' and 'waking' as we do. I find that particular claim dubious, or at the very least difficult to establish, partly for reasons I will return to later.⁷ The reason he gives in his lectures is even less persuasive. In *Sense and Sensibilia* he denies the philosopher's contention that there is no 'qualitative difference' between normal waking experience and dream experience. He argues that actually being presented to the Pope is not 'qualitatively indistinguishable' from dreaming that I am being presented to the Pope on the grounds that:

After all, we have the phrase 'a dream-like quality': some waking experiences are said to have this dream-like quality, and some artists and writers occasionally try to impart it, usually with scant success, to their works. But of course, if the fact here alleged *were* a fact, the phrase would be perfectly meaningless, because applicable to everything. If dreams were not 'qualitatively' different from waking experiences, then *every* waking experience would be like a dream; the dream-like quality would be, not difficult to capture, but impossible to avoid.⁸

Someone who believed that our ability to tell that we are not dreaming on particular occasions is guaranteed by the very meaningfulness of certain English expressions would

be sufficient for saying 'This is a C' comes to be 'accompanied or followed in definite circumstances by another special and distinctive feature or complex of features, which makes it seem desirable to revise our ideas' we will 'draw a distinction between "This looks like a C, but in fact is only a dummy, &c." and "This is a real C (live, genuine, &c.)" . . . If the special distinctive feature is one which does not have to manifest itself in *any* definite circumstances (on application of some specific test, after some limited lapse of time, &c.) then it is not a suitable feature on which to base a distinction between "real" and "dummy, imaginary, &c." (OM, 57).

⁷ See pp. 72 ff.

⁸ *Sense and Sensibilia*, pp. 48-9.

perhaps feel little need to describe those 'procedures' carefully or to explain exactly how they work. He would already be convinced that they must work, so even without a detailed description of those 'procedures' he would seem to be directly in conflict with Descartes's reasoning. For Descartes at the end of his first *Meditation* it is impossible to know that we are not dreaming, so it might look as if the whole issue between him and Austin turns on the cogency of this appeal to the meaningfulness of the expression 'a dream-like quality' or to our ability to use and to contrast the words 'dreaming' and 'waking' as we do. But in fact I think Austin's real opposition to the sceptical philosophical conclusion is to be found elsewhere.

Descartes's conclusion rests on the general requirement that we must know that we are not dreaming if we are to know anything at all about the world about us. That requirement is what renders inadequate any tests or procedures for determining that one is not dreaming; one would have to know that one was not simply dreaming that one was performing the test, and not dreaming that one was performing any of the other tests used to determine *that*, and so on. For Austin it is precisely in insisting on that strong general condition for knowledge that the real distortion or unreasonableness comes in. If it is not in general a condition of knowing things about the world around us that we must know that we are not dreaming, not only will Descartes have failed to show that we can never know that we are not dreaming (and that there can be no 'procedures' of the kind Austin has in mind), he will not have begun to show that we cannot know anything about the world around us either. Without Descartes's condition for knowledge, philosophical scepticism about the external world would be completely disarmed. Austin attacks what is really the heart of Descartes's position. Can it be shown that in insisting on his strong general condition for knowledge of the world Descartes is violating or abandoning the ordinary conditions or standards of knowledge?

I have already said that a moment's reflection seems enough to convince us that Descartes's condition is not in fact a condition of knowledge in everyday or scientific life.

After thinking about philosophical scepticism for some time we often tend to forget or distort what we actually do in everyday life, but if we insist on returning to a realistic account of how we actually behave there seems little doubt that we do not in fact impose that general condition on our knowledge-claims.

For example, suppose I remark to the bird-lovers at my cocktail party that a goldfinch has just appeared in the garden. 'Really? How do you know?', one of them asks, and I reply that I just saw it hop from one limb to another in that large pine tree. 'How do you know you're not dreaming?', asks another, in what would obviously be no better than a feeble attempt at a joke. There is no reason to take what he says seriously, and none of us would. We do not regard it as a threat to my knowledge. Suppose I am in a bird-recognition contest. I examine my specimen carefully, noting its differences from birds of similar but distinct species, and I announce that I now know that this one is a goldfinch. Could one of the judges at that point ask me how I know I didn't simply dream it, and then reject my answer because I cannot give a satisfactory defence? That would be perfectly outrageous, and I would feel no necessity to answer the question in order for my original claim to stand.

These are trivial examples, but the inappropriateness of insisting on Descartes's condition does not stem from the relative unimportance of the knowledge in question. Even when it matters a great deal, when it is literally a matter of life or death, as in a court of law, it is simply not true that the dream-possibility is always allowed as a relevant consideration for the claim to know some particular thing. If I testify on the witness stand that I spent the day with the defendant, that I went to the museum and then had dinner with him, and left him about midnight, my testimony under normal circumstances would not be affected in any way by my inability to answer if the prosecutor were then to ask 'How do you know you didn't dream the whole thing?'. The question is outrageous; it has no tendency to undermine my knowledge. It is nothing more than the desperate reaction of a hard-pressed lawyer with no case. Nor do we ever expect

to find a careful report of the procedures and results of an elaborate experiment in chemistry followed by an account of how the experimenter determined that he was not simply dreaming that he was conducting the experiment. No such thing was in question; the issue is never raised, let alone settled. The point is obvious, and the multiplication of further examples is unnecessary. It seems clear that we simply do not insist on fulfilling Descartes's condition in order to know things in real life. Nor do we insist that someone cure every conceivable illness in less than two minutes in order to be regarded as a physician.

Of course it is sometimes relevant to ask how or whether we know we are not dreaming. When that is a relevant criticism of a claim to know something, failure to answer the question satisfactorily would imply that we do not know what we claim to know. If I am lying half-awake in bed early in the morning after a late night and seem to hear someone calling my name from outside the window, I might not be sure whether there really is someone out there or I am only dreaming that I hear the call. I do not know whether there is someone out there or not. But from the fact that that possibility is sometimes relevant it does not follow that on every occasion we must know that it does not obtain if we are to know anything about the world around us. When the alarm-clock has sounded and I have reached out and turned it off and got out of bed and gone over to the window and opened the curtains and found my friend calling and gesticulating in the garden, there is no question at that point that I might be dreaming or that I should check to see whether I am dreaming before I can know that he is really there—even though I can truly say to him that I didn't know he was there a few minutes ago because I didn't know whether or not I was dreaming.

As Austin puts it in the case of the goldfinch:

Knowing it's a 'real' goldfinch isn't in question in the ordinary case when I say I know it's a goldfinch: reasonable precautions only are taken. But when it *is* called in question, in *special* cases, then I make sure it's a real goldfinch in ways essentially similar to those in which I made sure it was a goldfinch. . . . (OM, 56.)

That there are such ways of making sure that it is a real goldfinch does not of course guarantee that we can always tell that it is, nor is it a proof against 'miracles or outrages of nature' (OM, 56). Something might still go wrong, something completely unexpected might happen to the bird, but that by itself is no bar to saying, or having been right to say, that we know it is a real goldfinch.

It is only in *special* circumstances that certain kinds of possibilities are relevant to claims to know something. Austin makes the point in connection with claims to know something about the mind or feelings of another person. Worries about deception, or about whether the person is sufficiently like us to be feeling what we would feel, or about whether he is quite inadvertently behaving as he does, all arise only in 'special cases'. Again, there are (more or less roughly) established procedures for dealing with such cases when they arise, but:

These special cases where doubts arise and require resolving, are contrasted with the normal cases which hold the field¹ *unless* there is some special suggestion that deceit &c., is involved, and deceit, moreover, of an intelligible kind in the circumstances, that is, of a kind that can be looked into because motive, &c., is specially suggested. There is no suggestion that I *never* know what other people's emotions are, nor yet that in particular cases I might be wrong for no special reason or in no special way.

[¹Austin's footnote: 'You cannot fool all of the people all of the time' is 'analytic'.] (OM, 81.)

Austin's stress here on the need for special reasons to doubt when questions of 'reality' are at issue is not the same as the earlier point that there must always be some 'special basis' for the doubt or question 'But is it a *real* one?'. That requirement was expressed as the demand that the critic have 'some "reason for suggesting" that it isn't real, in the sense of some specific way, or limited number of specific ways, in which it is suggested that this experience or item may be phoney' (OM, 55). Descartes in his reasoning meets that requirement as stated: he specifies dreaming as the way the experience might be 'phoney'. But here Austin is arguing that even if the way the experience or item might be 'phoney' has been specified, the doubt or question 'But is it a *real*

using the expressions 'I know' and 'I promise' as we do in fact use them. (OM, 66.)

It would be no easy matter to give a precise formulation of the requirement that there must be some special reason to think a certain possibility might obtain before the raising of that possibility is allowed as a relevant criticism of a claim to know something. Is it enough for there simply to be such a reason, or must someone actually have that reason, and raise the possibility for that reason? How concrete or specific must the reason be, and how good a reason must there be for thinking something is amiss in the particular case? I want to leave aside all such questions of detail and ask about the conflict between *any* requirement along those general lines and the condition that Descartes insists on for knowledge about the world around us. I therefore want to grant everything Austin says about what sort of thing does actually happen when ordinary people are asked 'How do you know?', and everything else that could be discovered about how we respond to the questions or would-be challenges of others with respect to our knowledge. I want to concentrate on the question: do such facts about our everyday and scientific practices show that Descartes's reasoning deviates from our everyday procedures and standards for acquiring and assessing knowledge?

It certainly looks as if Descartes could not be right in insisting that we must rule out the dream-possibility in order to know something about the world around us if Austin is right about how the raising of such possibilities can work against our knowledge in everyday life. If there must be some special reason for suggesting or suspecting that one is dreaming before that reason for doubt is even allowed as relevant in everyday life, the most that is true of the dream-possibility with respect to our knowledge of the world is that it must be known not to obtain whenever there is some special reason to think it might obtain. That is to say, if there is some special or concrete reason to believe that one might be dreaming, one cannot know some particular thing about the world around us unless one knows that one is not dreaming. That is obviously weaker than Descartes's general requirement

one?' is relevant to the original knowledge-claim and must be answered only if there is some special reason for suggesting that that specified possibility might obtain. It is not simply that the critic of the knowledge-claim must specify some way in which knowledge would not be present on the occasion in question; he must also have some reason for thinking or suggesting that the possible deficiency he has in mind might be present on that occasion. In the absence of such a reason—that is, in the normal or non-special case—knowing it is a real goldfinch, for example, is not in question. The 'reasonable precautions' said to be taken in the ordinary case are precautions against only those possibilities that there is some special reason to think might obtain in that case. Whether or not such possibilities obtain is all that is in question.

The need for a special reason for doubt is also present when we cite authorities or rely on the testimony of others—a rich source of knowledge not much studied by philosophers.

Naturally, we are judicious: we don't say we know (at second hand) if there is any special reason to doubt the testimony: but there has to be *some* reason. It is fundamental in talking (as in other matters) that we are entitled to trust others, except in so far as there is some concrete reason to distrust them. (OM, 50.)

The same holds for any other possibility of error or mistake. There can be no doubt that human beings are inherently liable to be mistaken in particular claims to know things—and not just things we know 'at second hand' or from testimony. But the question 'How do you know?' is not a successful challenge if it is based only on such general human fallibility. That is not to deny that knowledge precludes error or mistake.

'When you know you can't be wrong' is perfectly good sense. You are prohibited from saying 'I know it is so, but I may be wrong', just as you are prohibited from saying 'I promise I will, but I may fail'. If you are aware you may be mistaken, you ought not to say you know. . . . But of course, being aware that you may be mistaken doesn't mean merely being aware that you are a fallible human being: it means that you have some concrete reason to suppose that you may be mistaken in this case . . . It is naturally *always* possible ('humanly' possible) that I may be mistaken or may break my word, but that by itself is no bar against

which says that one cannot know any particular thing about the world around us unless one knows that one is not dreaming. Descartes's reasoning imposes a condition on knowledge of the world which must be fulfilled in every case, whether there is any special reason to believe one might be dreaming or not. The weaker requirement says that that condition must be fulfilled only in some cases, when the 'special reason' condition is also fulfilled, but that otherwise the dream-possibility is not even relevant to our claims to know things about the world around us.

Another way to put the difference is that the weaker requirement allows for the possibility of knowledge of the world in a way that Descartes's requirement does not. I have tried to show how Descartes's requirement is strong enough to make knowledge of the world around us impossible; it precludes fulfilment of the very condition it holds to be necessary for knowledge of the world. For all the weaker requirement says, one *could* know things about the world around us without knowing that one is not dreaming. With no special reason to think one might be dreaming, that alleged possibility is simply not in question at all, so the possibility Descartes invokes would present no obstacle to knowing things about the world in those cases. It therefore looks as if the sceptical reasoning cannot succeed if only the weaker requirement, and not Descartes's condition for knowledge, is what is true of our ordinary conception of knowledge.

Descartes's reasoning itself cannot be said to fulfil only the weaker requirement. He considers his knowledge of the world around him in general by considering the particular case of his sitting by the fire with a piece of paper in his hand. That single case is chosen to serve as a representative of all of our knowledge of the world. It could sustain a quite general conclusion about all of our knowledge of the world only if it were a perfectly normal case, without special features. If Descartes had had some special reason to doubt the deliverances of his senses at that particular time and place—if it were early morning, for example, and he was not quite sure whether he was fully awake or not—his verdict could not support the kind of general conclusion he draws from it about our sensory knowledge in general. It would just

be one special, non-representative case in which, as it turns out, he fails to know. But if there is nothing special about the case, if he had no special reason to think he might be dreaming at that particular time and place, his challenge 'How do I know I am not dreaming?' will have no special basis. The possibility will be raised without any special or concrete reason for supposing in this case that he might be dreaming, so it would seem that it must violate what Austin says is a condition of its being a serious or even a relevant challenge to a knowledge-claim in everyday life.

Despite this apparently obvious conflict between our everyday practices and Descartes's requirement I still want to press the question whether all the facts about how we speak and respond to the questions or would-be challenges of others in everyday life are enough to show that Descartes in his reasoning deviates from our everyday standards and procedures and changes or distorts the meaning of the word 'know' or any of the other words he uses. I think those facts would not have that anti-sceptical consequence if a certain conception of everyday life, and hence a certain conception of the relation between the philosophical problem of the external world and what goes on in everyday life, were correct. In looking at the significance of those facts I will therefore be looking from a different angle at the significance of the philosophical problem itself and at what, if anything, it can reveal about our position and procedures in everyday life. If the philosophical sceptic's conception of everyday life is intelligible, everything that goes on in everyday life and in science would be compatible with the literal truth of the conclusion that no one knows anything about the world around us.

If only Austin's weaker, and not Descartes's stronger, requirement is true of our ordinary conception of knowledge, it nevertheless should state a truth about *knowledge*. What I mean by such an obvious remark is this. If we have a conception of knowledge that we employ in everyday life prior to and independently of all philosophizing, and if Descartes or some other philosopher is to be shown to have changed or distorted that conception in the course of his philosophizing, it must be shown that it is that very conception that he

has distorted, and that that conception is a conception of knowledge. Austin's weaker requirement, for example, will state a condition of knowledge only if it implies that in those cases in which there is no special reason to think one might be dreaming, and one fulfils all the other conditions for knowing, one does indeed *know* something about the world around us without knowing that one is not dreaming. I stress the point because I think that speaking very strictly, or on a certain conception of our linguistic and other behaviour in everyday life, facts of the kind cited by Austin do not actually have that implication.

What Austin reminds us of are facts of speech, of linguistic usage. When he describes what sort of thing does actually happen when ordinary people are asked 'How do you know?' he tells us what people say, and what conditions must be fulfilled in order for them to say it, or to be speaking correctly in saying it. In the passages already quoted, for example, he says:

If there is no definite lack, which you are at least prepared to specify on being pressed, then it's silly (outrageous) just to go on saying 'That's not enough'.

Knowing it's a 'real' goldfinch isn't in question in the ordinary case when I say I know it's a goldfinch: we don't say we know . . . if there is any special reason to doubt . . .

If you are aware you may be mistaken, you ought not to say you know . . .

[Its being always possible that I may be mistaken] is no bar against using the expressions 'I know' and 'I promise' as we do in fact use them.

I have italicized the crucial words to show that in each case what is in question is how certain expressions are or should be used. Similar facts are appealed to when Austin says:

we are often right to say we *know* even in cases where we turn out subsequently to have been mistaken. (OM, 66.)
we may be perfectly justified in saying we know or we promise, in spite of the fact that things 'may' turn out badly, and it's a more or less serious matter for us if they do. (OM, 69.)

I have said that I want to grant all facts of the kind Austin here describes. But in order to move from such facts about

the use of the expressions 'I know' or 'He knows' or 'How do you know?' to conclusions about knowledge we must at least know *why* those expressions are used as they are on those occasions.

There are also facts about our reactions to the raising of various possibilities in ordinary circumstances that seem to support Austin's conclusion. I pointed out in earlier examples that we simply ignore, perhaps with some embarrassment, the party guest's question how I know I didn't dream the goldfinch in the garden. We find the bird-contest judge's challenge incomprehensible, and do not regard it as affecting in any way the truth or reasonableness of my identification. We immediately throw out of court the desperate prosecutor's ridiculous challenge and proceed as if it never had been made. We do not expect, and would be astonished to find, dream-elimination tests appended to laboratory reports; they would not affect one way or the other our acceptance of the results reported. When something serious is in question we usually find it silly (outrageous) for someone to persist in asking how we know we are not dreaming and to insist that we don't know what we thought we knew until we can answer the question satisfactorily. These all seem to me undeniable facts of everyday life.

The question is whether all this linguistic and other behaviour is generated by or warranted by or even required by our everyday conception of knowledge. Whether that is so or not depends on why we behave in those ways in everyday life. It is admittedly bizarre, silly, outrageous, perhaps even incomprehensible, to raise the dream-possibility as a criticism of ordinary claims to know things in everyday and scientific life, but exactly what kind of outrageousness or inappropriateness is it? What is its source? Is it derived from our very conception of knowledge itself? Is anyone who raises the possibility in normal circumstances necessarily violating or rejecting the everyday meaning of the word 'know'?

These questions arise because there are two apparently distinct questions that can be asked about what someone says. We can ask whether it is true, or we can ask whether it was appropriately or reasonably said. The two questions do

not always get the same answer; certainly it is possible for them to differ. All the conditions sufficient for appropriate or reasonable utterance can be fulfilled when what is said is not literally true. The distinction even more obviously can be made in the other direction; there are countless things that are now true which no one is now in a position reasonably to assert or believe—many people are busily engaged in trying to find out what some of them are. I do not mean to suggest that there is or must be some sort of conflict or opposition, let alone an unbridgeable gap, between reasonable utterance and truth. There is nothing in the distinction itself which suggests that the truth is forever—or ever—beyond us. Normally, we believe, the conditions of reasonable utterance coincide with the conditions of truth. We usually take it for granted that what we are in a good position to assert is in fact true and what we are in a good position to deny is in fact false. In trying to find out whether or not some particular thing is true we try to get into the best possible position for accepting and asserting it, or for rejecting and denying it. The point is only that the two sets of conditions can be distinguished. From the fact that someone carefully, reasonably and appropriately asserts something on a particular occasion it does not directly follow that what he says is true, and from the fact that someone quite inappropriately and with no good reason says something it does not immediately follow that what he says is false. This holds just as much for assertions or denials of knowledge as for other assertions and denials.

For example, suppose I am at a party and my host asks me if I know whether my friend John, who was ill last week, will be coming to the party. I reply that I know he will be there, and when asked how I know I explain that he has now recovered, I have just talked to him on the telephone and he said he was coming right over; there is someone at the party he is interested in talking to and he wouldn't miss it for anything. Suppose further that John is well-known to be generally trustworthy, reliable, and also a careful, sober driver—and he doesn't live very far away. All this puts my assertion that I know John will be at the party beyond criticism. There could hardly be more favourable grounds

for claiming knowledge about something not currently under my direct observation. Suppose now that John for some reason unknown to me nevertheless fails to show up at the party. My saying that he would be there, in fact that I knew he would be there, was justified, reasonable, and appropriate in the circumstances, but it has turned out that what I said is not true. John is not at the party, so I did not know what I said. I knew. The best possible conditions for asserting something, or that I knew something, did not coincide with the conditions under which what I said is true.

Suppose that as I am leaving the party at the end of the evening and John has still not appeared my host turns on me and says 'You should be more careful about what you claim to know. You said you knew John would be here and he isn't. You didn't know any such thing!' I think we find that response simply outrageous. It is absurd and improper and completely unjustified. It is difficult to find the right words to describe its degree of insensitivity and social obtuseness. It perhaps even shows incomprehension of how and why we claim to know things in the face of the normal vicissitudes of life. But aside from the unreasonable abuse and insensitivity conveyed by it, the remark cannot be said to be totally false or without foundation. Part of what the host said was 'You said you knew John would be here and he isn't. You didn't know any such thing', and that is, at least, the literal truth, however harsh or inappropriate it might have been for him to say it. I did say I knew, and I didn't know. What the cruel host says is an accurate description of my position.

It is clear that the host's remarks are outrageous or unreasonable as criticisms of *me*, or of my having *said* that I know John will be there. My response when asked whether I knew John would be at the party was justified, reasonable, appropriate, and perfectly proper. It is not open to the kind of attack the host tries to subject it to. But what is invulnerable to those absurd attacks is my act of *saying* something, and also perhaps my coming to believe or to accept something. My asserting it is beyond criticism even if what I assert is (of course unknown to me) not true. And the host's remark about the state of my knowledge is true even if his

was due to the falsity of what I claimed to know. That necessary condition of knowledge was unfulfilled even though no one at the time was in a position to know that it was unfulfilled, and no one at the time was in a position appropriately or reasonably to criticize my claim on that basis. Perhaps the same is true of other necessary conditions of knowledge.

Imagine a slightly different sequence of events at the party from the one described so far. Suppose that, as soon as I had hung up the telephone from talking with John and had said that I knew he would be at the party, the boorish host had said 'But do you really know he'll be here? After all, how do you know he won't be struck down by a meteorite on the way over? You don't know he won't be.' We find this at least as outrageous as his response in the other story after the truth about John and the meteorite was known. Not only is this 'challenge' as unfair and as inappropriate as that other response, it is difficult to understand why he even brings up such a consideration at this point and thinks it is a relevant criticism. His doing so would normally suggest that he thinks there have been a lot of meteorites hitting the earth lately in this general area, some of them rather big and capable of causing harm. If that were so, perhaps I should have thought of it and considered it—or at least if I didn't know about it my ignorance might threaten my claim to know John would be there. But in the absence of any such special reason the 'challenge' seems just as outrageous as it did in the other story.

My act of asserting that I knew John would be at the party was made on just about the most favourable grounds one can have for claiming to know things. It is no reflection on me or on my saying what I did that I had not ruled out or even thought of the meteorite possibility. But once the question is asked, however inappropriately, can it be said that I do know that that possibility will not obtain? It seems to me that it cannot. When I hung up the telephone I do not think I could be described as knowing that John will not be hit by a meteorite. As it turned out, of course, I did not know it because it wasn't true—he *was* hit by a meteorite. But even if things had turned out differently—even if John had

making it is outrageous, unreasonable, and unjustified. So even if we know that a certain attempt to criticize a knowledge-claim is outrageous or unreasonable or would not be listened to in everyday life, we cannot immediately infer that the knowledge-claim does not suffer from the deficiency stated in the criticism, or that the person does nevertheless know what he claims to know. Whether that is so or not will depend on the nature and source of the outrageousness or inappropriateness in question. The inappropriately-asserted objection to the knowledge-claim might not be an outrageous violation of the conditions of knowledge, but rather an outrageous violation of the conditions for the appropriate assessment and acceptance of *assertions* of knowledge. John's being at the party is an admitted condition of my knowing that he will be at the party—since he is not there I do not know what I said I knew—but his not being there does not warrant any and all criticisms of my saying or believing what I did. My saying and believing in those circumstances is beyond criticism. But still I did not know.

Suppose that in fact the reason John never arrived at the party was that just as he stepped out his front door he was struck down by a meteorite, the only one of lethal size known to have hit the earth in a century and a half. Of course there was no special reason—no reason at all—for me or for anyone else to expect such a thing. Certainly I never thought of it (nor of its not happening) when I said I knew John would be at the party, and it is safe to say that I knew John would be at the party did either. If the news of John's fate had been conveyed to all of us just as I was about to leave the party, the host's parting shot as I described it a moment ago would be if anything even more outrageous and inappropriate. When we all discover why John never got to the party and hence why I failed to know, and we see how bizarre and how completely unforeseeable his not getting there was, it is even more difficult to understand what the host could think he was doing in saying to me 'You said you knew John would be here and he isn't. You didn't know any such thing!' But I think it cannot be denied that one thing the host was doing was speaking the truth.

My failure to know in this case as originally described

actually arrived at the party—I do not think it was true when I hung up the telephone that I knew John would not be hit by a meteorite. So again, part of what the host says is true. I did not know any such thing. But still I said I knew John would be at the party.

I want to be careful here. I want to emphasize that I am not saying that in this second case I do not know that John will be at the party because I do not know when hanging up the telephone that he will not be struck down by a meteorite on the way over. I am concerned at the moment not so much with the truth about my state of knowledge in the example as with the question of how we arrive at conclusions about the state of my knowledge on the basis of admitted facts about how we speak and assess our assertions of knowledge and respond to the assertions and attempted criticisms of others. We agree that the imagined response of the host is outrageous and unjustified; it is something he has no specific reason to bring up as a possible objection to what I have said. All I am saying at the moment is that it does not follow directly from the admitted outrageousness of his introducing that possibility that my ruling out the meteorite possibility is simply *not* a condition of my knowing that John will be at the party. Its being a necessary condition of my knowledge is so far at least compatible with the host's remarks' being inappropriate or outrageous, just as John's being at the party is a necessary condition of my knowing he will be there even if it is inappropriate or outrageous for the host to say that I didn't know he would be there because in fact that condition was not fulfilled. A necessary condition of knowledge might remain unfulfilled even though it would be outrageous for anyone to assert that it is or inappropriate for anyone to criticize my knowledge-claim on that basis. The inappropriateness or outrageousness might have some source other than the falsity of what is said or implied about knowledge.

I emphasize a mere possibility here—that a certain conclusion about knowledge *might* not be true even granted the facts about how we speak and how we react to the speech of others—because I think philosophers who have investigated knowledge in Descartes's way have a conception of everyday

life and everyday speech that would get between those facts of usage and that conclusion about knowledge. That is not to say that their conception must be correct, or even fully intelligible, but it should at least make us more cautious in inferring directly from a reasonable, justified, even exemplary assertion of knowledge on a particular occasion to the conclusion that on that occasion all the conditions of knowledge are in fact fulfilled. If it is possible for the conditions sufficient for appropriate or reasonable utterance to be fulfilled even though what is said is not literally true—and it does seem that that is possible—someone might be fully justified in saying he knows some particular thing about the world around him without its being true that he does know that thing. In particular, when someone claims to know something about the world without asking himself about or even thinking of a certain possibility, and that possibility, if realized, would mean that he does not know what he claims to know, he might fail to know in that situation precisely because he has not eliminated that possibility. If there were no special reason for him to consider that possibility, he might nevertheless be fully justified in saying he knows. I was fully justified in saying I knew John would be at the party even though I did not think of the possibility that a meteorite might strike him; there was no reason to think such a bizarre event might occur.

On the conception I have in mind, the requirement that there must be some 'special reason' for thinking a certain possibility might obtain in order for that possibility to be relevant to a particular knowledge-claim would be seen as a requirement on the appropriate or reasonable assertion of knowledge, but not necessarily as a requirement on knowledge itself. In the absence of such a 'special reason', one might perhaps be fully justified in saying 'I know that *p*' even though it is not true that one knows that *p*. Descartes reaches his sceptical conclusion about our knowledge of the world around us on the basis of a condition he holds is necessary for the truth of 'I know that *p*'. To show that our everyday notion of knowledge contains no such condition, but only the weaker requirement that would enable us to know things about the world without knowing we are not dreaming

as long as there were no special reason to think we might be dreaming, it would have to be shown that when no such 'special reason' is present, 'I know that *p*' will sometimes be true and not just justifiably asserted.

How is it to be shown that that weaker requirement, or any other description of the way we actually speak and respond to the assertions of others, does in fact state a condition of knowledge, as opposed to a condition of appropriately or justifiably saying that one knows? As long as it is even intelligible to suppose that there is a logical gap between the fulfilment of the conditions for appropriately making and assessing assertions of knowledge on the one hand, and the fulfilment of the conditions for the truth of those assertions on the other, evidence from usage or from our practice will not establish a conclusion about the conditions of knowledge. The charge of violating or altering the meaning of the word 'know' (or any other word) can therefore be laid at the sceptical philosopher's doorstep on the basis of such evidence only if a certain conception of meaning, a certain conception of everyday speech, and a certain conception of the relation between them, can all be shown to be mistaken, perhaps even unintelligible. Rejecting such a conception would involve much more than the simple rejection of an isolated and idiosyncratic 're-definition' of knowledge, and even much more than a simple denial of the initially starting conclusion that no one can know anything about the world around us. That is why I think discovering the source of Descartes's requirement might reveal something deep and important.

Descartes and other philosophers who have examined knowledge in the same way and have been led to sceptical conclusions are fully aware that the kinds of doubts or criticisms they raise in their philosophical investigations would not always be appropriately raised in everyday or scientific activity. That in itself does not show that they must be changing or misunderstanding the meaning of the word 'know' or any other words. By invoking the conception I have just mentioned, they would attribute the inappropriateness to something other than the notion of knowledge itself.

Descartes, for example, insists that his procedure of asking

what is beyond all possible doubt is to be followed only in the philosophical investigation of human knowledge and not in everyday life.

... we are to make use of this doubt only when we are engaged in contemplating the truth. For, as regards the conduct of our life, we are frequently obliged to follow opinions which are merely probable, because the opportunities for action would in most cases pass away before we could deliver ourselves from our doubts. (HR, 219-20.)

C. I. Lewis, another philosopher whose Cartesian examination of knowledge proceeds by the raising of possible doubts to ordinary claims to know things about the world, writes:

To quibble about such doubts will not, in most cases, be common sense. But we are not trying to weigh the degree of theoretical dubiety which common-sense practicality should take account of, but to arrive at an accurate analysis of knowledge.⁹

Both accounts stress the contrast between the practical and the theoretical, or between what is appropriate or required in action and what is appropriate or required in knowing the truth. The standards or procedures we follow in everyday life find their source in the exigencies of action and in the general conditions under which actions must be performed. In the case of action, unlike that of belief and knowledge, truth is not the only important consideration. Actions take place at different times and in changing conditions, so what is a perfectly reasonable thing to do in one situation is not equally reasonable—in fact, might be quite outrageous—to do in another. What is or is not an appropriate or reasonable thing to do is determined by the situation at hand, by one's aims or interests at the moment, by one's appraisal of the situation, and, as Descartes emphasizes, by the time at one's disposal. It would be silly to stand for a long time in a quickly filling bus trying to decide on the absolutely best place to sit. Since sitting somewhere in the bus is better than standing, although admittedly not as good as sitting in the best of all possible seats, the best thing to do is to sit down quickly. In general, our actions

⁹ C. I. Lewis, *An Analysis of Knowledge and Valuation* (La Salle, Illinois, 1946), p. 180.

are more likely to succeed and to produce satisfaction to the extent to which we can deliberate longer and more rigorously and with more and more information, but it is part of the very nature of practical life that we often cannot carry that process far enough to give us the kind of certainty we would otherwise like to have. We do the best we can in the circumstances.

These platitudes show that there is no general answer to the question of how certain we should be before we act, or what possibilities of failure we should be sure to eliminate before doing something. It will vary from case to case, and in each case it will depend on how serious it would be if the act failed, how important it is for it to succeed by a certain time, how it fares in competition on these and other grounds with alternative actions which might be performed instead, and so on. This holds just as much for the action of saying something, or saying that you know something, or ruling out certain possibilities before saying that you know something, as for other kinds of action. There is no general answer to the question of which possibilities we should rule out before we assert something or say we know it to be true. Checking our beliefs or justifying our claims to know something is itself something we do, and the desirability or reasonableness of doing it beyond a certain point must always be weighed against the desirability or reasonableness of doing all those other things incompatible with it. It is a practical question how much time, effort and ingenuity we should spend on supporting and checking our beliefs, so we might easily find that it would be silly or outrageous on a particular occasion to go on trying to eliminate a certain possibility. That is to say, it would be silly or outrageous in the circumstances to act in that way.

The doubts or possibilities considered by the philosopher investigating human knowledge are not put forward as relevant to such practical questions as whether to assert something or to say that you know it or to raise an objection to what someone else has said. They are thought relevant only to the question of whether one knows something—whether it is true that one knows—and not whether it is appropriate or reasonable to say that one knows. And if the dream-

possibility, for example, is a possibility that one must know not to obtain if one is to know some particular thing about the world around us, then one will simply not know that thing about the world if one has not been able to eliminate that possibility—even though it might be completely inappropriate or unreasonable on particular occasions in everyday life to insist on ruling out that possibility before saying that one knows.

One way to bring out what I think is the sceptical philosopher's conception of everyday life in relation to his epistemological project is to consider in some detail the following story adapted from an example of Thompson Clarke's.¹⁰ Suppose that in wartime people must be trained to identify aircraft and they are given a quick, uncomplicated course on the distinguishing features of different planes and how to recognize them. They learn from their manuals, for example, that if a plane has features x , y , and w it is an E , and if it has x , y , and z it is an F . A fully-trained and careful spotter on the job will not say that a particular plane is an F until he has found all three features, x , y , z . If at a certain point he has found only x and y and cannot yet tell what other features the plane has got, he does not know whether it is an F or an E . Once he finds that it also has feature z he can report that the plane in the sky is an F . He might even be asked how he knows it is an F and reply 'Because it has x , y , z '. He has observed the plane in the sky very carefully, he has followed his training to the letter, and he is right that it has x , y , z . There seems no doubt that he knows the plane is an F .

Suppose that there are in fact some other airplanes, G s, which also have features x , y , z . The trainees were never told about them because it would have made the recognition of F s too difficult; it is almost impossible to distinguish an F from a G from the ground. The policy of simplifying the whole operation by not mentioning G s in the training manual might be justified by the fact that there are not many of them, or that they are only reconnaissance planes, or that in some other ways they are not as directly dangerous as F s;

¹⁰ Thompson Clarke, 'The Legacy of Scepticism', *The Journal of Philosophy*, 1972, pp. 759 ff.

practical-purposes. After all, there are more important things in wartime—and even in peacetime everyday life—than knowledge.

I think the sceptical philosopher sees our position in everyday life as analogous to that of the airplane-spotters. There might be very good reasons why we do not normally eliminate or even consider countless possibilities which nevertheless strictly speaking must be known not to obtain if we are to know the sorts of things we claim to know. We therefore cannot conclude simply from our having carefully and conscientiously followed the standards and procedures of everyday life that we thereby know the things we ordinarily claim to know. The philosophical investigation of our knowledge is concerned with whether and how it is true that we know, whether and how the conditions necessary and sufficient for our knowing things about the world are fulfilled. Descartes's argument turned on its being a condition of our knowing any particular thing about the world around us that we know we are not dreaming, and on this conception the admitted fact that we do not insist on eliminating that possibility in everyday life does not show that we do not need to eliminate it in order to know things about the world. The well-trained airplane-spotter is not required to rule out the possibility that the plane he sees with features x , y , z is a G ; nor do his teachers or his fellow spotters insist on that possibility's being eliminated. But we recognize that it is nevertheless a condition of knowing that the plane is an F on the basis of x , y , z that one know that it is not a G . Facts about the way we speak and the procedures we follow in everyday life do not show that the sceptical philosopher has misunderstood or distorted the nature of knowledge if this conception of our everyday practices and procedures is correct.

The point is worth stressing. Many people are apparently disposed to think that if the philosopher holds that a certain condition must be met in order to know something, and we do not insist on that condition's being met in everyday life, then the philosopher simply *must* be imposing new or higher standards on knowledge or changing the meaning of the word 'know' or some other word. But if our position in everyday

it does not matter as much whether they fly over our territory.

When we are given this additional information I think we immediately see that even the most careful airplane-spotter does not know that a plane he sees is an F even though he knows that it has x , y , z . For all he knows, it might be a G . Just as he did not know the plane was an F when he had found only features x and y —for all he knew then, it might have been an E —so he does not know now that it is an F because all the features he has now found are also present on another kind of airplane. Of course there would be no point in telling him that he does not know; the same good reasons for not even mentioning G s in the training manual would still apply. In saying that he does not know it is an F we would be making no criticism of his performance; he has followed his training perfectly and conscientiously. We ourselves might not even care whether the plane in the sky is an F or a G ; it is precisely because G s are not worth worrying about in the same way that the manual was written as it was. But I think it must nevertheless be admitted that the spotter does not know that the plane is an F .

In saying that he does not know the plane is an F I do not mean to deny that he can be said to know 'for all practical purposes'. Whether it is an F or a G does not matter much; that is why the training could afford to ignore the differences between them. All things considered, it is best to have a policy of not distinguishing between the two kinds of planes when deciding what to do, how to respond to their presence. So as a contribution to the war effort his recognition is beyond criticism. We might even be tempted to say something like 'As far as his training goes, he knows it is an F ', or 'He knows that according to his manual it is an F '. But if we know the facts about G s I think we cannot say simply 'He knows it is an F '. When I say that we cannot say that he knows it is an F I mean that we recognize that that would not be true. We recognize that he does not know it is an F even though there is absolutely nothing to be gained by pointing his ignorance out to him or to anyone else. For all practical purposes we can accept his saying that he knows it's an F . He can perhaps be said to know-for-all-

life is like that of the airplane-spotters that is not so. When we who know the facts about Gs say that the careful spotter does not really know that the plane he sees is an F I do not think we are imposing on him or inventing for ourselves some new and unreasonably strict conception of knowledge. If we explained the situation to the spotter himself (which admittedly would not help the war effort) he too would agree that he did not know whether the plane was an F or a G. Just as he recognized earlier that the presence of x and y alone was not enough to settle the question whether the plane was an F, so he would see, with the new information, that even x , y , and z are not enough. He would see that he has to do more in order to know that it is an F. The fact that there is nothing more that a mere airplane-spotter can do—that it is almost impossible to distinguish an F from a G from the ground—would not alter that judgement. He would see that with the resources currently available to him he simply cannot know whether a plane in the sky is an F or not. But in coming to that conclusion he would not have altered the conception of knowledge with which he began. He originally understood the word 'know' and applied that conception of knowledge fully reasonably and justifiably in particular cases, but (as we knew all along and he now would come to realize) he never knew on any of those occasions that the plane in the sky was an F.

That is how I think the philosopher who investigates human knowledge sees the relation between what he concludes about knowledge and the way we speak about knowledge in everyday life. We do not ordinarily insist on the dream-possibility's being ruled out unless there is some special reason to think it might obtain; the philosopher insists that it must always be known not to obtain in order to know anything about the world around us. But on his understanding of everyday life that difference is not to be explained by the philosopher's insisting on or inventing a conception of knowledge stricter or more demanding than that of the scientist or the lawyer or the plain man. Rather he claims to share with all of us one and the same conception of knowledge—that very conception that operates in everyday and scientific life.

One thing the sceptical philosopher can appeal to to show that he does not introduce a new or extraordinary conception of knowledge into his investigation, I think, is the ease with which we all acknowledge, when presented with the case, that Descartes ought to know that he is not dreaming if he is to know that he is sitting by the fire with a piece of paper in his hand. The force we feel in the sceptical argument when we first encounter it is itself evidence that the conception of knowledge employed in the argument is the very conception we have been operating with all along. If we become even half persuaded that Descartes really should eliminate the dream-possibility, I think we do not have the sense that the knowledge for which that is now felt to be required is something different from the knowledge expressed in Descartes's original conviction that he knew he was sitting by the fire with a piece of paper in his hand. Nor do we think it is something different from the kind of knowledge we take ourselves to seek and to possess in everyday life. That is why the sceptical argument can seem to threaten our everyday knowledge. We are originally inclined to respond to it in the way the careful airplane-spotter would respond to the news about Gs. We realize that, strictly speaking, we must be able to eliminate the dream-possibility if we are to know anything about the world around us.

But of course we are also strongly inclined to reject the sceptical reasoning because what it would require of us deviates so radically from what we require of ourselves and others in everyday life. The sceptical philosopher has an explanation of that difference. There is a single conception of knowledge at work both in everyday life and in the philosophical investigation of human knowledge, but that conception operates in everyday life under the constraints of social practice and the exigencies of action, co-operation and communication. The practical social purposes served by our assertions and claims to know things in everyday life explain why we are normally satisfied with less than what, with detachment, we can be brought to acknowledge are the full conditions of knowledge. From the detached point of view—when only the question of whether we know is at issue—our interests and assertions in everyday life are seen

an F. There is a real difference between them. We from our more detached position will agree that the careless spotter does not know, and that the careful spotter did not know when he had found only x and y , but we will also agree that the careful spotter does not know the plane is an F even when he has also found it has z . 'He knows it is an F' is always false, given the facts about Gs and the only way spotters can find out things about the planes in the sky.¹¹

On this conception it is possible for a perfectly meaningful expression to be appropriately and justifiably applied in certain situations and for its negation to be equally appropriately and justifiably applied in others, even though what is said in each of the positive applications is never true. When those within the situation say 'He does not know it is an F' on some particular occasion they could be said to be relying on a distinction between that occasion and those occasions on which the conditions normally appropriate for asserting 'He knows it is an F' are fulfilled. What is important for the action of saying something is whether it is one sort of occasion or the other. When we, outside the restricted practical context, say that even the careful spotter does not know the plane is an F we are not simply drawing the same contrast. When we deny that knowledge is present we are not distinguishing the spotter's claim to know it is an F from those cases in which the conditions normally appropriate for the assertion of 'He knows it is an F' are fulfilled. We know that it is a case of just that kind; we know he is justified in saying what he does. Our grounds for denying he knows are different. We are distinguishing his position from one in which the conditions of knowledge are fulfilled—conditions of the truth of 'He knows it is an F'. But the fact that we say what we do on different grounds from those available to speakers within the restricted practice does not show that the notion of knowledge we use is different from theirs. When

¹¹ Of course the example will seem unrealistically restricted. Someone could presumably tell a plane was an F if he flew up next to it and read the label on the side, or shot it down and took it apart and checked its serial-number on a list at the factory. It is not impossible to know in that way that a particular plane is an F. But spotters confined to distant sightings in the sky can never know of any plane that it is an F. The point is that that limitation need have no effect on the meaning of what they say or their reasons for saying it.

as restricted in certain ways. Certain possibilities are not even considered, let alone eliminated, certain assumptions are shared and taken for granted and so not examined, and our claims are made and understood as if they were restricted to the particular issues that have explicitly arisen. In the context of the war effort no one has any reason to challenge the careful airplane-spotter's claim to know that the plane he sees is an F. Within the restricted range of possibilities he was trained to consider he has chosen the right one. But we from our detached position can see that his 'knowledge' is confined or restricted to that range. He has been fully competent in doing what he must do, but he does not really know that the plane in the sky is an F.

I have said that we who know the facts about Gs and are in that way detached from the airplane-spotter's context would say that he does not know the plane is an F. Our verdict about his lack of knowledge is not arrived at on the same sorts of grounds as those within the situation might have for saying the very same thing. For those within the spotters' context, there is a contrast between the cases they describe as knowing and those they describe as not knowing. When a plane first appears in the sky, for example, the spotter might say 'I don't know yet what sort of plane it is. It has got x and y , but that is all I can see, so I don't know whether it is an F or not. It might be an E'. Those who are waiting at headquarters to act on his report will have to say that he does not know yet whether it is an F. After he gets a better look and notices that the plane also has z he is no longer in doubt. 'It's an F', he says, 'I know it is. It has also got z , and that rules out the possibility that it is an E'. It can now be reported to headquarters that he knows it is an F, so appropriate action can be taken. There is obviously a real difference between the earlier and the later state of affairs.

A similar contrast exists between the report made by a careful spotter and that of a less conscientious trainee who finds features x and y and simply guesses that z is probably present too, or who concludes without further thought from the presence of x and y that the plane is an F. The careful spotter would be said by his colleagues and superiors to know, the careless spotter not to know, that the plane is

they in their (justified) ignorance think that all the conditions of knowledge are in fact satisfied in the careful spotter's case, they mean by 'He knows it is an F' what we mean by it, but they are simply mistaken (through no fault of theirs). The careful spotter's case does indeed differ in easily discernible ways from the cases in which they say of someone 'He does not know it is an F'. Everyone within that practice can be aware of that difference. What we, in our detached position, realize is that that difference is not the difference between knowing and not knowing.

It is in this way, I think, that the sceptical philosopher would reply to any argument that starts from the premiss that each of a pair of expressions *S* and *not-S* is meaningfully applied on different occasions and reaches the conclusion that both *S* and *not-S* must sometimes apply truly to such occasions. That 'paradigm-case argument' had a brief vogue at the height of linguistic philosophy in the 1950s. Something like it seems to be appealed to in Austin's rhetorical question 'How could we use and contrast the words 'waking' and 'dreaming' as we do if there were not recognized ways of telling on particular occasions that we are not dreaming?'. But the argument fails because it takes no account of how and why the expressions we use come to be applied to the different sorts of occasions to which we apply them. There can be real and easily discernible differences between two sorts of occasions, and we might apply an expression, or its negation, to an occasion on the basis of just such discernible features. But if certain widely-shared but unexamined assumptions are what make it possible or desirable for us to proceed in that way, or if certain restrictions are in force which limit our interest simply to drawing a particular distinction between the two kinds of occasion, then although we will be marking a real difference between the occasion to which we apply *S* and that to which we apply *not-S*, it will not follow that the distinction we draw is in fact the distinction between *S*'s applying truly to a particular occasion and its not so applying. 'He knows it is an F' was appropriately applied to the airplane-spotters in situations differing in clearly recognizable ways from those in which 'He does not know it is an F' was correctly applied. But

the difference drawn between those two sorts of situations within that practice was not the difference between knowing and not knowing. Even in the former cases what the careful spotter said was false; 'He knows it is an F' is never true under the conditions described.

If our own more general practices of gaining and assessing knowledge in everyday life also operated under a similar set of practical constraints or restrictions it looks as if it would also be possible for no one to know anything about the world around us even though our ordinary procedures are followed to the letter and our claims to know things are often beyond criticism. At least this much can be said: no anti-sceptical conclusion to the contrary could be drawn simply from the fact that we use the expressions 'I know . . .', 'He knows . . .', etc., as we do in fact use them. One would then be in a strong position to defend the sceptical conclusion against any objection to the effect that it distorts the meanings of the very words in which it is expressed since it conflicts with obvious facts about how those words are ordinarily used. The evidence from usage would not support that conclusion about meaning on the conception of the relation between meaning and use that I have tried to identify.¹²

¹² Peter Unger rightly insists on the importance of this distinction in his defence of scepticism. He identifies a class of terms he calls 'absolute terms' (like 'flat' and 'empty') which are appropriately applied on many occasions even though they are never literally true of any of the things to which they are applied. That is shown to be no bar to our using and understanding the expressions as we do. For Unger the same holds for 'certain' and, since knowledge implies certainty, for 'know' as well. Our use of such terms is therefore compatible with the literal truth of scepticism. I agree that Unger's conception of the relation between meaning and use can be used to defend scepticism against the charge that it conflicts with the facts of usage. I do not agree that Unger can *establish* scepticism on the basis of his theory of 'absolute terms' alone. I think his argument to show that no one is ever undogmatically certain of (and hence doesn't know) anything about the world makes essential use of a step that is equivalent in force to Descartes's requirement that we must know we are not dreaming if we are to know anything about the world around us. Without that requirement, the 'absolute-ness' of 'certain' and 'know' will not yield the sceptical conclusion. And I have tried to show here that with that requirement we have all we need to generate the sceptical conclusion, so the doctrine of 'absolute terms' is not needed. For a slightly fuller version of this appreciation and criticism of Unger see my review of P. Unger, *Ignorance: a Case for Scepticism* (Oxford, 1975) in *The Journal of Philosophy*, 1977.

Pressing for the precise source of linguistic oddity or inappropriateness and

I have defended the sceptical conclusion against a certain line of attack in order to begin to reveal what I think is the sceptical philosopher's conception of everyday life and everyday assertions lying behind it. He sees those assertions as restricted in certain ways relative to what, with detachment, we can all recognize to be the full conditions of their truth. So now we are led to the question whether that conception is correct, or even fully intelligible. When we begin to understand how it would vindicate scepticism I think we feel it cannot be correct. We see that that conception must somehow be rejected if we are to show how and why the elimination of the dream-possibility is not always required for knowledge of the world around us. That would begin to account for what I feel to be the depth and importance of the sceptical reasoning. Coming to terms with it would eventually involve a great deal more than simply deciding whether somebody knows something in a particular case, or even whether anybody knows anything about the world around us. A whole way of thinking of ourselves and of our practices in everyday life, and perhaps even the possibility of our getting a certain kind of detached understanding of ourselves, would be at issue. On this conception of epistemology there is much more at stake than the question of what knowledge is, or whether and how we know things.

The idea of ourselves and of our relation to the world that lies behind the sceptical reasoning seems to me deeply powerful and not easily abandoned. As long as it is even an intelligible way of thinking the sceptical conclusion will seem to be defensible against attack. In trying to give expression to the idea it is natural to resort to what seems like nothing more than the merest platitudes. If that is so, trying to avoid scepticism by throwing over the old conception will not be easy—it will involve denying what seem to be obvious truths.

distinguishing between meaning and use in the way both Unger and I here rely on form the basis of a fundamental criticism of linguistic philosophy and a quite general theory of language and communication in the important work of H. P. Grice, to which I have been much indebted in my thinking on these and other matters. See his 'The Causal Theory of Perception', *The Aristotelian Society: Supplementary Volume XXXV*, 1961; 'Logic and Conversation' in D. Davidson and G. Harman (eds), *The Logic of Grammar* (Belmont, California, 1975); and the not-yet-published William James Lectures delivered at Harvard University in 1967.

The simplest way to put the idea that lies behind our concern with knowledge is that the world around us that we claim to know about exists and is the way it is quite independently of its being known or believed by us to be that way. It is an objective world. In fact, of course, much of the world we claim to know about was here long before we were and some of it will remain after we have gone. In many cases what we believe or think we know about the world does not require anyone's knowing or believing anything in order for it to be true. If I believe that there is a mountain more than five thousand metres high on the continent of Africa, for example, what I believe will be true, or false, depending solely on the heights of the mountains in Africa. Whether anyone knows or believes or has any special reason to suspect anything about those mountains is not part of what I believe when I believe there is a mountain more than five thousand metres high. If I do not know what to believe and I ask or wonder whether there are any mountains in Africa more than five thousand metres high, my question has an answer which is completely independent of anyone's knowing or believing or being in a position to assert anything. It is quite independent of whether any human or other animate beings have ever existed. What I ask or come to believe concerns only the distance above sea-level of certain mountains.

Of course, I would not have come to believe or even to understand what I now believe unless people had existed and had come to assert things and to know things. But what I believe or understand, viz., that there is a mountain more than five thousand metres high in Africa, does not itself require any such things in order to be true. No statement of precisely *what* I understand, therefore no account of what 'There is a mountain more than five thousand metres high in Africa' means, will include anything about human beings or human knowledge or human thought. In particular it will not include anything about whether that sentence itself is or can be known to be true or could be reasonably asserted in certain circumstances. That would introduce an extraneous reference to human beings or human knowledge into a statement solely about the non-human world.

I am trying to express a conception of the independence of the world, of the idea that the world is there quite independently of human knowledge and belief, that I think we all understand. It embodies a conception of objectivity, of things being a certain way whether anyone is affected by them or interested in them or knows or believes anything about them or not. There seems to be nothing in the conception itself to imply that knowledge or reasonable belief about the objective world is impossible, or that what we can discover or know things about is, or must always be, something different from the objective world so conceived. Quite the contrary. In seeking knowledge we are trying to find out what is true, to ascertain how the world is in this or that respect. Was the suspect in Cleveland that night? Does sitting in draughts contribute to catching cold? Is there a mountain more than five thousand metres high in Africa? What we want to find out in each case is what is true, what the objective facts of the matter are. And what we aspire to and eventually claim to know is the objective truth or falsity of, for example, 'There is a mountain more than five thousand metres high in Africa'. What we aspire to and eventually claim to know is something that holds quite independently of our knowing it or of our being in a position reasonably to assert it. That is the very idea of objectivity.

Many of the things we ask or believe or want to know about do involve human knowledge, human belief and human reasoning. We ask whether anyone knows or has any good reason to believe that sitting in a draught contributes to catching cold, and if so how they know it or what the reasons are. We believe that much more is now known about the properties of matter than was known two hundred years ago. We believe that the causes of cancer are still unknown. What we ask or believe or claim knowledge about in these cases do involve human beings and human knowledge and human thought. They are questions or assertions about what we might call the human world, as opposed to that non-human part of the world that would have been the way it is whether any human beings had existed or behaved in certain ways or not.

Even here, I think, with respect to knowledge and other

human institutions, we have the same conception of objectivity. We want to know whether it is objectively true that somebody knows or has good reason to believe that sitting in a draught contributes to catching cold. In saying that the causes of cancer are still unknown we take ourselves to be making a statement about the present state of human knowledge, and we think human knowledge is in whatever state it is in with respect to the causes of cancer quite independently of our now knowing or being in a position reasonably to assert that it is. Of course we do assert what we do about the present state of human knowledge because we believe we know or have good reason to assert that no one knows the causes of cancer, but we do not regard our being in the position to make that assertion as itself part of what we know or assert when we say that no one knows the causes of cancer. Most facts of human knowledge and belief are in that respect as objective and as independent of anyone's knowing what they are as are the facts about mountains in Africa. If it is an objective fact that the causes of cancer are not known at present then in stating that fact or claiming to know it I am stating or claiming to know an objective fact about human knowledge. And if I try to find out whether anyone knows of any connection between draughts and the common cold I am trying to ascertain what the objective facts about that aspect of human knowledge really are.

Looked at in this way, if I say that I myself know a certain thing, or if I ask or wonder whether I do, what I am saying or asking about will be true or not depending on the present state of my own knowledge with respect to that thing. So when I ask whether I really know that the witness was in Cleveland that night, or when Descartes asks whether he knows that he is sitting by the fire with a piece of paper in his hand, we are enquiring into the present state of our own knowledge with respect to the matter in question. We seek a certain kind of understanding of our state or our relation to the facts—what might be called an objective understanding of our position. Whether someone (even ourselves) knows a certain thing is in that respect as objective a matter of fact as whether there is a mountain of a certain height in Africa, and what we seek is knowledge of whether or not that objective

matter obtains, and perhaps in addition some understanding of how the conditions necessary and sufficient for its obtaining have been fulfilled.

That is just how we understand the position of the airplane-spotters. When the careful spotter says 'That plane is an F' he is saying something only about the identity of that plane in the sky—something that would be true or false whether any spotters were watching or not. And when he says 'I know it is an F' he is stating something about his relation to that objective fact. He used the manual and his observation of the plane to bring him into the position he is now in, but in saying he knows it is an F he is not saying anything about the manual or about his observations. If asked how he knows it is an F, he will say that he saw that it had x , y , z , and he knows that the manual says that any plane with x , y , z is an F. But when he says 'I know it is an F', he is not saying simply 'I know it has x , y , z ' and 'I know that according to the manual it is an F'.

We know that the careful spotter does not know that the plane he sees in the sky is an F. But we can agree that he does know that according to the manual it is an F. So the question of whether he knows what kind of plane it is is not the same as the question whether he knows what the manual says it is. A reflective airplane-spotter in his spare time might be expected to be aware of that distinction, just as we are. Of course, believing what he has been told in his training, he thinks the two questions get the same answer. But in asking himself how he knows what kinds of planes there are in the sky he would see that the manual and his observations are all he has got to go on, and he would admit that if the manual were incorrect in certain ways he would not know everything he now thinks he knows. This thought need not be relevant to the war effort; knowledge or truth are not the only values in time of war. But if he does think he knows that some planes in the sky are Fs, he will appeal to the correctness of the manual to explain that knowledge to himself or others. In the story as told, that assumption is not true—it is because we know that the manual is not correct that we know that he does not know that the plane is an F. But even if the assumption of the correctness of the manual

were true, the reflective spotter would see that its being true is required for his knowing what kinds of planes there are in the sky in the way he does.

Whether the manual is correct or not is itself an objective fact. In this case we outsiders know it is not correct. The spotter who relies on the manual regards it as correct; and he can see that its objective correctness is essential to his knowing. It is because he believes in its correctness that he thinks he knows the plane is an F. We who have a more objective understanding of the spotter's position know that he does not know. We are in a position that he is not in with respect to one of the facts essential to his knowing. We are therefore in a better position for determining whether 'He knows it is an F' is objectively true or not. The reflective spotter thinks it is true, and he thinks he can explain how his knowledge is possible. If we told him what we know about Fs he would realize that he had not been in the best position for determining whether he knows or for explaining how his knowledge is possible. Even without our help, if it occurred to him that the manual might not be correct, he could see that he was not in the best position he could be in for explaining his putative knowledge. He would see that checking the reliability of the manual would put him in a better position for determining whether what he says when he says 'I know it is an F' is objectively true. It would give him a more objective understanding of his position.

The sceptical philosopher's conception of our own position and of his quest for an understanding of it is parallel to this reflective airplane-spotter's conception. It is a quest for an objective or detached understanding and explanation of the position we are objectively in. What is seen to be true from a detached 'external' standpoint might not correspond to what we take to be the truth about our position when we consider it 'internally', from within the practical contexts which give our words their social point. Philosophical scepticism says the two do not correspond; we never know anything about the world around us, although we say or imply that we do hundreds of times a day.

I think we do have a conception of things being a certain way quite independently of their being known or believed

82 *Philosophical Scepticism and Everyday Life*

or said to be that way by anyone. I think the source of the philosophical problem of the external world lies somewhere within just such a conception of an objective world or in our desire, expressed in terms of that conception, to gain a certain kind of understanding of our relation to the world. But in trying to describe that conception I think I have relied on nothing but platitudes we would all accept—not about specific ways we all now believe the world to be, but just the general idea of what an objective world or an objective state of affairs would be. If those platitudes about objectivity do indeed express the conception of the world and our relation to it that the sceptical philosopher relies on, and if I am right in thinking that scepticism can be avoided only if that conception is rejected, it will seem that in order to avoid scepticism we must deny platitudes we all accept. I believe this sometimes has happened in philosophy.

But perhaps the commonplaces I have appealed to, if they really are uncontroversial, do not manage to express the full conception of objectivity and of everyday life that the philosopher relies on in his sceptical reasoning. Perhaps there is a way of taking them so that they express no philosophical conception at all, and so do not generate or exacerbate the philosophical problem of the external world. I want now to explore several different ways in which that might be thought to be true. They are in effect different ways of trying to explain what the philosophical problem of the external world amounts to.

07/20/09 11:52 AM