

*seats* 3<sup>00</sup>

# *The Problem of Knowledge*

---

A. J. AYER

1956

PENGUIN BOOKS

## SCEPTICISM AND CERTAINTY

(i) *Philosophical scepticism*

I HAVE said that what the philosophical sceptic calls in question is not the way in which we apply our standards of proof, but these standards themselves. But not all questioning of accepted canons of evidence is philosophical. There was a time when people believed that examining the entrails of birds was a way of discovering whether a certain course of action would be propitious, whether, for example, the occasion was favourable for joining battle. Then any sceptic who doubted the value of such a method of divination would have been questioning an accepted canon of evidence. And it is now agreed that he would in fact have been right. But the justification for his doubt would have been not philosophical, but scientific. It might have been the case that these so-called omens were systematically connected with the events which they were supposed to presage: but experience shows otherwise. In the same way, a medieval doubter might have raised the question whether the failure to survive a trial by ordeal was a trustworthy indication of guilt. He, too, would have been challenging a recognized method of proof; and his scepticism would have been justified. But, again, it would have been justified on scientific grounds. It is a matter of empirical fact that the innocent, no less than the guilty, are susceptible to physical injury and death.

The peculiarity of the philosopher's doubts is that they are not in this way connected with experience. Experience does indeed show that such reputed sources of knowledge as memory or perception or testimony are fallible. But the philosophical sceptic is not concerned, as a scientist would be, with distinguishing the conditions in which these sources are likely

to fail from those in which they can normally be trusted. Whereas the enlightened thinker who casts doubt upon the reliability of omens is suggesting that they do not yield good enough results, that this method of prognostication does not reach a standard which other methods could, and perhaps do, satisfy, the philosophical sceptic makes no such distinction: his contention is that any inference from past to future is illegitimate. Similarly, he will maintain not merely that there are circumstances in which a man's senses are liable to deceive him, as when he is suffering from some physiological disorder, but rather that it is to be doubted whether the exercise of sense-perception can in any circumstances whatever afford proof of the existence of physical objects. He will argue not merely that memory is not always to be trusted, but that there is no warrant for supposing that it ever is: the doubt which he raises is whether we can ever be justified in inferring from present experiences to past events. In questioning one's right to believe in the experiences of others he will not be content with producing empirical evidence to show how easily one may be mistaken; so far from encouraging us to be more circumspect, his argument is designed to show that however circumspect we are it makes no difference: it puts the thoughts and feelings of others behind a barrier which it is impossible that one should ever penetrate.

The fact that this type of scepticism is so undiscriminating in its scope, that it rains alike on the just and the unjust, has been thought to expose it to an easy refutation. Just as, to use a simile of Ryle's, 'there can be false coins only where there are coins made of the proper materials by the proper authorities',<sup>1</sup> so, it is argued, there can be times when our senses deceive us only if there are times when they do not. A perception is called illusory by contrast with other perceptions which are veridical: therefore to maintain that all perceptions must be illusory would be to deprive the word 'illusory' of

1. G. Ryle, *Dilemmas* (Cambridge, 1954), p. 95.

its meaning. This rejoinder would not, indeed, be fatal to a more moderate sceptic who held, not that all perceptions are bound to be illusory, but only that we can never really know that any are not; but he too is exposed to a similar objection. For how, it may be asked, could we ever discover that any appearances were deceptive unless we knew that some were trustworthy? From a distance, or in a dim light, I may mistake the shape or colour of the thing that I am looking at; I may confuse one object with another; in exceptional conditions, I may even think that I am perceiving something when there is nothing there at all: but I should not know that I made these errors unless I were in a position to correct them. From close at hand and in a stronger light I can see what the colour and shape of the thing really are, and knowing this I am enabled to infer that I saw them wrongly before. I learn that I have had a hallucination because the further course of my experience assures me that the object which I thought I saw does not exist. In the same way, the only reason that I have for thinking that I suffer from errors of memory is that what I seem to remember sometimes runs counter to other historical evidence which I am entitled to accept: my only reason for supposing that I am wrong about the experiences of others is that I make judgements about them which are inconsistent with what I subsequently discover to be right.

This argument is not decisive. It is true that no judgements of perception would be specially open to distrust unless some were trustworthy; but this is not a proof that we cannot be mistaken in trusting those that we do. Even granting that it makes no sense to say that all our perceptions are delusive, any one of them still may be. We have to make good our claim to know that some particular ones are not. And the same applies to the other types of judgement which the sceptic impugns. From the fact that our rejection of some of them is grounded on our acceptance of others it does not follow that those that we accept are true.

Nevertheless the argument does show that these general forms of scepticism can find no justification in experience. A historian who is distrustful of one of his authorities may have his suspicions confirmed by finding that the reports which this authority gives conflict with the evidence that is available from other sources; if these sources are numerous and independent, and if they agree with one another, he will be reasonably confident that their account of the matter is correct. But if his doubts embraced every statement which referred to the past, there would be no such way of confirming them; for all the relevant evidence would be equally suspect. In the same way, a scientist who is sceptical of the truth of some particular hypothesis may justify himself by showing that it is at variance with some well-established theory. But for someone who maintains that all inductive reasoning is illegitimate there are no well-established theories; there are theories which have not as yet been confuted, but they are not considered any more worthy of credence than those that have; nor, on this view, does the fact that a theory has been falsified make it any the less likely to hold good in future cases. It is, indeed, a matter of experience that general hypotheses do meet with counter-instances; and it might therefore seem that the view that all inductive reasoning is illegitimate had some empirical support. But this conclusion would be mistaken; or rather, it would misrepresent the sceptic's standpoint. His thesis is not that every theory, or hypothesis, will eventually break down, but that the accumulation of favourable instances, however long continued, affords us no good reason for believing it. And clearly the validity of this contention is independent of the actual course of our experience.

If experience cannot justify the sceptic, neither can it refute him. Psychologically, indeed, he may receive encouragement from the fact that by following our accepted standards of proof we sometimes arrive at beliefs which turn out to be

false: it would be hard for him to get a hearing if the procedures which he questions never led us astray. But it is not essential to his position that this be so. All that he requires is that errors should be possible, not that they should actually occur. For his charge against our standards of proof is not that they work badly; he does not suggest that there are others which would work better. The ground on which he attacks them is that they are logically defective; or if not defective, at any rate logically questionable.

When we claim the right to be sure of the truth of any given statement, the basis of the claim may be either that the statement is self-evident, or that its truth is directly warranted by our experience, or that it is validly derivable from some other statement, or set of statements, of which we have the right to be sure. Accordingly, if such claims are to be challenged, it may be argued either that the statements which we take as requiring no further proof, beyond an appeal to intuition or experience, are themselves not secure, or that the methods of derivation which we regard as valid may not really be so. These lines of argument do not exclude each other, and both have been pursued. It has been queried whether we can ever be in a position to say of any statement that there is no doubt about its truth; and this query extends to the validity even of deductive reasoning: for if nothing is certain, then it is not certain that one statement follows from another. But our justification for deriving statements from one another is put in question chiefly in the cases where the transition is not deductive, or at least not obviously so. There is, or has been thought to be, a general problem of induction which concerns the validity of all types of factual inference: but, as we have noted, there are also special problems concerning our right to pass from one sort of statement to another; they raise such questions as whether, or how, we are justified in making assertions about physical objects on the basis of our sense-experiences, or in attributing experiences to others on the

evidence of their behaviour, or in regarding our memories as giving us knowledge of the past. It is by forcing us to consider questions of this sort that the sceptic performs his main service to philosophy. But before attempting to examine them it may be well for us first to discuss the problem of certainty; the question whether there are any statements whose truth can be established beyond the possibility of doubt.

### (ii) *The quest for certainty*

The quest for certainty has played a considerable part in the history of philosophy: it has been assumed that without a basis of certainty all our claims to knowledge must be suspect. Unless some things are certain, it is held, nothing can be even probable. Unfortunately it has not been made clear exactly what is being sought. Sometimes the word 'certain' is used as a synonym for 'necessary' or for '*a priori*'. It is said, for example, that no empirical statements are certain, and what is meant by this is that they are not necessary in the way that *a priori* statements are, that they can all be denied without self-contradiction. Accordingly, some philosophers take *a priori* statements as their ideal. They wish, like Leibniz, to put all true statements on a level with those of formal logic or pure mathematics; or, like the existentialists, they attach a tragic significance to the fact that this cannot be done. But it is perverse to see tragedy in what could not conceivably be otherwise; and the fact that all empirical statements are contingent, that even when true they can be denied without self-contradiction, is itself a matter of necessity. If empirical statements had the formal validity which makes the truths of logic unassailable they could not do the work that we expect of them; they would not be descriptive of anything that happens. In demanding for empirical statements the safeguard of logical necessity, these philosophers have failed to see that they would thereby rob them of their factual content.

Neither is this the only way in which their ideal of *a priori* statements fails them. Such statements are, indeed, unassailable, in the sense that, if they are true, there are no circumstances in which they could have been false. One may conceive of a world in which they had no useful application, but their being useless would not render them invalid: even if the physical processes of addition or subtraction could for some reason not be carried out, the laws of arithmetic would still hold good. But from the fact that *a priori* statements, if they are true, are unassailable in this sense, it does not follow that they are immune from doubt. For, as we have already remarked, it is possible to make mistakes in mathematics or in logic. It is possible to believe an *a priori* statement to be true when it is not. And we have seen that it is vain to look for an infallible state of intuition, which would provide a logical guarantee that no mistake was being made. Here too, it may be objected that the only reason that we have for concluding that any given *a priori* statement is false is that it contradicts some other which is true. That we can discover our errors shows that we have the power to correct them. The fact that we sometimes find ourselves to be mistaken in accepting an *a priori* statement, so far from lending favour to the suggestion that all those that we accept are false, is incompatible with it. But this still leaves it open for us to be at fault in any particular case. There is no special set of *a priori* statements of which it can be said that just these are beyond the reach of doubt. In very many instances the doubt would not, indeed, be serious. If the validity of some logical principle is put in question, one may be able to find a way of proving or disproving it. If it be suggested that the proof itself is suspect, one may obtain reassurance by going over it again. When one has gone over it again and satisfied oneself that there is nothing wrong with it, then to insist that it may still not be valid, that the conclusion may not really have been proved, is merely to pay lip-service to human fallibility. The doubt is

maintained indefinitely, because nothing is going to count as its being resolved. And just for this reason it is not serious. But to say that it is not serious is not logically to exclude it. There can be doubt so long as there is the possibility of error. And there must be the possibility of error with respect to any statement, whether empirical or *a priori*, which is such that from the fact that someone takes it to be so it does not follow logically that it is so. We have established this point in our discussion of knowledge, and we have seen that it is not vitiated by the fact that in the case of *a priori* statements there may be no other ground for accepting them than that one sees them to be true.

Philosophers have looked to *a priori* statements for security because they have assumed that inasmuch as these statements may themselves be certain, in the sense of being necessary, they can be certainly known. As we have seen, it may even be maintained that only what is certainly true can be certainly known. But this, it must again be remarked, is a confusion. *A priori* statements can, indeed, be known, not because they are necessary but because they are true and because we may be entitled to feel no doubt about their truth. And the reason why we are entitled to feel no doubt about their truth may be that we can prove them, or even just that we can see them to be valid; in either case there is an appeal to intuition, since we have at some point to claim to be able to see the validity of a proof. If the validity of every proof had to be proved in its turn, we should fall into an infinite regress. But to allow that there are times when we may justifiably claim the right to be sure of the truth of an *a priori* statement is not to allow that our intuitions are infallible. One is conceded the right to be sure when one is judged to have taken every reasonable step towards making sure: but this is still logically consistent with one's being in error. The discovery of the error refutes the claim to knowledge; but it does not prove that the claim was not, in the circumstances, legitimately made. The claim

to know an *a priori* statement is satisfied only if the statement is true; but it is legitimate if it has the appropriate backing, which may, in certain cases, consist in nothing more than the statement's appearing to be self-evident. Even so, it may fail: but if such claims were legitimate only when there was no logical possibility of error, they could not properly be made at all.

Thus, if the quest for certainty is simply a quest for knowledge, if saying that a statement is known for certain amounts to no more than saying that it is known, it may find its object in *a priori* statements, though not indeed in them uniquely. If, on the other hand, it is a search for conditions which exclude not merely the fact, but even the possibility, of error, then knowledge of *a priori* statements does not satisfy it. In neither case is the fact that these *a priori* statements may themselves be certain, in the sense of being necessary, relevant to the issue. Or rather, as we have seen, it is relevant only if we arbitrarily decide to make it so.

### (iii) '*I think, therefore I am*'

The attempt to put knowledge on a foundation which would be impregnable to doubt is historically associated with the philosophy of Descartes. But Descartes, though he regarded mathematics as the paradigm of knowledge, was aware that its *a priori* truths are not indubitable, in the sense that he required. He allowed it to be possible that a malignant demon should deceive him even with respect to those matters of which he was the most certain.<sup>1</sup> The demon would so work upon his reason that he took false statements to be self-evidently true. The hypothesis of there being such an arch-deceiver is indeed empty, since his operations could never be detected: but it may be regarded as a picturesque way of

1. René Descartes, *Meditations on the First Philosophy*, Meditation I.

expressing the fact that intuitive conviction is not a logical guarantee of truth. The question which Descartes then raises is whether, of all the propositions which we think we know, there can be any that escape the demon's reach.

His answer is that there is one such proposition: the famous *cogito ergo sum*: I think, therefore I am.<sup>1</sup> The demon might perhaps have the power to make me doubt whether I was thinking, though it is difficult to see what this would come to; it is not clear what such a state of doubt would be. But even allowing that the expression 'I am doubting whether I am thinking' describes a possible situation, the doubt must be unwarranted. However much he can shake my confidence, the demon cannot deceive me into believing that I am thinking when I am not. For if I believe that I am thinking, then I must believe truly, since my believing that I am thinking is itself a process of thought. Consequently, if I am thinking, it is indubitable that I am thinking, and if it is indubitable that I am thinking, then, Descartes argues, it is indubitable that I exist, at least during such times as I think.

Let us consider what this argument proves. In what sense is the proposition that I think, and consequently that I exist, shown to be indubitable? It is not a question for psychology. The suggestion is not that it is physically impossible to doubt that one is thinking, but rather that it somehow involves a logical impossibility. Yet while there may be some question about the meaning that one should attach to the statement that I doubt whether I am thinking, it has not been shown to be self-contradictory. Nor is the statement that I am thinking itself the expression of a necessary truth. If it seems to be necessary, it is because of the absurdity of denying it. To say 'I am not thinking' is self-stultifying since if it is said intelligently it must be false: but it is not self-contradictory. The proof that it is not self-contradictory is that it might have been true. I am now thinking but I might easily not have

1. *Vide Meditation II and Discourse on Method*, part IV.

been. And the same applies to the statement that I exist. It would be absurd for me to deny that I existed. If I say that I do not exist, it must be false. But it might not have been false. It is a fact that I exist, but not a necessary fact.

Thus neither 'I think' nor 'I exist' is a truth of logic: the logical truth is only that I exist if I think. And we have seen that even if they were truths of logic they would not for that reason be indubitable. What makes them indubitable is their satisfying a condition which Descartes himself does not make explicit, though his argument turns upon it. It is that their truth follows from their being doubted by the person who expresses them. The sense in which I cannot doubt the statement that I think is just that my doubting it entails its truth: and in the same sense I cannot doubt that I exist. There was therefore no need for Descartes to derive '*sum*' from '*cogito*'; for its certainty could be independently established by the same criterion.

But this certainty does not come to very much. If I start with the fact that I am doubting, I can validly draw the conclusion that I think and that I exist. That is to say, if there is such a person as myself, then there is such a person as myself, and if I think, I think. Neither does this apply only to me. It is obviously true of anyone at all that if he exists he exists and that if he thinks he thinks. What Descartes thought that he had shown was that the statements that he was conscious, and that he existed, were somehow privileged, that, for him at least, they were evidently true in a way which distinguished them from any other statements of fact. But this by no means follows from his argument. His argument does not prove that he, or anyone, knows anything. It simply makes the logical point that one sort of statement follows from another. It is of interest only as drawing attention to the fact that there are sentences which are used in such a way that if the person who employs them ever raises the question whether the statements which they express are true, the

answer must be yes. But this does not show that these statements are in any way sacrosanct, considered in themselves.

Yet surely I can be certain that I am conscious, and that I exist. Surely my evidence for this could not be stronger than it is. But again it is not clear what is being claimed when it is said that these things are certain or that one can be certain of them. Perhaps only that I know that they are so, and of course I do. But these are not the only facts that I know, nor, as it sometimes appears to be suggested, is my knowing them a condition of my knowing anything else. It is conceivable that I should not have been self-conscious, which is to say that I should not know that I existed; but it would not follow that I could not know many other statements to be true. In theory, I could know any of the innumerable facts which are logically independent of the fact of my existing. I should indeed know them without knowing that I knew them, though not necessarily without knowing that they were known: my whole conception of knowledge would be impersonal. Perhaps this is a strange supposition, but it is not self-contradictory.

But while in the case of other facts which I may reasonably claim to know, it is at least conceivable that the evidence which I have for them should be even stronger than it is, surely the fact that I exist and the fact that I am conscious stand out for the reason that in their case the evidence is perfect. How could I possibly have better evidence than I do for believing that I am conscious, let alone for believing that I exist? This question is indeed hard to answer, but mainly because it seems improper in these cases to speak of evidence at all. If someone were to ask me How do you know that you are conscious? What evidence have you that you exist? I should not know how to answer him: I should not know what sort of answer was expected. The question would appear to be a joke, a parody of philosophical cautiousness. If it were seriously pressed, I might become indignant: What do you

mean, how do I know that I exist? I am here, am I not, talking to you? If a 'philosophical' answer were insisted on, it might be said that I proved that I existed and that I was conscious by appealing to my experience. But not then to any particular experience. Any feeling or perception that I cared to instance would do equally well. When Hume looked for an impression of his self, he failed to find one: he always stumbled instead upon some particular perception.<sup>1</sup> He allowed that others might be luckier, but in this he was ironical. For the point is not that to have an experience of one's self is to perform a remarkably difficult feat of introspection: it is that there is nothing that would count as having an experience of one's self, that the expression 'having an experience of one's self' is one for which there is no use. This is not to say that people are not self-conscious, in the sense that they conceive of things as happening to themselves. It is that the consciousness of one's self is not one experience among others, not even, as some have thought, a special experience which accompanies all the others. And this is not a matter of psychology but of logic. It is a question of what self-consciousness is understood to mean.

If there is no distinctive experience of finding out that one is conscious, or that one exists, there is no experience at all of finding out that one is not conscious, or that one does not exist. And for this reason it is tempting to say that sentences like 'I exist', 'I am conscious', 'I know that I exist', 'I know that I am conscious' do not express genuine propositions. That Mr A exists, or that Mr A is conscious, is a genuine proposition; but it may be argued that it is not what is expressed by 'I exist' or 'I am conscious', even when I am Mr A. For although it be true that I am Mr A, it is not necessarily true. The word 'I' is not synonymous with 'Mr A' even when it is used by Mr A to refer to himself. That he is Mr A, or that

<sup>1</sup>. David Hume, *A Treatise of Human Nature*, Book 1, part iv, section vi.

he is identifiable in any other manner, is an empirical statement which may be informative not only to others, but also in certain circumstances to Mr A himself, for instance if he has lost his memory. It cannot therefore be reasoned that because one may succeed in expressing genuine propositions by replacing the 'I' in such sentences as 'I am conscious' or 'I exist' by a noun, or descriptive phrase, which denotes the person concerned, these sentences still have a factual meaning when this replacement is not made.

All the same it is not difficult to imagine circumstances in which they would have a use. 'I am conscious' might be said informatively by someone recovering from a swoon. If I had been presumed to be dead there might be a point in my proclaiming that I still existed. On recovering consciousness after some accident or illness, I might make this remark even to myself, and make it with a sense of discovery. Just as there are moments between sleep and waking when one may seriously ask oneself if one is awake, so there are states of semi-consciousness in which saying 'I exist' answers a genuine question. But what information does this answer give? If I have occasion to tell others that I exist, the information which they receive is that there exists a man answering to some description, whatever description it may be that they identify me by; it would not be the same in every case. But when I tell myself that I exist, I do not identify myself by any description: I do not identify myself at all. The information which I convey to myself is not that there exists a person of such and such a sort, information which might be false if I were mistaken about my own identity or character. Yet I am in fact a person of such and such a sort. There is nothing more to me than what can be discovered by listing the totality of the descriptions which I satisfy. This is merely an expression of the tautology that if a description is complete there is nothing left to be described. But can it not be asked what it is that one is describing? The answer is that this question makes sense only

as a request for further description: it implies that the description so far given is incomplete, as in fact it always will be. But then if, in saying that I exist, I am not saying anything about a description's being satisfied, what can I be saying? Again it is tempting to answer that I am saying nothing.

Yet this would not be correct. Even when it is not doing duty for a description, nor coupled with one, the demonstrative 'I' may have a use. In the case which we envisaged, the case of a return to consciousness, it signals the presence of some experience or other. It does not, however, characterize this experience in any way. It merely points to the existence of whatever it is, in the given circumstances, that makes its own use possible. And since it is a contingent fact that any such situation does exist, the assertion which simply serves to mark it may be held to be informative. The sentence 'I exist', in this usage, may be allowed to express a statement which like other statements is capable of being either true or false. It differs, however, from most other statements in that if it is false it can not actually be made. Consequently, no one who uses these words intelligently and correctly can use them to make a statement which he knows to be false. If he succeeds in making the statement, it must be true.

It is, therefore, a peculiar statement; and not only peculiar but degenerate. It is degenerate in the way that the statements which are expressed by such sentences as 'this exists' or 'this is occurring now' are degenerate. In all these cases the verbs which must be added to the demonstratives to make a grammatical sentence are sleeping partners. The work is all done by the demonstrative: that the situation, to which it points, exists, or is occurring, is a condition of the demonstrative's use. It is for this reason that any statement of this sort which is actually expressed must be true. It is not necessarily true, since the situation to which the demonstrative points might not have existed; it is logically possible that the condition for

this particular use of the demonstrative should not have obtained. It is, however, like an analytic statement in that, once we understand the use of the demonstrative, here functioning as subject, the addition of the predicate tells us nothing further. Divorced from its context the whole statement has no meaning. Taken in context it is informative just as drawing attention to whatever it may be that the demonstrative is used to indicate. It approximates, therefore, to a gesture or to an ejaculation. To say 'I exist' or 'this is occurring now' is like saying 'look!' or pointing without words. The difference is that, in the formulation of the indicative sentence, the existential claim is made explicit; and it is because of this that the sentence may be said to express a statement, whereas the ejaculation or the gesture would not: one does not speak of ejaculations or gestures as being true or false. But there is no difference in the information conveyed.

Thus we see that the certainty of one's own existence is not, as some philosophers have supposed, the outcome of some primary intuition, an intuition which would have the distinctive property of guaranteeing the truth of the statement on which it was directed. It is indeed the case that if anyone claims to know that he exists, or that he is conscious, he is bound to be right. But this is not because he is then in some special state of mind which bestows this infallibility upon him. It is simply a consequence of the purely logical fact that if he is in any state whatever it follows that he exists; if he is in any conscious state, whatever it follows that he is conscious. He might exist without knowing it; he might even be conscious without knowing it, as is presumably the case with certain animals: there is at any rate no contradiction in supposing them to be conscious without supposing them to be conscious of themselves. But, as we have seen, if anyone does claim to know that he exists or that he is conscious, his claim must be valid, simply because its being valid is a condition of its being made. This is not to say, however, that he, or

anyone, knows any description of himself, or his state of consciousness, to be true. To know that one exists is not, in this sense, to know anything about oneself any more than knowing that *this* exists is knowing anything about *this*. Knowing that I exist, knowing that this is here, is having the answer to a question which is put in such a form that it answers itself. The answer is meaningful only in its context, and in its context the condition of its being meaningful is its being true. This is the ground for saying that statements like 'I exist' are certain, but it is also the proof of their degeneracy: they have nothing to say beyond what is implied in the fact that they have a reference.

(iv) *Are any statements immune from doubt?*

If our aim is never to succumb to falsehood, it would be prudent for us to abstain from using language altogether. Our behaviour might still be hesitant or misguided but it is only with the use of language that truth and error, certainty and uncertainty, come fully upon the scene. It is only such things as statements or propositions, or beliefs or opinions, which are expressible in language, that are capable of being true or false, certain or doubtful. Our experiences themselves are neither certain nor uncertain; they simply occur. It is when we attempt to report them, to record or forecast them, to devise theories to explain them, that we admit the possibility of falling into error, or for that matter of achieving truth. For the two go together: security is sterile. It is recorded of the Greek philosopher Cratylus that, having resolved never to make a statement of whose truth he could not be certain, he was in the end reduced simply to wagging his finger. An echo of his point of view is to be found in the disposition of some modern philosophers to regard the expression of purely demonstrative statements like 'this here now' as the ideal limit to which all narrative uses of language should approach. It is a matter

in either case of gesticulating towards the facts without describing them. But it is just their failure to describe that makes these gestures defective as a form of language. Philosophers have been attracted by the idea of a purely demonstrative use of words because they have wanted to make the best of both worlds. They have sought as it were to merge their language with the facts it was supposed to picture; to treat its signs as symbols, and yet bestow upon them the solidity which belongs to the facts themselves, the facts being simply there without any question of doubt or error arising. But these aims are incompatible. Purely demonstrative expressions are in their way secure; but only because the information which they give is vanishingly small. They point to something that is going on, but they do not tell us what it is.

Some philosophers, however, have thought that they could go further than this. They have thought it possible to find a class of statements which would be both genuinely informative and at the same time logically immune from doubt. The statements usually chosen for this rôle contain a demonstrative component, but they are not wholly demonstrative; they contain also a descriptive component which is supposed to characterize some present state of the speaker, or some present content of his experience. The sort of example that we are offered is 'I feel a headache' or 'this looks to me to be red' or 'this is louder than that', where 'this' and 'that' refer to sounds that I am actually hearing or, more ambitiously, 'it seems to me that this is a table' or 'I seem to remember that such and such an event occurred'. Such statements may be false as well as true: nor is their truth a condition of their being made. I may, for example, be lying when I say that I feel a headache. But while I may be lying and so deceive others, I cannot, so it is maintained, myself be in any doubt or in any way mistaken about the fact. I cannot be unsure whether I feel a headache, nor can I think that I feel a headache when I do not. And the same applies to the other examples. In all

cases, so it is alleged, if one misdescribes the nature of one's present experience, one must be doing so deliberately. One must be saying something which one knows for certain to be false.

Since the only way in which any statement of fact can be discovered either to be true or false is by someone's having some experience, these statements which are supposed, as it were, to photograph the details of our experiences seem to occupy a privileged position: for it would appear that it is their truth or falsehood that provides the test for the validity of all the others. For this reason they have sometimes been described as basic statements, or basic propositions. Or rather, it has been assumed that there must be some statements the recognition of whose truth or falsehood supplies the natural terminus to any process of empirical verification; and statements which are descriptive of the present contents of experiences are selected as the most worthy candidates. The reason why they are so distinguished is that it is thought that they alone are directly and conclusively verifiable; of all statements which have a descriptive content they alone are not subject to any further tests. If they were subject to further tests the process of verification would not terminate with them. But where else, then, could it terminate? So these experiential statements, as we may call them, are taken as basic because they are held to be 'incorrigible'.

To say that these statements are incorrigible is not, however, to say that one's assessment of their truth or falsehood can not ever be revised. Or if it does imply this, it is an error. Suppose that, feeling a headache, I write down in my diary the sentence 'I feel a headache'. To-morrow when I read this entry I may seem to remember that I did not make it seriously; and so I may decide that the statement which it expressed was false. In the circumstances envisaged this decision would be wrong; but this does not mean that I am not free to make it, or to revise it in its turn. But, it may be said, the

statement which you subsequently reject is not the same as the one you originally accepted. The statement which is expressed by the sentence 'I feel a headache now' is different from the statement which is expressed by the sentence 'I felt a headache then' even though the pronoun refers to the same person in each case and 'now' and 'then' refer to the same moment. Now there is indeed a sense in which these sentences do have different meanings; the correct translation of one of them into a different language would not be a correct translation of the other. Granted that their reference is the same, the difference in their form shows that they are uttered at different times. But I think it would be wrong to conclude that they expressed different statements; for the state of affairs which makes what is expressed by either of them true is one and the same. Moreover, it seems strange to say that when I verify a prediction about the course of my experience, the statement which I actually verify is different from the statement which embodies the prediction, since one is expressed by a sentence in the present and the other by a sentence in the future tense. Yet this would follow from the assumption that if two sentences differ in this way the statements which they express cannot be the same. I think, therefore, that this assumption is to be rejected, and consequently that experiential statements are not incorrigible in the sense that once they have been discovered to be true they cannot subsequently be denied. Clearly, if we have discovered them to be true, we shall be in error if we subsequently deny them: all that I am now maintaining is that it is an error which it is within our power to make.

But in what sense then is it at all plausible to claim that these statements are incorrigible? Only, I think, in the sense that one's grounds for accepting them may be perfect. It is, therefore, misleading to talk of a class of incorrigible, or indubitable, statements as though 'being incorrigible' or 'being indubitable' were properties which belonged to statements in themselves. The suggestion is rather that there is a

class of statements which in certain conditions only cannot be doubted; statements which are known incorrigibly when they are made by the right person in the right circumstances and at the right time. Thus, in my view at least, the sentences 'he has a headache', when used by someone else to refer to me, 'I shall have a headache', used by me in the past with reference to this moment, and 'I have a headache' all express the same statement; but the third of these sentences alone is used in such conditions as make it reasonable for me to claim that the statement is incorrigibly known. What is 'incorrigible' in this case is the strength of the basis on which I put the statement forward: not in the sense that the existence of such a basis cannot be denied or doubted by other persons, or by myself at other times, but that given its existence—and it is fundamental to the argument that I *am* given it—then, independently of all other evidence, the truth of the statement is perfectly assured. It is in this sense only that the statement may be regarded as not being subject to any further tests: a claim which may seem more modest when it is remarked that even if I am given a conclusive basis for accepting the truth of what I say in such conditions, the gift is immediately withdrawn. The conditions change; the experience is past; and I am left free to doubt or deny that I ever had it, and so again to put in question the truth of the statement which for a moment I 'incorrigibly' knew.

The ground, then, for maintaining that, while one is having an experience, one can know with absolute certainty the truth of a statement which does no more than describe the character of the experience in question is that there is no room here for anything short of knowledge: there is nothing for one to be uncertain or mistaken about. The vast majority of the statements which we ordinarily make assert more than is strictly contained in the experiences on which they are based: they would indeed be of little interest if they did not. For example, I am now seated in a vineyard: and I can fairly claim

to know that there are clusters of grapes a few feet away from me. But in making even such a simple statement as 'that is a bunch of grapes', a statement so obvious that in ordinary conversation, as opposed, say, to an English lesson, it would never be made, I am in a manner going beyond my evidence. I can see the grapes: but it is requisite also that in the appropriate conditions I should be able to touch them. They are not real grapes if they are not tangible; and from the fact that I am having just these visual experiences, it would seem that nothing logically follows about what I can or cannot touch. Neither is it enough that I can see and touch the grapes: other people must be able to perceive them too. If I had reason to believe that no one else could, in the appropriate conditions, see or touch them, I should be justified in concluding that I was undergoing a hallucination. Thus, while my basis for making this assertion may be very strong, so strong indeed as to warrant a claim to knowledge, it is not conclusive; my experience, according to this argument, could still be what it is even though the grapes which I think that I am perceiving really do not exist. But suppose now that I make an even less ambitious statement: suppose that I assert merely that I am seeing what now looks to me to be a bunch of grapes, without the implication that there is anything really there at all; so that my statement would remain true even if I were dreaming or suffering a complete hallucination. How in that case could I possibly be wrong? What other people may experience, or what I myself may experience at other times, does not affect the issue. My statement is concerned only with what appears to me at this moment, and to me alone: whether others have the same impression is irrelevant. I may indeed be using words eccentrically. It may be that it is not correct in English to describe what I seem to be seeing as a bunch of grapes. But this, so it is argued, does not matter. Even if my use of words be unconventional, what I mean to express by them must be true.

(v) *Public and private uses of language*

But this implies not only that the experience which I am describing is private, in the sense that it is mine and not anybody else's, but also that I am giving a private description of it. No doubt the words in which I express my statement are drawn from common speech. No doubt it can be understood by others as well as by myself: we have even allowed that it could be made by others, though they would not, like me, be qualified to make it incorrigibly. But if, provided that I am not lying, my statement must be true however I express it, then even though I am using words which belong to a public language, and using them correctly, there is a sense in which my use of them is private. It is private inasmuch as the meaning of my words is supposed to be fixed entirely by the character of the experience I am using them to indicate, independently of any public standard of usage. This point may not have been made clear in our examples, just because they have been chosen so as to be publicly intelligible. For if I say that I am now seeing what looks to me to be a bunch of grapes, the expression 'looks to me to be a bunch of grapes' may well be understood to mean 'looks to me as a bunch of grapes normally does look', not only to me but to any normal observer; and in that case the question how it normally looks is relevant to the truth of what I am saying. If I were mistaken, as I might be, in supposing that the standard appearance of a bunch of grapes was anything like this, my statement would be false. But the assumption is that my statement remains true even though what I describe as looking like a bunch of grapes does not by conventional standards merit this description. And this means that I am using the expression 'what looks to me to be a bunch of grapes' simply to refer to the content of this experience, whatever it may be. This is not indeed how I normally should use this expression, but it is the way in which I am required to use it if my statement is to be incorrigible. In fact it is an

expression which has a conventional use, but in so far as it serves merely to characterize this momentary, private experience, any other expression which I had chosen to invent for the purpose would have done just as well. Its business being merely to record an episode in my private history, no one else can be in a position to say that my use of it is incorrect.

At this point, however, some philosophers would object that this is not a possible use of language.<sup>1</sup> Whether or not the signs which I employ to record the ways things look to me have a conventional use, they must, if they are to function as descriptive symbols, be endowed with meaning: and they cannot be endowed with meaning unless they are used in accordance with a rule. But rules are public. There are objective tests for deciding whether they are being kept or broken. I can be right or wrong in saying that this looks to me like a bunch of grapes because I have ways of finding out how bunches of grapes are supposed to look: there is a public standard to which I can appeal. But if I do no more than affix an arbitrary label to some experience that I am having, I have no way of testing whether the label is correctly attached or not. There will, indeed, be no meaning in saying that its attachment is either correct or incorrect; and in that case it only masquerades as a label. It is not a symbol of anything at all. I am not bound to employ signs which are familiar to others: I can devise and use a private code. But though the materials of my language may be private, in the sense that only I employ them, its use cannot be: if it is to be a genuine language, it must function in the way that a public language does. It must be teachable to others whether or not it is ever actually taught: there must be means available to them as well as to me of deciding whether I observe its rules. But these conditions would not be met if my words served merely to label my experiences.

1. *Vide* my symposium with R. Rhees, 'Can there be a Private Language?', *Supplementary Proceedings of the Aristotelian Society*, vol. xxviii.

I do not think that this objection can be sustained. I shall not here discuss the more general question how far, and in what sense, one's private experiences are communicable; it will arise at a later stage when we come to consider the problems connected with one's knowledge of the minds of others.<sup>1</sup> For the present I wish only to maintain that whether or not my descriptions of my experiences are intelligible to others, their being so is not a condition of their being intelligible to myself. I agree that if I am to give my words a descriptive meaning, I must use them in accordance with some set of rules. My words must do more than simply point at my experiences: if a word applies to something it must apply to it not merely as being *this* but as being something of a certain sort. But it is not necessary that the question whether I keep or break my rules should be subject to a social check. Admittedly, if I cannot go beyond the sequence of my private feelings and impressions, if I am, as it were, in the position of one who is watching a cinema show with no power of identifying what he sees except by correlating one fleeting image with another, the means which I have for assuring myself that my use of words is consistent will be limited: I have in fact only my memory to rely on. And then it may be asked how the accuracy of my memory is itself to be tested. Only by comparing one memory with another. But is this a genuine test? Am I not then, as Wittgenstein suggests, like a man who buys several copies of the morning paper in order to assure himself that what it says is true?<sup>2</sup>

But with any use of language the same difficulty arises. Suppose that I wish to make sure that I am employing the name of some colour correctly and that, not simply trusting to my memory, I consult a colour-atlas. To profit by it, I must be able to recognize the signs and samples which it contains. I must be able to see that such and such a mark upon the page

1. *Vide ch. 5, section iv.*

2. *Philosophical Investigations*, 1. 265, p. 93.

is an inscription of the word I am concerned with; I must be able to tell whether such and such a colour which I am seeing or remembering is the same as the one with which the atlas links the word. If I have recourse to the testimony of others, I must be able to identify the shapes that they write down or the noises that they make. No doubt mistakes can always occur; but if one never accepted any identification without a further check, one would never identify anything at all. And then no descriptive use of language would be possible. But if one can recognize a word on a page, a sign made by some other person, the person himself and countless other objects, all without further ado, why should one not as immediately recognize one's own feelings and sensations? And why in that case should one not be able to describe them in accordance with certain rules of one's own? It would no doubt be an advantage if one's adherence to these rules were capable of being publicly checked, but it does not seem to be essential.

(vi) *Are mistakes about one's own immediate experience only verbal?*

For those who have the use of language, there is an intimate connection between identifying an object and knowing what to call it. Indeed on many occasions one's recognizing whatever it may be is simply a matter of one's coming out with the appropriate word. Of course the word must be meant to designate the object in question, but there are not, or need not be, two separate processes, one of fixing the object and the other of labelling it. The intention is normally to be found in the way in which the label is put on. There is, however, a sense in which one can recognize an object without knowing how to describe it. One may be able to place the object as being of the same sort as such and such another, or as having appeared before on such and such occasions, although one forgets what it is called or even thinks that it is

called something which it is not. To a certain extent this placing of the object is already a fashion of describing it: we are not now concerned with the cases where recognition, conceived in terms of adaptive behaviour, is independent of the use of any symbols at all: but our finding a description of this sort is consistent with our ignoring or infringing some relevant linguistic rule. And this can happen also when the rule is of one's own making, or at least constituted by one's own practice. When the usage which they infringe is private, such lapses can only be exceptional; for unless one's practice were generally consistent, there would be no rule to break: but it is to be envisaged that they should now and then occur.

If this is so, one can be mistaken, after all, in the characterization of one's present experience. One can at least misdescribe it in the sense that one applies the wrong word to it; wrong because it is not the word which by the rules of one's language is correlated with an 'object' of the sort in question. But the reply to this may be that one would then be making only a verbal mistake. One would be misusing words, but not falling into any error of fact. Those who maintain that statements which describe some feature of one's present experience are incorrigible need not deny that the sentences which express them may be incorrectly formulated. What they are trying to exclude is the possibility of one's being factually mistaken.

But what is supposed to be the difference in this context between a verbal and a factual mistake? The first thing to remark is that we are dealing with words which, though general in their application, are also ostensive: that is, they are meant to stand for features of what is directly given in experience. And with respect to words of this kind, it is plausible to argue that knowing what they mean is simply a matter of being disposed to use them on the right occasions, when these are presented. It then appears to follow that to be in doubt as to the nature of something which is given, to

wonder, for example, what colour this looks to me to be, is to be in doubt about the meaning of a word. And, correspondingly, to misdescribe what is given is to misuse a word. If I am not sure whether this looks crimson, what I am doubting is whether 'crimson' is the right word to describe this colour: if I resolve this doubt wrongly I have used the word 'crimson' when I should not or failed to use it when I should. This example is made easier to accept because the word 'crimson' has a conventional use. It is harder to see how I can use a word improperly when it is I alone who set the standard of propriety: my mistake would then have to consist in the fact that I had made an involuntary departure from some consistent practice which I had previously followed. In any event, it is argued, my mistake is not factual. If I were to predict that something, not yet presented to me, was going to look crimson, I might very well be making a factual mistake. My use of the word 'crimson' may be quite correct. It properly expresses my expectation: only the expectation is not in fact fulfilled. But in such a case I venture beyond the description of my present experience: I issue a draft upon the facts which they may refuse to honour. But for them to frustrate me I must put myself in their power. And this it is alleged I fail to do when I am merely recording what is directly given to me. My mistakes then can only be verbal. Thus we see that the reason why it is held to be impossible to make a factual error in describing a feature of one's present experience is that there is nothing in these circumstances which is allowed to count as one's being factually mistaken.

Against this, some philosophers would argue that it is impossible to describe anything, even a momentary private experience, without venturing beyond it. If I say that what I seem to see is crimson, I am saying that it bears the appropriate resemblance in colour to certain other objects. If it does not so resemble them I have classified it wrongly, and in doing so I have made a factual mistake. But the answer to

this is that merely from the statement that a given thing looks crimson, it cannot be deduced that anything else is coloured or even that anything else exists. The fact, if it be a fact, that the colour of the thing in question does not resemble that of other things which are properly described as crimson does indeed prove that in calling it crimson I am making a mistake; I am breaking a rule which would not exist unless there were, or at any rate could be, other things to which the word applied. But in saying that this is crimson, I am not explicitly referring to these other things. In using a word according to a rule, whether rightly or wrongly, I am not talking about the rule. I operate it but I do not say how it operates. From the fact that I have to refer to other things in order to show that my description of something is correct, it does not follow that my description itself refers to them. We may admit that to describe is to classify; but this does not entail that in describing something one is bound to go beyond it, in the sense that one actually asserts that it is related to something else.

Let us allow, then, that there can be statements which refer only to the contents of one's present experiences. Then, if it is made a necessary condition for being factually mistaken that one should make some claim upon the facts which goes beyond the content of one's present experience, it will follow that even when these statements misdescribe what they refer to the error is not factual; and then there appears no choice but to say that it is verbal. The question is whether this ruling is to be accepted.

The assumption which lies behind it is that to understand the meaning of an ostensive word one must be able to pick out the instances to which it applies. If I pick out the wrong instances, or fail to pick out the right ones, I show that I have not learned how to use the word. If I hesitate whether to apply it to a given case, I show that I am so far uncertain of its meaning. Now there is clearly some truth in this assumption. We should certainly not say that someone knew the

meaning of an ostensive word if he had no idea how to apply it; more than that, we require that his use of it should, in general, be both confident and right. But this is not to say that in every single case in which he hesitates over the application of the word, he must be in doubt about its meaning. Let us consider an example. Suppose that two lines of approximately the same length are drawn so that they both come within my field of vision and I am then asked to say whether either of them looks to me to be the longer, and if so which. I think I might very well be uncertain how to answer. But it seems very strange to say that what, in such a case, I should be uncertain about would be the meaning of the English expression 'looks longer than'. It is not at all like the case where I know which looks to me the longer, but having to reply in French, and speaking French badly, I hesitate whether to say 'plus longue' or 'plus large'. In this case I am uncertain only about the proper use of words, but in the other surely I am not. I know quite well how the words 'looks longer than' are used in English. It is just that in the present instance I am not sure whether, as a matter of fact, either of the lines does look to me to be longer than the other.

But if I can be in doubt about this matter of fact, I can presumably also come to the wrong decision. I can judge that this line looks to me to be longer than that one, when in fact it does not. This would indeed be a curious position to be in. Many would say that it was an impossible position, on the ground that there is no way of distinguishing between the way things look to someone and the way he judges that they look. After all he is the final authority on the way things look to him, and what criterion is there for deciding how things look to him except the way that he assesses them? But in allowing that he may be uncertain how a thing looks to him, we have already admitted this distinction. We have drawn a line between the facts and his assessment, or description, of

them.<sup>1</sup> Even so, it may be objected, there is no sense in talking of there being a mistake unless it is at least possible that the mistake should be discovered. And how could it ever be discovered that one had made a mistake in one's account of some momentary, private experience? Clearly no direct test is possible. The experience is past: it cannot be produced for re-inspection. But there may still be indirect evidence which would carry weight. To return to our example, if I look at the lines again, it may seem quite clear to me that A looks longer than B, whereas I had previously been inclined to think that B looked longer than A, or that they looked the same length. This does not prove that I was wrong before: it may be that they look to me differently now from the way they did then. But I might have indirect, say physiological, evidence that their appearance, that is the appearance that they offer to me, has not changed. Or I may have reason to believe that in the relevant conditions things look the same to certain other people as they do to me: and then the fact that the report given by these other people disagrees with mine may have some tendency to show that I am making a mistake. In any event it is common ground that one can misdescribe one's experience. The question is only whether such misdescription is always to be taken as an instance of a verbal mistake. My contention is that there are cases in which it is more plausible to say that the mistake is factual.

If I am right, there is then no class of descriptive statements which are incorrigible. However strong the experiential basis on which a descriptive statement is put forward, the possibility of its falsehood is not excluded. Statements which do no more than describe the content of a momentary, private experience achieve the greatest security because they run the smallest risk. But they do run some risk, however small, and because of this they too can come to grief. Complete security

1. Yes, but it may still be argued that his assessment, when he reaches it, *settles* the question. The point is whether a meaning can be given to saying that he decides wrongly. I suggest that it can.

is attained only by statements like 'I exist' which function as gesticulations. But the price which they pay for it is the sacrifice of descriptive content.

We are left still with the argument that some statements must be incorrigible, if any are ever to be verified. If the statements which have been taken as basic are fallible like all the rest, where does the process of verification terminate? The answer is that it terminates in someone's having some experience, and in his accepting the truth of some statement which describes it, or, more commonly, the truth of some more far-reaching statement which the occurrence of the experience supports. There is nothing fallible about the experience itself. What may be wrong is only one's identification of it. If an experience has been misidentified, one will be misled into thinking that some statement has been verified when it has not. But this does not mean that we never verify anything. There is no reason to doubt that the vast majority of our experiences are taken by us to be what they are; in which case they do verify the statements which are construed as describing them. What we do not, and can not, have is a logical guarantee that our acceptance of a statement is not mistaken. It is chiefly the belief that we need such a guarantee that has led philosophers to hold that some at least of the statements which refer to what is immediately given to us in experience must be incorrigible. But, as I have already remarked, even if there could be such incorrigible statements, the guarantee which they provided would not be worth very much. In any given case it would operate only for a single person and only for the fleeting moment at which he was having the experience in question. It would not, therefore, be of any help to us in making lasting additions to our stock of knowledge.

In allowing that the descriptions which people give of their experiences may be factually mistaken, we are dissociating having an experience from knowing that one has it. To know

that one is having whatever experience it may be, one must not only have it but also be able to identify it correctly, and there is no necessary transition from one to the other; not to speak of the cases when we do not identify our experiences at all, we may identify them wrongly. Once again, this does not mean that we never know, or never really know, what experiences we are having. On the contrary it is exceptional for us not to know. All that is required is that we should be able to give an account of our experiences which is both confident and correct; and these conditions are very frequently fulfilled. It is no rebuttal of our claim to knowledge that, in this as in other domains, it may sometimes happen that we think we know when we do not.

The upshot of our argument is that the philosopher's ideal of certainty has no application. Except in the cases where the truth of a statement is a condition of its being made, it can never in any circumstances be logically impossible that one should take a statement to be true when it is false; and this holds good whatever the statement may be, whether, for example, it is itself necessary or contingent. It would, however, be a mistake to express this conclusion by saying, lugubriously or in triumph, that nothing is really certain. There are a great many statements the truth of which we rightly do not doubt; and it is perfectly correct to say that they are certain. We should not be bullied by the sceptic into renouncing an expression for which we have a legitimate use. Not that the sceptic's argument is fallacious; as usual his logic is impeccable. But his victory is empty. He robs us of certainty only by so defining it as to make it certain that it cannot be obtained.

#### (vii) *How do we know?*

One reason why it is plausible to maintain that statements which do no more than describe the contents of present ex-

periences are incorrigible is that we are not required to vindicate our claims to know that they are true. It would seem absurd to ask someone how he knew that he was in pain or how he knew that what he was seeing looked to him to be of such and such a colour. For what better answer could he give than that these just were the experiences that he was having? This is not to say that there can not be independent evidence for the truth of such statements. Without it people other than the speaker would have no reason for accepting them, neither would he himself at other times. In certain cases, as we have seen, he may even use it to check the accuracy of his description of some present experience. But so long as he is actually having the experience in question, the independent evidence that there may be for its existence plays for him a subordinate rôle. His claim to know what the experience is, though it is subject to correction, is not considered to be in need of any external support.

In the ordinary way, however, the statements of fact which we claim to know are not limited to the description of our present experiences. If they refer to them at all they also refer beyond them, and in most instances they do not ostensibly refer to them at all. Even in the case of these statements we may not always be able to say how we know that they are true, but at least it is always pertinent to put the question; if no answer is obtained, the claim to knowledge becomes suspect, though it may still be upheld. To give an answer is to put forward some other statement which supports the statement of which knowledge is claimed; it is implied that this second statement is itself known to be true. Again, it may be asked how this is known, and then a third assertion may be made which supports the second. And so the process may continue until we reach a statement which we are willing to accept without a further reason. Not that it is theoretically impossible that a further reason should be found. It is just that at a certain point we decide that no further reason is required.

Thus, to ask how a statement is known to be true is to ask what grounds there are for accepting it. The question is satisfactorily answered if the grounds themselves are solid and if they provide the statement with adequate support. But here a distinction must be drawn between asking what grounds there are for accepting a given statement and asking what grounds a particular person actually has for accepting it. For example, if I am asked how I know that the earth is round, I may reply by giving the scientific evidence; in so doing I shall probably not refer to any experiences that I myself have had. But the question may also be interpreted as asking not so much how this is known as how *I* know it: and if I construe it in this way my reply will take a different form. I may mention some source from which I derived the information, some book that I have read or some person who has instructed me; I may perhaps be able to add that I have myself made some of the relevant observations, such as that of watching a ship disappear over the horizon. It may well be, however, that I cannot now recall any particular occasion on which I was informed that the earth was round, or any particular observations that I have made which go to prove it. Yet I may still say that I know this to be so, on the ground that it is common knowledge. My personal licence for the statement may be lost, but by consulting the right authorities, or by carrying out certain experiments, I can easily get it renewed. In this case, as in a great many others, I answer the question how I know by referring not to experiences that I have actually had but rather to experiences that I could have if I chose.

Since nothing is known unless somebody knows it, there is a ground for saying that the first type of answer to the question 'How do you know?' reduces to the second. Having justified a claim to knowledge by testing the scientific, or historical, evidence, one may then be asked how these supporting statements themselves are warranted. If the question

is pressed far enough, it seems that the answer must at some point take the form of saying that someone has actually observed whatever it may be. Further, since it is my claim that is being challenged, must I not end by referring, not just to observations that have been made by someone or other, but to experiences of my own? But here, as we have just seen, this second type of answer reverts to the first. For it will seldom be the case that the appropriate reference is to any particular experience that I either am, or remember, having. Nearly always, it will be a matter of claiming that I should have certain experiences if I took the proper steps. But here the point of saying that I should have these experiences is just that the facts are so; in other words, that the statements which they would verify are true. It may be held even that these two claims are equivalent, on the ground that every statement of fact is ultimately reducible to statements about possible, if not actual, experiences. Whether this is so or not is a question into which we shall have to enter later on.

However this may be, it is clear that when, as is commonly the case, a statement is accredited on the basis of certain others, their support of it must be genuine; the passage from evidence to conclusion must be legitimate. And it is at this point that the sceptic attacks. He produces arguments to show that the steps which we presume to be legitimate are not so at all. It will be found that most of our claims to knowledge are thereby put in question, and not merely our claims to knowledge but even our claims to rational belief.

(viii) *Doubts about factual reasoning:  
the problem of induction*

The range of this scepticism varies. It may be applied to all proof whatsoever or, somewhat less generally, only to all forms of experimental proof. In the second case it gives rise to the notorious problem of induction. This problem can be set out

very simply. Inductive reasoning is taken to cover all the cases in which we pass from a particular statement of fact, or set of particular statements of fact, to a factual conclusion which they do not formally entail. The inference may be from particular instances to a general law, or proceed directly by analogy from one particular instance to another. In all such reasoning we make the assumption that there is a measure of uniformity in nature; or, roughly speaking, that the future will, in the appropriate respects, resemble the past. We think ourselves entitled to treat the instances which we have been able to examine as reliable guides to those that we have not. But, as Hume pointed out, this assumption is not demonstrable; the denial that nature is uniform, to whatever degree may be in question, is not self-contradictory. Neither, as Hume also saw, is there any means of showing, without logical circularity, that the assumption is even probable. For the only way of showing that it was probable would be to produce evidence which confirmed it, and it is only if there are fair samples in nature that any evidence can be confirmatory. But whether there are fair samples in nature is just the point at issue. The same considerations apply if we seek to justify some more specific hypothesis, or would-be law of nature. Unless it is treated as a definition, in which case the problem is merely transferred to that of making sure that the definition is ever satisfied, such a proposition will not be demonstrable; the denial of it will not be self-contradictory. And once again the arguments which are meant to show that it is probable will themselves invoke the assumption that inductive reasoning is to be relied on. There are those, indeed, who think that this difficulty can be circumvented by basing their assessments of the probability of hypotheses on an *a priori* theory of probability: and much ingenious work has been done towards this end. It seems to me, however, that it has been done in vain. For the *a priori* theory of probability is just a mathematical calculus of chances. And I do not see how from a purely

formal calculus it is possible to derive any conclusion at all about what is in fact likely to happen. The calculus can indeed be used in conjunction with empirical premises: but then the justification of these empirical premises brings back the very difficulties that the appeal to the *a priori* calculus was intended to avoid.

For the most part, attempts to solve the problem of induction have taken the form of trying to fit inductive arguments into a deductive mould. The hope has been, if not to turn problematic inference into formal demonstration, at least to make it formally demonstrable that the premises of an inductive argument can in many cases confer a high degree of probability upon its conclusion. It has been thought that this could be achieved by bringing in additional premises about the constitution of the world. Logically the selection of these principles involves considerable difficulties; merely to invoke the uniformity of nature, or a law of universal causation, will not be enough. But even if we suppose the logical requirements to be somehow met, it seems clear that this enterprise must fail. For if these principles are to do the work that is expected of them, they must themselves be empirical hypotheses; and so once again the original problem returns with the question how they are to be justified.

Some philosophers of science attempt to rule out these questions altogether by saying that they arise out of a misconception of scientific method. In their view, scientists do not employ inductive reasoning; or rather, in so far as they do employ it, it is only one of the means by which they arrive at their hypotheses; they are not, or do not need to be, concerned with its validity. For what matters to them is the worth of the hypothesis itself, not the way in which it has come to be believed. And the process of testing hypotheses is deductive. The consequences which are deduced from them are subjected to empirical verification. If the result is favourable the hypothesis is retained; if not, it is modified or rejected and

another one adopted in its place. But even if this is the correct account of scientific method it does not eliminate the problem of induction. For what would be the point of testing a hypothesis except to confirm it? Why should a hypothesis which has failed the test be discarded unless this shows it to be unreliable; that is, except on the assumption that having failed once it is likely to fail again? It is true that there would be a contradiction in holding both that a hypothesis had been falsified and that it was universally valid: but there would be no contradiction in holding that a hypothesis which had been falsified was the more likely to hold good in future cases. Falsification might be regarded as a sort of infantile disease which even the healthiest hypotheses could be depended on to catch. Once they had had it there would be a smaller chance of their catching it again. But this is not in fact the view that we take. So far from approaching nature in the spirit of those gamblers at roulette who see in a long run of one colour a reason for betting on the other, we assume in general that the longer a run has been the more it is likely to continue. But how is this assumption to be justified? If this question could be answered, the problem of induction would be solved.

It does not seem, however, that it can be answered. What is demanded is a proof that what we regard as rational procedure really is so; that our conception of what constitutes good evidence is right. But of what kind is this proof supposed to be? A purely formal proof would not be applicable, and anything else is going to beg the question. For instance, it is often said that the ground for trusting scientific methods is simply that they work; the predictions which they lead us to make most commonly turn out to be true. But the fact is only that they have worked up to now. To say that they work is, in this context, to imply that they will go on working in the future. It is tacitly to assume that the future can in this matter be relied on to resemble the past. No doubt this assumption is correct, but there can be no way of proving it without its

being presupposed. So, if circular proofs are not to count, there can be no proof. And the same applies to any other assumption which might be used to guarantee the reliability of inductive reasoning. A proof which is formally correct will not do the work, and a proof which does the work will not be formally correct.

This does not mean that the use of scientific method is irrational. It could be irrational only if there were a standard of rationality which it failed to meet; whereas in fact it goes to set the standard: arguments are judged to be rational or irrational by reference to it. Neither does it follow that specific theories or hypotheses cannot be justified. The justification of a hypothesis is to be found in the evidence which favours it. But if someone chooses to deny that the fact that a hypothesis has been so favoured is a ground for continuing to trust it, he cannot be refuted; or rather he can be refuted only by reference to the standards which he questions, or rejects. No proof that we are right can be forthcoming: for at this stage nothing is going to be allowed to count as such a proof.

Thus, here again the sceptic makes his point. There is no flaw in his logic: his demand for justification is such that it is necessarily true that it cannot be met. But here again it is a bloodless victory. When it is understood that there logically could be no court of superior jurisdiction, it hardly seems troubling that inductive reasoning should be left, as it were, to act as judge in its own cause. The sceptic's merit is that he forces us to see that this must be so.

#### (ix) *The pattern of sceptical arguments*

There is, however, a special class of cases in which the problems created by the sceptic's logic are not so easily set aside. They are those in which the attack is directed, not against factual inference as such, but against some particular forms of it in which we appear to end with statements of a different

category from those with which we began. Thus doubt is thrown on the validity of our belief in the existence of physical objects, or scientific entities, or the minds of others, or the past, by an argument which seeks to show that it depends in each case upon an illegitimate inference. What is respectively put in question is our right to make the transition from sense-experiences to physical objects, from the world of common sense to the entities of science, from the overt behaviour of other people to their inner thoughts and feelings, from present to past. These are distinct problems, but the pattern of the sceptic's argument is the same in every case.

The first step is to insist that we depend entirely on the premises for our knowledge of the conclusion. Thus, it is maintained that we have no access to physical objects otherwise than through the contents of our sense-experiences, which themselves are not physical: we infer the existence of scientific entities, such as atoms and electrons, only from their alleged effects: another person's mind is revealed to us only through the state of his body or by the things he says and does: the past is known only from records or through our memories, the contents of which themselves belong to the present. Relatively to our knowledge of the evidence, our knowledge of the conclusion must in every case be indirect: and logically this could not be otherwise.

The second step in the argument is to show that the relation between premises and conclusion is not deductive. There can be no description of our sense-experiences, however long and detailed, from which it follows that a physical object exists. Statements about scientific entities are not formally deducible from any set of statements about their effects, nor do statements about a person's inner thoughts and feelings logically follow from statements about their outward manifestations. However strong the present evidence for the existence of certain past events may be, it is not demonstrative. There would be no formal contradiction in admitting the

existence of our memory-experiences, or of any other of the sources of our knowledge of the past, and yet denying that the corresponding past events had ever taken place.

But then, the argument proceeds, these inferences are not inductive either. Assuming inductive inference to be legitimate at all, it carries us, to use a phrase of Hume's, from instances of which we have experience to those of which we have none.<sup>1</sup> But here it is essential that these instances of which we in fact have no experience should be such as we are capable of experiencing. Let it be granted, in spite of the problem of induction, that on the basis of what we do experience we are sometimes entitled to infer the existence of unobserved events: our reliance on argument will then be a substitute for the direct observations which, for some practical reason, we are unable to make. The position is quite different when the things whose existence we are claiming to infer not merely are not given to us in experience but never could be. For what foundation could there be in such a case for our inductive arguments and how could their success be tested? Some philosophers even consider it to be nonsensical to assert the existence of an object which could not, at least in principle, be observed; and clearly no amount of inductive evidence can warrant a meaningless conclusion. But even if one does not go so far as to call such conclusions meaningless, it must be admitted, according to this argument, that they can have no inductive backing. Experimental reasoning can carry us forward at a given level; on the basis of certain sense-experiences it allows us to predict the occurrence of other sense-experiences; from observations of the way a person is behaving it allows us to infer that his future behaviour will take such and such a course. What it does not permit us is to jump from one level to another; to pass from premises concerning the contents of our sense-experiences to conclusions

1. *Vide David Hume, A Treatise of Human Nature, Book I, part III.*

about physical objects, from premises concerning other people's overt behaviour to conclusions about their minds.

The last step is to argue that since these inferences cannot be justified either deductively or inductively, they cannot be justified at all. We are not entitled even to make the elementary move of inferring from our present experiences to the existence of past events, or, admitting the whole range of our experiences, to arrive at the existence of physical objects: and assuming that we had sufficient warrant for believing in the existence of the physical objects which make up the world of common sense, we still should not be entitled to make the transition from these to the entities of science, or from any physical phenomena to the existence of other minds. It would indeed be hard to find even a philosopher who was willing to accept these consequences. It is scarcely to be imagined that anyone should seriously maintain that we had no right whatsoever to be sure, or even moderately confident, of anything concerning physical objects, or the minds of others, or the past. But even if he shrinks from carrying his argument to what appears to be its logical conclusion, the sceptic may still insist that it presents a question for us to answer. No doubt we do know what he says we cannot know; we are at least called upon to explain how it is possible that we should.

The problem which is presented in all these cases is that of establishing our right to make what appears to be a special sort of advance beyond our data. The level of what, for the purposes of the problem, we take to be data varies; but in every instance they are supposed to fall short, in an uncompromising fashion, of the conclusion to which we look to them to lead us. For those who wish to vindicate our claim to knowledge, the difficulty is to find a way of bridging or abolishing this gap.

Concern with the theory of knowledge is very much a matter of taking this difficulty seriously. The different ways of trying to meet it mark out different schools of philosophy,

or different methods of attacking philosophical questions. Apart from the purely sceptical position, which sets the problem, there are four main lines of approach. It is interesting that each of them consists in denying a different step in the sceptic's argument.

First, Naïve Realism. The naïve realist denies the first step of all. He will not allow that our knowledge of the various things which the sceptic wishes to put beyond our reach is necessarily indirect. His position is that the physical objects which we commonly perceive are, in a sense to be explained, directly 'given' to us, that it is not inconceivable that such things as atoms and electrons should also be directly perceived, that at least in certain favourable instances one can inspect the minds of others, that memory makes us directly acquainted with the past. The general attitude displayed is that of intuitionism. It is in the same spirit that philosophers maintain that they intuit moral values, or try to justify induction by claiming the power of apprehending necessary connections between events. But of course it is possible to take up the naïve realist's position on any one of these questions, without being committed to it on the others.

Secondly, Reductionism. The reductionist allows the first step in the sceptic's argument, but denies the second. Although his philosophical temper is diametrically opposed to that of the naïve realist, or indeed to intuitionism in any form, they have this much in common. Both of them try to close the gap which the sceptic relies on keeping open. But whereas the naïve realist does so by bringing the evidence up to the conclusion, the reductionist's policy is to bring the conclusion down to the level of the evidence. His view, which we shall presently examine, is that physical objects are logically constructed out of the contents of our sense-experiences, just as the entities of science are nothing over and above their so-called effects. In the same way, he holds that statements which appear to be about the minds of others are equivalent

to statements about their physical manifestations, and that statements which appear to be about the past are equivalent to statements about what are ordinarily regarded as records of the past, that is to statements about the present and future. Thus the conclusion, being brought down to the level of the evidence, is presented in every case as being deducible from it. It is again to be noted that one may take a reductionist view of any one of these questions without being bound to apply it to the others.

Thirdly, we have what may be called the Scientific Approach. This is the position of those who admit the first two steps in the sceptic's argument but deny the third. Unlike their predecessors, they accept the existence of the gap between evidence and conclusion, but they hold that it can be bridged by a legitimate process of inductive reasoning. Thus they will maintain that physical objects, though not directly observable in the way the naïve realists suppose, can be known to us indirectly as the causes of our sensations, just as the existence of scientific entities can be inferred from their effects, without our having to identify the two. On this view, the deliverances of memory, and other records, make the existence of the past an overwhelmingly probable hypothesis. Knowing that we ourselves have inner thoughts and feelings, we can attribute them to others by analogy.

Finally, there is the method of Descriptive Analysis. Here one does not contest the premises of the sceptic's argument, but only its conclusion. No attempt is made either to close or to bridge the gap: we are simply to take it in our stride. It is admitted that the inferences which are put in question are not deductive and also that they are not inductive, in the generally accepted sense. But this, it is held, does not condemn them. They are what they are, and none the worse for that. Moreover, they can be analysed. We can, for example, show in what conditions we feel confident in attributing certain experiences to others: we can evaluate different types of record:

we can distinguish the cases in which our memories or perceptions are taken to be reliable from those in which they are not. In short, we can give an account of the procedures that we actually follow. But no justification of these procedures is necessary or possible. One may be called upon to justify a particular conclusion, and then one can appeal to the appropriate evidence. But no more in these cases than in the case of the more general problem of induction, can there be a proof that what we take to be good evidence really is so. And if there cannot be a proof, it is not sensible to demand one. The sceptic's problems are insoluble because they are fictitious.

(x) *Remarks on the different methods of answering the sceptic*

I do not wish to say, at this stage, that any one of these approaches is, or is not, correct. If any such judgement can be made, it must follow an examination of the various problems. Except that we shall not enter into the philosophy of science, we shall deal with each of them in detail. Though we have seen that they exhibit a common pattern, there are sufficient differences between them for it to be by no means certain that a single type of answer will be appropriate in every case: we may find that a method which works well in one instance works badly in another. Again, this is a matter for particular investigation. There are, however, one or two general remarks which it may be useful to make before we enter into the details of our enquiry.

First, as to naïve realism. The strength of the naïve realist lies in his allegiance to common sense. What he knows, he knows; the arguments which go to show that he may not know it after all do not affect him; by denying the first of the sceptic's premises he absolves himself from considering the rest. Neither will he allow any tampering with the subject-matter

of his knowledge. Physical objects are physical objects, minds are minds, the past is the past. But while such truisms may be a useful corrective to the extravagances of more imaginative philosophers, they are not philosophically enlightening. In this, as in other fields, the failing of intuitionism is that it offers us no account of the way in which things are known. It may seem to offer an account, but the account is spurious. For to say that something is known by intuition or, as the naïve realist might put it, by direct acquaintance, is not at all to say *how* it is known. The addition of the explanatory phrase serves only to deny an explanation. It is justified only in the cases, if there are any, where no answer to the question how one knows is to be expected.

If the naïve realist tends to be too plain a man, the reductionist is hardly plain enough. Being willing to follow his arguments wherever they lead, he is not deterred by any appearance of paradox. To identify such things as atoms and electrons or, in another field, unconscious mental processes, with their alleged effects is not, indeed, unduly paradoxical: and perhaps the same can be said of the reduction of physical objects, like chairs and tables, to the contents of our sense-experiences, though here already there may be a protest on the part of common sense. But to maintain that when we appear to be speaking about the minds of others we are really speaking only about their bodies will seem to most people to be obviously false: while the view that all apparent references to the past are really references to the present, or future, is on the face of it preposterous. It is to be noted, however, that the reductionist does not embrace these paradoxes for their own sakes. He is convinced by argument that unless statements about the past, or physical objects, or the minds of others, are construed in this way, we can have no reason whatsoever for believing them to be true. He therefore accepts these analyses as the only alternative to outright scepticism. Since the consequences are so strange, one may suspect that

there is something wrong with the argument. But even if the reductionist can be refuted, his errors are instructive. He takes us on a philosophical journey while the naïve realist, secure in the possession of his property, is content to stay at home.

The scientific approach, as I have called it, is valuable to the extent that one does not merely insist that factual inferences from one level to another are legitimate but seriously tries to meet the arguments which go to show that they are not. If this can be achieved, the only task that remains is to show, in each case, exactly how evidence and conclusion are related. At this point, the third of our methods develops into the fourth, the method of descriptive analysis. The difference between them is important so long as it remains an open question whether the procedures, which sustain our claim to knowledge, do or do not require a proof of their legitimacy. If it can be shown that they do not, in a way that satisfactorily disposes of the sceptic's alleged disproof, then it does not greatly matter whether we regard the need for analysis as superseding the demand for justification, or whether we make the justification consist in the analysis. Assuming this to be the result, the analytic method profits by being the heir of all the rest. But it comes into its inheritance only when most of the difficult work is done. It is a weakness of some contemporary philosophers that they allow it to succeed too soon.

Having said so much in general about the questions which confront us, it is time that we developed the argument for particular cases. We shall begin with the problem of perception.