

Karl Popper

The Problem of Induction (1953, 1974)

For a brief formulation of the problem of induction we can turn to Born, who writes: '. . . no observation or experiment, however extended, can give more than a finite number of repetitions'; therefore, 'the statement of a law - B depends on A - always transcends experience. Yet this kind of statement is made everywhere and all the time, and sometimes from scanty material.' 1

In other words, the logical problem of induction arises from (1) Hume's discovery (so well expressed by Born) that it is impossible to justify a law by observation or experiment, since it 'transcends experience'; (2) the fact that science proposes and uses laws 'everywhere and all the time'. (Like Hume, Born is struck by the 'scanty material', i.e. the few observed instances upon which the law may be based.) To this we have to add (3) the principle of empiricism which asserts that in science only observation and experiment may decide upon the acceptance or rejection of scientific statements, including laws and theories.

These three principles, (1), (2), and (3), appear at first sight to clash; and this apparent clash constitutes the logical problem of induction.

Faced with this clash, Born gives up (3), the principle of empiricism (as Kant and many others, including Bertrand Russell, have done before him), in favour of what he calls a 'metaphysical principle'; a metaphysical principle which he does not even attempt to formulate; which he vaguely describes as a 'code or rule of craft'; and of which I have never seen any formulation which even looked promising and was not clearly untenable.

But in fact the principles (1) to (3) do not clash. We can see this the moment we realize that the acceptance by science of a law or of a theory is tentative only; which is to say that all laws and theories are conjectures, or tentative hypotheses (a position which I have sometimes called 'hypotheticism'); and that we may reject a law or theory on the basis of new evidence, without necessarily discarding the old evidence which originally led us to accept it. (I do not doubt that Born and many others would agree that theories are accepted only tentatively. But the widespread belief in induction shows that the far-reaching implications of this view are rarely seen.)

The principle of empiricism (3) can be fully preserved, since the fate of a theory, its acceptance or rejection, is decided by observation and experiment - by the results of tests. So long as a theory stands up to the severest tests we can design, it is accepted;

if it does not, it is rejected. But it is never inferred, in any sense, from the empirical evidence. There is neither a psychological nor a logical induction. Only the falsity of the theory can be inferred from empirical evidence and this inference is a purely deductive one.

Hume showed that it is not possible to infer a theory from observation statements; but this does not affect the possibility of refuting a theory by observation statements. The full appreciation of this possibility makes the relation between theories and observations perfectly clear.

This solves the problem of the alleged clash between the principles (1), (2), and (3), and with it Hume's problem of induction

Hume's problem of induction has almost always been badly formulated by what may be called the philosophical tradition. I will first give a few of these bad formulations, which I shall call the traditional formulations of the problem of induction. I shall replace them, however, by what I regard as better formulations.

Typical examples of formulations of the problem of induction that are both traditional and bad are the following.

What is the justification for the belief that the future will resemble the past? What is the justification of so-called inductive inferences?

By an inductive inference is here meant an inference from repeatedly observed instances to some as yet unobserved instances. It is of comparatively minor significance whether such an inference from the observed to the unobserved is, from the point of view of time, predictive or retrodictive; whether we infer that the sun will rise tomorrow or that it did rise 100,000 years ago. Of course, from a pragmatic point of view, one might say that it is the predictive type of inference which is the more important. No doubt usually it is.

There are various other philosophers who also regard as misconceived this traditional problem of induction. Some say that it is misconceived because no justification is needed for inductive inference; no more in fact than for deductive inference. Inductive inference is inductively valid just as deductive inference is deductively valid. I think it was Professor Strawson who was the first to say this.

I am of a different opinion. I hold with Hume that there simply is no such logical entity as an inductive inference; or, that all so-called inductive inferences are logically invalid - and even inductively invalid, to put it more sharply [see the end of this selection]. We have many examples of deductively valid inferences, and even some

partial criteria of deductive validity; but no example of an inductively valid inference exists. 2 And I hold, incidentally, that this result can be found in Hume, even though Hume, at the same time, and in sharp contrast to myself, believed in the psychological power of induction; not as a valid procedure, but as a procedure which animals and men successfully make use of, as a matter of fact and of biological necessity.

I take it as an important task to make clear, even at the cost of some repetition, where I agree and where I disagree with Hume.

I agree with Hume's opinion that induction is invalid and in no sense justified. Consequently neither Hume nor I can accept the traditional formulations which uncritically ask for the justification of induction; such a request is uncritical because it is blind to the possibility that induction is invalid in every sense, and therefore unjustifiable.

I disagree with Hume's opinion (the opinion incidentally of almost all philosophers) that induction is a fact and in any case needed. I hold that neither animals nor men use any procedure like induction, or any argument based on the repetition of instances. The belief that we use induction is simply a mistake. It is a kind of optical illusion.

What we do use is a method of trial and the examination of error; however misleadingly this method may look like induction, its logical structure, if we examine it closely, totally differs from that of induction. Moreover, it is a method which does not give rise to any of the difficulties connected with the problem of induction.

Thus it is not because induction can manage without justification that I am opposed to the traditional problem; on the contrary, it would urgently need justification. But the need cannot be satisfied. Induction simply does not exist, and the opposite view is a straightforward mistake.

There are many ways to present my own non-inductivist point of view. Perhaps the simplest is this. I will try to show that the whole apparatus of induction becomes unnecessary once we admit the general fallibility of human knowledge or, as I like to call it, the conjectural character of human knowledge.

Let me point this out first for the best kind of human knowledge we have; that is, for scientific knowledge. I assert that scientific knowledge is essentially conjectural or hypothetical.

Take as an example classical Newtonian mechanics. There never was a more successful theory. If repeated observational success could establish a theory, it would have established Newton's theory. Yet Newton's theory was superseded in the field of

astronomy by Einstein's theory, and in the atomic field by quantum theory. And almost all physicists think now that Newtonian classical mechanics is no more than a marvellous conjecture, a strangely successful hypothesis, and a staggeringly good approximation to the truth.

I can now formulate my central thesis, which is this. Once we fully realize the implications of the conjectural character of human knowledge, then the problem of induction changes its character completely: there is no need any longer to be disturbed by Hume's negative results, since there is no need any longer to ascribe to human knowledge a validity derived from repeated observations. Human knowledge possesses no such validity. On the other hand, we can explain all our achievements in terms of the method of trial and the elimination of error. To put it in a nutshell, our conjectures are our trial balloons, and we test them by criticizing them and by trying to replace them - by trying to show that there can be better or worse conjectures, and that they can be improved upon. The place of the problem of induction is usurped by the problem of the comparative goodness or badness of the rival conjectures or theories that have been proposed.

The main barrier to accepting the conjectural character of human knowledge, and to accepting that it contains the solution of the problem of induction, is a doctrine which may be called the commonsense theory of human knowledge or the bucket theory of the human mind. 3

I think very highly of common sense. In fact, I think that all philosophy must start from commonsense views and from their critical examination.

For our purposes here I want to distinguish two parts of the commonsense view of the world and draw attention to the fact that they clash with one another.

The first is commonsense realism; this is the view that there is a real world, with real people, animals and plants, cars and stars in it. I think that this view is true and immensely important, and I believe that no valid criticism of it has ever been proposed. [See also selection 17 below.]

A very different part of the commonsense view of the world is the commonsense theory of knowledge. The problem is the problem of how we get knowledge about the world. The commonsense solution is: by opening our eyes and ears. Our senses are the main if not the only sources of our knowledge of the world.

This second view I regard as thoroughly mistaken, and as insufficiently criticized (in spite of Leibniz and Kant). I call it the bucket theory of the mind, because it can be summed up by the diagram overleaf.

What allegedly enters the bucket through our senses are the elements, the atoms or molecules, of knowledge. Our knowledge then consists of an accumulation, a digest, or perhaps a synthesis of the elements offered to us by our senses.

Both halves of commonsense philosophy, commonsense realism and the commonsense theory of knowledge, were held by Hume; he found, as did Berkeley before him, that there is a clash between them. For the commonsense theory of knowledge is liable to lead to a kind of anti-realism. If knowledge results from sensations, then sensations are the only certain elements of knowledge, and we can have no good reason to believe that anything but sensation exists.

Hume, Berkeley, and Leibniz were all believers in a principle of sufficient reason. For Berkeley and Hume the principle took the form: if you do not have sufficient reasons for holding a belief, then this fact is itself a sufficient reason for abandoning this belief. Genuine knowledge consisted for both Berkeley and Hume essentially of belief, backed by sufficient reasons: but this led them to the position that knowledge consists, more or less, of sensations on their own.

Thus for these philosophers the real world of common sense does not really exist; according to Hume, even we ourselves do not fully exist. All that exist are sensations, impressions, and memory images. [See also selection 22, section 1, below.]

This anti-realistic view can be characterized by various names, but the most usual name seems to be 'idealism'. Hume's idealism appeared to him to be a strict refutation of commonsense realism. But though he felt rationally obliged to regard commonsense realism as a mistake, he himself admitted that he was in practice quite unable to disbelieve in commonsense realism for more than an hour.

Thus Hume experienced very strongly the clash between the two parts of commonsense philosophy: realism, and the commonsense theory of knowledge. And although he was aware that emotionally he was unable to give up realism, he looked on this fact as a mere consequence of irrational custom or habit; he was convinced that a consistent adherence to the more critical results of the theory of knowledge ought to make us abandon realism. 4 Fundamentally, Hume's idealism has remained the mainstream of British empiricism.

Hume's two problems of induction -- the logical problem and the psychological problem -- can best be presented, I think, against the background of the commonsense theory of induction. This theory is very simple. Since all knowledge is supposed to be the result of past observation, so especially is all expectational knowledge such as that the sun will rise tomorrow, or that all men are bound to die, or that bread nourishes. All this has to be the result of past observation.

It is to Hume's undying credit that he dared to challenge the commonsense view of induction, even though he never doubted that it must be largely true. He believed that induction by repetition was logically untenable - that rationally, or logically, no amount of observed instances can have the slightest bearing upon unobserved instances. This is Hume's negative solution of the problem of induction, a solution which I fully endorse.

But Hume held, at the same time, that although induction was rationally invalid, it was a psychological fact, and that we all rely on it.

Thus Hume's two problems of induction were:

(1) The logical problem: Are we rationally justified in reasoning from repeated instances of which we have had experience to instances of which we have had no experience?

Hume's unrelenting answer was: No, we are not justified, however great the number of repetitions may be. And he added that it did not make the slightest difference if, in this problem, we ask for the justification not of certain belief, but of probable belief. Instances of which we have had experience do not allow us to reason or argue about the probability of instances of which we have had no experience, any more than to the certainty of such instances.

(2) The following psychological question: How is it that nevertheless all reasonable people expect and believe that instances of which they have had no experience will conform to those of which they have had experience? Or in other words, why do we all have expectations, and why do we hold on to them with such great confidence, or such strong belief?

Hume's answer to this psychological problem of induction was: Because of 'custom or habit'; or in other words, because of the irrational but irresistible power of the law of association. We are conditioned by repetition; a conditioning mechanism without which, Hume says, we could hardly survive.

My own view is that Hume's answer to the logical problem is right and that his answer to the psychological problem is, in spite of its persuasiveness, quite mistaken.

The answers given by Hume to the logical and psychological problems of induction lead immediately to an irrationalist conclusion. According to Hume, all our knowledge, especially all our scientific knowledge, is just irrational habit or custom, and it is rationally totally indefensible.

Hume himself thought of this as a form of scepticism; but it was rather, as Bertrand Russell pointed out, an unintended surrender to irrationalism. It is an amazing fact that a peerless critical genius, one of the most rational minds of all ages, not only came to disbelieve in reason, but became a champion of unreason, of irrationalism.

Nobody has felt this paradox more strongly than Bertrand Russell, an admirer and, in many respects, even a late disciple of Hume. Thus in the Hume chapter in *A History of Western Philosophy*, published in 1946, Russell says about Hume's treatment of induction: 'Hume's philosophy ... represents the bankruptcy of eighteenth-century reasonableness' and, 'It is therefore important to discover whether there is any answer to Hume within a philosophy that is wholly or mainly empirical. If not, there is no intellectual difference between sanity and insanity. The lunatic who believes that he is a poached egg is to be condemned solely on the ground that he is in a minority ...'

Russell goes on to assert that if induction (or the principle of induction) is rejected, 'every attempt to arrive at general scientific laws from particular observations is fallacious, and Hume's scepticism is inescapable for an empiricist.'

And Russell sums up his view of the situation created by the clash between Hume's two answers, by the following dramatic remark:

'The growth of unreason throughout the nineteenth century and what has passed of the twentieth is a natural sequel to Hume's destruction of empiricism.' 5

This last quotation of Russell's goes perhaps too far. I do not wish to overdramatize the situation; and although I sometimes feel that Russell is right in his emphasis, at other moments I doubt it.

Yet the following quotation from Professor Strawson seems to me to support Russell's grave opinion: '[If] . . . there is a problem of induction, and . . . Hume posed it, it must be added that he solved it . . . [;] our acceptance of the "basic canons" [of induction] ... is forced upon us by Nature.... Reason is, and ought to be the slave of the passions.' 6

However this may be, I assert that I have an answer to Hume's psychological problem which completely removes the clash between the logic and the psychology of knowledge; and with it, it removes all of Hume's and Strawson's reasoning against reason.

My own way of avoiding Hume's irrationalist consequences is very simple. I solve the psychological problem of induction (and also such formulations as the pragmatic problem) in a manner which satisfies the following 'principle of the primacy of the logical solution', or, more briefly, the 'principle of transference'. The principle runs

like this: the solution of the logical problem of induction, far from clashing with those of the psychological or pragmatic problems, can, with some care, be directly transferred to them. As a result, there is no clash, and there are no irrationalist consequences.

The logical problem of induction itself needs some reformulation to start with.

First, it must be formulated in terms not only of 'instances' (as by Hume) but of universal regularities or laws. Regularities or laws are presupposed by Hume's own term 'instance'; for an instance is an instance of something - of a regularity or of a law. (Or, rather, it is an instance of many regularities or many laws.)

Secondly, we must widen the scope of reasoning from instances to laws so that we can take heed also of counterinstances.

In this way, we arrive at a reformulation of Hume's logical problem of induction along the following lines:

Are we rationally justified in reasoning from instances or from counterinstances of which we have had experience to the truth or falsity of the corresponding laws or to instances of which we have had no experience?

This is a purely logical problem. It is essentially merely a slight extension of Hume's logical problem of induction formulated here earlier, in section v.

The answer to this problem is: as implied by Hume, we certainly are not justified in reasoning from an instance to the truth of the corresponding law. But to this negative result a second result, equally negative, may be added: we are justified in reasoning from a counterinstance to the falsity of the corresponding universal law (that is, of any law of which it is a counterinstance). Or in other words, from a purely logical point of view, the acceptance of one counterinstance to 'All swans are white' implies the falsity of the law 'All swans are white' - that law, that is, whose counterinstance we accepted. Induction is logically invalid; but refutation or falsification is a logically valid way of arguing from a single counterinstance to - or, rather, against - the corresponding law.

This shows that I continue to agree with Hume's negative logical result; but I extend it.

This logical situation is completely independent of any question of whether we would, in practice, accept a single counterinstance - for example, a solitary black swan - in refutation of a so far highly successful law. I do not suggest that we would necessarily be so easily satisfied; we might well suspect that the black specimen before us was not

a swan. And in practice, anyway, we would be most reluctant to accept an isolated counterinstance. But this is a different question [see section IV of selection 10 below]. Logic forces us to reject even the most successful law the moment we accept one single counterinstance.

Thus we can say: Hume was right in his negative result that there can be no logically valid positive argument leading in the inductive direction. But there is a further negative result; there are logically valid negative arguments leading in the inductive direction: a counterinstance may disprove a law.

Hume's negative result establishes for good that all our universal laws or theories remain for ever guesses, conjectures, hypotheses. But the second negative result concerning the force of counterinstances by no means rules out the possibility of a positive theory of how, by purely rational arguments, we can prefer some competing conjectures to others.

In fact, we can erect a fairly elaborate logical theory of preference - preference from the point of view of the search for truth.

To put it in a nutshell, Russell's desperate remark that if with Hume we reject all positive induction, 'there is no intellectual difference between sanity and insanity' is mistaken. For the rejection of induction does not prevent us from preferring, say, Newton's theory to Kepler's, or Einstein's theory to Newton's: during our rational critical discussion of these theories we may have accepted the existence of counterexamples to Kepler's theory which do not refute Newton's, and of counterexamples to Newton's which do not refute Einstein's. Given the acceptance of these counterexamples we can say that Kepler's and Newton's theories are certainly false; whilst Einstein's may be true or it may be false: that we don't know. Thus there may exist purely intellectual preferences for one or the other of these theories; and we are very far from having to say with Russell that all the difference between science and lunacy disappears. Admittedly, Hume's argument still stands, and therefore the difference between a scientist and a lunatic is not that the first bases his theories securely upon observations while the second does not, or anything like that. Nevertheless we may now see that there may be a difference: it may be that the lunatic's theory is easily refutable by observation, while the scientist's theory has withstood severe tests.

What the scientist's and the lunatic's theories have in common is that both belong to conjectural knowledge. But some conjectures are much better than others; and this is a sufficient answer to Russell, and it is sufficient to avoid radical scepticism. For since it is possible for some conjectures to be preferable to others, it is also possible for our conjectural knowledge to improve, and to grow. (Of course, it is possible that a theory

that is preferred to another at one time may fall out of favour at a later time so that the other is now preferred to it. But, on the other hand, this may not happen.)

We may prefer some competing theories to others on purely rational grounds. It is important that we are clear what the principles of preference or selection are.

In the first place they are governed by the idea of truth. We want, if at all possible, theories which are true, and for this reason we try to eliminate the false ones.

But we want more than this. We want new and interesting truth. We are thus led to the idea of the growth of informative content, and especially of truth content. That is, we are led to the following principle of preference: a theory with a great informative content is on the whole more interesting, even before it has been tested, than a theory with little content. Admittedly, we may have to abandon the theory with the greater content, or as I also call it, the bolder theory, if it does not stand up to tests. But even in this case we may have learned more from it than from a theory with little content, for falsifying tests can sometimes reveal new and unexpected facts and problems. [See also selection 13 below.]

Thus our logical analysis leads us direct to a theory of method, and especially to the following methodological rule: try out, and aim at, bold theories, with great informative content; and then let these bold theories compete, by discussing them critically and by testing them severely.

My solution of the logical problem of induction was that we may have preferences for certain of the competing conjectures; that is, for those which are highly informative and which so far have stood up to eliminative criticism. These preferred conjectures are the result of selection, of the struggle for survival of the hypotheses under the strain of criticism, which is artificially intensified selection pressure.

The same holds for the psychological problem of induction. Here too we are faced with competing hypotheses, which may perhaps be called beliefs, and some of them are eliminated, while others survive, anyway for the time being. Animals are often eliminated along with their beliefs; or else they survive with them. Men frequently outlive their beliefs; but for as long as the beliefs survive (often a very short time), they form the (momentary or lasting) basis of action.

My thesis is that this Darwinian procedure of the selection of beliefs and actions can in no sense be described as irrational. In no way does it clash with the rational solution of the logical problem of induction. Rather, it is just the transference of the logical solution to the psychological field. (This does not mean, of course, that we never suffer from what are called 'irrational beliefs'.)

Thus with an application of the principle of transference to Hume's psychological problem Hume's irrationalist conclusions disappear.

In talking of preference I have so far discussed only the theoretician's preference - if he has any; and why it will be for the 'better', that is, more testable, theory, and for the better tested one. Of course, the theoretician may not have any preference: he may be discouraged by Hume's, and my, 'sceptical' solution to Hume's logical problem; he may say that, if he cannot make sure of finding the true theory among the competing theories, he is not interested in any method like the one described - not even if the method makes it reasonably certain that, if a true theory should be among the theories proposed, it will be among the surviving the preferred, the corroborated ones. Yet a more sanguine or more dedicated or more curious 'pure' theoretician may well be encouraged, by our analysis, to propose again and again new competing theories in the hope that one of them may be true - even if we shall never be able to make sure of any one that it is true.

Thus the pure theoretician has more than one way of action open to him; and he will choose a method such as the method of trial and the elimination of error only if his curiosity exceeds his disappointment at the unavoidable uncertainty and incompleteness of all our endeavours.

It is different with him qua man of practical action. For a man of practical action has always to choose between some more or less definite alternatives, since even inaction is a kind of action.

But every action presupposes a set of expectations, that is, of theories about the world. Which theory shall the man of action choose? Is there such a thing as a rational choice?

This leads us to the pragmatic problems of induction which to start with, we might formulate thus:

(1) Upon which theory should we rely for practical action, from a rational point of view?

(2) Which theory should we prefer for practical action, from a rational point of view?

My answer to (1) is: from a rational point of view, we should not 'rely' on any theory, for no theory has been shown to be true, or can be shown to be true (or 'reliable').

My answer to (2) is: we should prefer the best tested theory as a basis for action.

In other words, there is no 'absolute reliance'; but since we have to choose, it will be 'rational' to choose the best tested theory. This will be 'rational' in the most obvious sense of the word known to me: the best tested theory is the one which, in the light of our critical discussion, appears to be the best so far; and I do not know of anything more 'rational' than a well-conducted critical discussion.

Since this point appears not to have got home I shall try to restate it here in a slightly new way, suggested to me by David Miller. Let us forget momentarily about what theories we 'use' or 'choose' or 'base our practical actions on', and consider only the resulting proposal or decision (to do X; not to do X; to do nothing; or so on). Such a proposal can, we hope, be rationally criticized; and if we are rational agents we will want it to survive, if possible, the most testing criticism we can muster. But such criticism will freely make use of the best tested scientific theories in our possession. Consequently any proposal that ignores these theories (where they are relevant, I need hardly add) will collapse under criticism. Should any proposal remain, it will be rational to adopt it.

This seems to me all far from tautological. Indeed, it might well be challenged by challenging the italicized sentence in the last paragraph. Why, it might be asked, does rational criticism make use of the best tested although highly unreliable theories? The answer, however, is exactly the same as before. Deciding to criticize a practical proposal from the standpoint of modern medicine (rather than, say, in phrenological terms) is itself a kind of 'practical' decision (anyway it may have practical consequences). Thus the rational decision is always: adopt critical methods that have themselves withstood severe criticism.

There is, of course, an infinite regress here. But it is transparently harmless.

Now I do not particularly want to deny (or, for that matter, assert) that, in choosing the best tested theory as a basis for action, we 'rely' on it, in some sense of the word. It may therefore even be described as the most 'reliable' theory available, in some sense of this term. Yet this is not to say that it is 'reliable'. It is 'unreliable' at least in the sense that we shall always do well, even in practical action, to foresee the possibility that something may go wrong with it and with our expectations.

But it is not merely this trivial caution which we must derive from our negative reply to the pragmatic problem (1). Rather, it is of the utmost importance for the understanding of the whole problem, and especially of what I have called the traditional problem, that in spite of the 'rationality' of choosing the best tested theory as a basis of action, this choice is not 'rational' in the sense that it is based upon good reasons in favour of the expectation that it will in practice be a successful choice: there can be no good reasons in this sense, and this is precisely Hume's result. On the

contrary, even if our physical theories should be true, it is perfectly possible that the world as we know it, with all its pragmatically relevant regularities, may completely disintegrate in the next second. This should be obvious to anybody today; but I said so 7 before Hiroshima: there are infinitely many possible causes of local, partial, or total disaster.

From a pragmatic point of view, however, most of these possibilities are obviously not worth bothering about because we cannot do anything about them: they are beyond the realm of action. (I do not, of course, include atomic war among those disasters which are beyond the realm of human action, although most of us think in just this way since we cannot do more about it than about an act of God.)

All this would hold even if we could be certain that our physical and biological theories were true. But we do not know it. On the contrary, we have very good reason to suspect even the best of them; and this adds, of course, further infinities to the infinite possibilities of catastrophe.

It is this kind of consideration which makes Hume's and my own negative reply so important. For we can now see very clearly why we must beware lest our theory of knowledge proves too much. More precisely, no theory of knowledge should attempt to explain why we are successful in our attempts to explain things.

Even if we assume that we have been successful - that our physical theories are true - we can learn from our cosmology how infinitely improbable this success is: our theories tell us that the world is almost completely empty, and that empty space is filled with chaotic radiation. And almost all places which are not empty are occupied either by chaotic dust, or by gases, or by very hot stars - all in conditions which seem to make the application of any physical method of acquiring knowledge impossible.

There are many worlds, possible and actual worlds, in which a search for knowledge and for regularities would fail. And even in the world as we actually know it from the sciences, the occurrence of conditions under which life, and a search for knowledge, could arise - and succeed - seems to be almost infinitely improbable. Moreover, it seems that if ever such conditions should appear, they would be bound to disappear again, after a time which, cosmologically speaking, is very short.

It is in this sense that induction is inductively invalid, as I said above. That is to say, any strong positive reply to Hume's logical problem (say, the thesis that induction is valid) would be paradoxical. For, on the one hand, if induction is the method of science, then modern cosmology is at least roughly correct (I do not dispute this); and on the other, modern cosmology teaches us that to generalize from observations taken, for the most part, in our incredibly idiosyncratic region of the universe would almost

always be quite invalid. Thus if induction is 'inductively valid' it will almost always lead to false conclusions; and therefore it is inductively invalid.

Knowledge without Authority (1960)

This part of my lecture might be described as an attack on empiricism, as formulated for example in the following classical statement of Hume's: 'If I ask why you believe any particular matter of fact . . ., you must tell me some reason; and this reason will be some other fact, connected with it. But as you cannot proceed after this manner, in infinitum, you must at last terminate in some fact, which is present to your memory or senses; or must allow that your belief is entirely without foundation.' 1

The problem of the validity of empiricism may be roughly put as follows: is observation the ultimate source of our knowledge of nature? And if not, what are the sources of our knowledge?

These questions remain, whatever I may have said about Bacon, and even if I should have managed to make those parts of his philosophy on which I have commented somewhat unattractive for Baconians and for other empiricists.

The problem of the source of our knowledge has recently been restated as follows. If we make an assertion, we must justify it; but this means that we must be able to answer the following questions.

'How do you know? What are the sources of your assertion?' This, the empiricist holds, amounts in its turn to the question,

'What observations (or memories of observations) underlie your assertion?'

I find this string of questions quite unsatisfactory.

First of all, most of our assertions are not based upon observations, but upon all kinds of other sources. 'I read it in The Times' or perhaps 'I read it in the Encyclopaedia Britannica' is a more likely and a more definite answer to the question 'How do you know?' than 'I have observed it' or 'I know it from an observation I made last year'.

'But', the empiricist will reply, 'how do you think that The Times or the Encyclopaedia Britannica got their information? Surely, if you only carry on your inquiry long enough, you will end up with reports of the observations of eyewitnesses (sometimes called "protocol sentences" or- by yourself- "basic statements"). Admittedly, the empiricist will continue, 'books are largely made from other books. Admittedly, a historian, for example, will work from documents. But ultimately, in the last analysis,

these other books, or these documents, must have been based upon observations. Otherwise they would have to be described as poetry, or invention, or lies, but not as testimony. It is in this sense that we empiricists assert that observation must be the ultimate source of our knowledge.'

Here we have the empiricist's case, as it is still put by some of my positivist friends.

I shall try to show that this case is as little valid as Bacon's; that the answer to the question of the sources of knowledge goes against the empiricist; and, finally, that this whole question of ultimate sources - sources to which one may appeal, as one might to a higher court or a higher authority - must be rejected as based upon a mistake.

First I want to show that if you actually went on questioning The Times and its correspondents about the sources of their knowledge, you would in fact never arrive at all those observations by eyewitnesses in the existence of which the empiricist believes. You would find, rather, that with every single step you take, the need for further steps increases in snowball-like fashion.

Take as an example the sort of assertion for which reasonable people might simply accept as sufficient the answer 'I read it in The Times'; let us say the assertion 'The Prime Minister has decided to return to London several days ahead of schedule'. Now assume for a moment that somebody doubts this assertion, or feels the need to investigate its truth. What shall he do? If he has a friend in the Prime Minister's office, the simplest and most direct way would be to ring him up; and if this friend corroborates the message, then that is that.

In other words, the investigator will, if possible, try to check, or to examine, the asserted fact itself, rather than trace the source of the information. But according to the empiricist theory, the assertion 'I have read it in The Times' is merely a first step in a justification procedure consisting in tracing the ultimate source. What is the next step?

There are at least two next steps. One would be to reflect that 'I have read it in The Times' is also an assertion, and that we might ask 'What is the source of your knowledge that you read it in The Times and not, say, in a paper looking very similar to The Times?' The other is to ask The Times for the sources of its knowledge. The answer to the first question may be 'But we have only The Times on order and we always get it in the morning', which gives rise to a host of further questions about sources which we shall not pursue. The second question may elicit from the editor of The Times the answer: 'We had a telephone call from the Prime Minister's office.' Now according to the empiricist procedure, we should at this stage ask next: 'Who is the gentleman who received the telephone call?' and then get his observation report;

but we should also have to ask that gentleman: 'What is the source of your knowledge that the voice you heard came from an official in the Prime Minister's office?', and so on.

There is a simple reason why this tedious sequence of questions never comes to a satisfactory conclusion. It is this. Every witness must always make ample use, in his report, of his knowledge of persons, places, things, linguistic usages, social conventions, and so on. He cannot rely merely upon his eyes or ears, especially if his report is to be of use in justifying any assertion worth justifying. But this fact must of course always raise new questions as to the sources of those elements of his knowledge which are not immediately observational.

This is why the programme of tracing back all knowledge to its ultimate source in observation is logically impossible to carry through: it leads to an infinite regress. (The doctrine that truth is manifest cuts off the regress. This is interesting because it may help to explain the attractiveness of that doctrine.)

I wish to mention, in parentheses, that this argument is closely related to another - that all observation involves interpretation in the light of our theoretical knowledge, or that pure observational knowledge, unadulterated by theory, would, if at all possible, be utterly barren and futile. [See the last paragraph of selection 11, section 1, below. 2]

The most striking thing about the observationalist programme of asking for sources - apart from its tediousness - is its stark violation of common sense. For if we are doubtful about an assertion, then the normal procedure is to test it, rather than to ask for its sources; and if we find independent corroboration, then we shall often accept the assertion without bothering at all about sources.

Of course there are cases in which the situation is different. Testing an historical assertion always means going back to sources; but not, as a rule, to the reports of eyewitnesses.

Clearly, no historian will accept the evidence of documents uncritically. There are problems of genuineness, there are problems of bias, and there are also such problems as the reconstruction of earlier sources. There are, of course, also problems such as: was the writer present when these events happened? But this is not one of the characteristic problems of the historian. He may worry about the reliability of a report, but he will rarely worry about whether or not the writer of a document was an eyewitness of the event in question, even assuming that this event was of the nature of an observable event. A letter saying 'I changed my mind yesterday on this question' may be most valuable historical evidence, even though changes of mind are

unobservable (and even though we may conjecture, in view of other evidence, that the writer was lying).

As to eyewitnesses, they are important almost exclusively in a court of law where they can be cross-examined. As most lawyers know, eyewitnesses often err. This has been experimentally investigated, with the most striking results. Witnesses most anxious to describe an event as it happened are liable to make scores of mistakes, especially if some exciting things happen in a hurry; and if an event suggests some tempting interpretation, then this interpretation, more often than not, is allowed to distort what has actually been seen.

Hume's view of historical knowledge was different: '... we believe', he writes in the Treatise, 'that CAESAR was kill'd in the senate-house on the ides of March . . . because this fact is establish'd on the unanimous testimony of historians, who agree to assign this precise time and place to that event. Here are certain characters and letters present either to our memory or senses; which characters we likewise remember to have been us'd as the signs of certain ideas; and these ideas were either in the minds of such as were immediately present at that action, and receiv'd the ideas directly from its existence; or they were deriv'd from the testimony of others, and that again from another testimony . . . 'till we arrive at those who were eye-witnesses and spectators of the event.' 3

It seems to me that this view must lead to the infinite regress described above. For the problem is, of course, whether 'the unanimous testimony of historians' is to be accepted, or whether it is, perhaps, to be rejected as the result of their reliance on a common yet spurious source. The appeal to 'letters present either to our memory or senses' cannot have any bearing on this or on any other relevant problem of historiography.

But what, then, are the sources of our knowledge?

The answer, I think, is this: there are all kinds of sources of our knowledge; but none has authority.

We may say that The Times can be a source of knowledge, or the Encyclopaedia Britannica. We may say that certain papers in the Physical Review about a problem in physics have more authority, and are more of the character of a source, than an article about the same problem in The Times or the Encyclopaedia. But it would be quite wrong to say that the source of the article in the Physical Review must have been wholly, or even partly, observation. The source may well be the discovery of an inconsistency in another paper, or say, the discovery of the fact that a hypothesis proposed in another paper could be tested by such and such an experiment; all these

non-observational discoveries are 'sources' in the sense that they all add to our knowledge.

I do not, of course, deny that an experiment may also add to our knowledge, and in a most important manner. But it is not a source in any ultimate sense. It has always to be checked: as in the example of the news in *The Times* we do not, as a rule, question the eyewitness of an experiment, but, if we doubt the result, we may repeat the experiment, or ask somebody else to repeat it.

The fundamental mistake made by the philosophical theory of the ultimate sources of our knowledge is that it does not distinguish clearly enough between questions of origin and questions of validity. Admittedly, in the case of historiography, these two questions may sometimes coincide. The question of the validity of an historical assertion may be testable only, or mainly, in the light of the origin of certain sources. But in general the two questions are different; and in general we do not test the validity of an assertion or information by tracing its sources or its origin, but we test it, much more directly, by a critical examination of what has been asserted - of the asserted facts themselves.

Thus the empiricist's questions 'How do you know? What is the source of your assertion?' are wrongly put. They are not formulated in an inexact or slovenly manner, but they are entirely misconceived: they are questions that beg for an authoritarian answer.

The traditional systems of epistemology may be said to result from yes-answers or no-answers to questions about the sources of our knowledge. They never challenge these questions or dispute their legitimacy; the questions are taken as perfectly natural, and nobody seems to see any harm in them.

This is quite interesting, for these questions are clearly authoritarian in spirit. They can be compared with that traditional question of political theory, 'Who should rule?', which begs for an authoritarian answer such as 'the best', or 'the wisest', or 'the people', or 'the majority'. (It suggests, incidentally, such silly alternatives as 'Who should be our rulers: the capitalists or the workers?', analogous to 'What is the ultimate source of knowledge: the intellect or the senses?') This political question is wrongly put and the answers which it elicits are paradoxical [see selection 25 below]. It should be replaced by a completely different question such as How can we organize our political institutions so that bad or incompetent rulers (whom we should try not to get, but whom we so easily might get all the same) cannot do too much damage? I believe that only by changing our question in this way can we hope to proceed towards a reasonable theory of political institutions.

The question about the sources of our knowledge can be replaced in a similar way. It has always been asked in the spirit of: 'What are the best sources of our knowledge - the most reliable ones, those which will not lead us into error, and those to which we can and must turn, in case of doubt, as the last court of appeal?' I propose to assume, instead, that no such ideal sources exist - no more than ideal rulers - and that all 'sources' are liable to lead us into error at times. And I propose to replace, therefore, the question of the sources of our knowledge by the entirely different question: 'How can we hope to detect and eliminate error?'

The question of the sources of our knowledge, like so many authoritarian questions, is a genetic one. It asks for the origin of our knowledge, in the belief that knowledge may legitimize itself by its pedigree. The nobility of the racially pure knowledge, the untainted knowledge, the knowledge which derives from the highest authority, if possible from God: these are the (often unconscious) metaphysical ideas behind the question. My modified question, 'How can we hope to detect error?' may be said to derive from the view that such pure, untainted and certain sources do not exist, and that questions of origin or of purity should not be confounded with questions of validity, or of truth. This view may be said to be as old as Xenophanes. Xenophanes knew that our knowledge is guesswork, opinion - doxa rather than episteme - as shown by his verses [quoted on p. 31 above]. Yet the traditional question of the authoritative sources of knowledge is repeated even today - and very often by positivists and by other philosophers who believe themselves to be in revolt against authority.

The proper answer to my question 'How can we hope to detect and eliminate error?' is, I believe, 'By criticizing the theories or guesses of others and - if we can train ourselves to do so - by criticizing our own theories or guesses.' (The latter point is highly desirable, but not indispensable; for if we fail to criticize our own theories, there may be others to do it for us.) This answer sums up a position which I propose to call 'critical rationalism'. It is a view, an attitude, and a tradition, which we owe to the Greeks. It is very different from the 'rationalism' or 'intellectualism' of Descartes and his school, and very different even from the epistemology of Kant. Yet in the field of ethics, of moral knowledge, it was approached by Kant with his principle of autonomy- This principle expresses his realization that we must not accept the command of an authority, however exalted, as the basis of ethics. For whenever we are faced with a command by an authority, it is for us to judge, critically, whether it is moral or immoral to obey. The authority may have power to enforce its commands, and we may be powerless to resist. But if we have the physical power of choice, then the ultimate responsibility remains with us. It is our own critical decision whether to obey a command; whether to submit to an authority.

Kant boldly carried this idea into the field of religion: '. . . in whatever way', he writes, 'the Deity should be made known to you, and even . . . if He should reveal Himself to you: it is you . . . who must judge whether you are permitted to believe in Him, and to worship Him.' 4

In view of this bold statement, it seems strange that in his philosophy of science Kant did not adopt the same attitude of critical rationalism, of the critical search for error. I feel certain that it was only his acceptance of the authority of Newton's cosmology - a result of its almost unbelievable success in passing the most severe tests - which prevented Kant from doing so. If this interpretation of Kant is correct, then the critical rationalism (and also the critical empiricism) which I advocate merely put the finishing touch to Kant's own critical philosophy. And this was made possible by Einstein, who taught us that Newton's theory may well be mistaken in spite of its overwhelming success.

So my answer to the questions 'How do you know? What is the source or the basis of your assertion? What observations have led you to it?' would be: 'I do not know: my assertion was merely a guess. Never mind the source, or the sources, from which it may spring - there are many possible sources, and I may not be aware of half of them; and origins or pedigrees have in any case little bearing upon truth. But if you are interested in the problem which I tried to solve by my tentative assertion, you may help me by criticizing it as severely as you can; and if you can design some experimental test which you think might refute my assertion, I shall gladly, and to the best of my powers, help you to refute it.'

This answer applies, strictly speaking, only if the question is asked about some scientific assertion as distinct from an historical one. If my conjecture was an historical one, sources (in the non-ultimate sense) will of course come into the critical discussion of its validity. Yet fundamentally, my answer will be the same, as we have seen.

It is high time now, I think, to formulate the epistemological results of this discussion. I will put them in the form of nine theses.

(1) There are no ultimate sources of knowledge. Every source, every suggestion, is welcome; and every source, every suggestion, is open to critical examination. Except in history, we usually examine the facts themselves rather than the sources of our information.

(2) The proper epistemological question is not one about sources; rather, we ask whether the assertion made is true - that is to say, whether it agrees with the facts. (That we may operate, without getting involved in antinomies, with the idea of

objective truth in the sense of correspondence to the facts, has been shown by the work of Alfred Tarski.) And we try to find this out, as well as we can, by examining or testing the assertion itself; either in a direct way, or by examining or testing its consequences.

(3) In connection with this examination, all kinds of arguments may be relevant. A typical procedure is to examine whether our theories are consistent with our observations. But we may also examine, for example, whether our historical sources are mutually and internally consistent.

(4) Quantitatively and qualitatively by far the most important source of our knowledge - apart from inborn knowledge - is tradition. Most things we know we have learnt by example, by being told, by reading books, by learning how to criticize, how to take and to accept criticism, how to respect truth.

(5) The fact that most of the sources of our knowledge are traditional condemns anti-traditionalism as futile. But this fact must not be held to support a traditionalist attitude: every bit of our traditional knowledge (and even our inborn knowledge) is open to critical examination and may be overthrown. Nevertheless, without tradition, knowledge would be impossible.

(6) Knowledge cannot start from nothing - from a tabula rasa - nor yet from observation. The advance of knowledge consists, mainly, in the modification of earlier knowledge. Although we may sometimes, for example in archaeology, advance through a chance observation, the significance of the discovery will usually depend upon its power to modify our earlier theories.

(7) Pessimistic and optimistic epistemologies are about equally mistaken. The pessimistic cave story of Plato is the true one, and not his optimistic story of anamnesis (even though we should admit that all men, like all other animals, and even all plants, possess inborn knowledge). But although the world of appearances is indeed a world of mere shadows on the walls of our cave, we all constantly reach out beyond it; and although, as Democritus said, the truth is hidden in the deep, we can probe into the deep. There is no criterion of truth at our disposal, and this fact supports pessimism. But we do possess criteria which, if we are lucky, may allow us to recognize error and falsity. Clarity and distinctness are not criteria of truth, but such things as obscurity or confusion may indicate error. Similarly coherence cannot establish truth, but incoherence and inconsistency do establish falsehood. And, when they are recognized, our own errors provide the dim red lights which help us in groping our way out of the darkness of our cave.

(8) Neither observation nor reason is an authority. Intellectual intuition and imagination are most important, but they are not reliable: they may show us things very clearly, and yet they may mislead us. They are indispensable as the main sources of our theories; but most of our theories are false anyway. The most important function of observation and reasoning, and even of intuition and imagination, is to help us in the critical examination of those bold conjectures which are the means by which we probe into the unknown.

(9) Every solution of a problem raises new unsolved problems; the more so the deeper the original problem and the bolder its solution. The more we learn about the world, and the deeper our learning, the more conscious, specific, and articulate will be our knowledge of what we do not know, our knowledge of our ignorance. For this, indeed, is the main source of our ignorance - the fact that our knowledge can only be finite, while our ignorance must necessarily be infinite.

We may get a glimpse of the vastness of our ignorance when we contemplate the vastness of the heavens: though the mere size of the universe is not the deepest cause of our ignorance, it is one of its causes. 'Where I seem to differ from some of my friends', F.P. Ramsey wrote in a charming passage, 'is in attaching little importance to physical size. I don't feel the least humble before the vastness of the heavens. The stars may be large but they cannot think or love; and these are qualities which impress me far more than size does. I take no credit for weighing nearly seventeen stone.'⁵ I suspect that Ramsey's friends would have agreed with him about the insignificance of sheer physical size; and I suspect that if they felt humble before the vastness of the heavens, this was because they saw in it a symbol of their ignorance.

I believe that it would be worth trying to learn something about the world even if in trying to do so we should merely learn that we do not know much. This state of learned ignorance might be a help in many of our troubles. It might be well for all of us to remember that, while differing widely in the various little bits we know, in our infinite ignorance we are all equal.

There is a last question I wish to raise.

If only we look for it we can often find a true idea, worthy of being preserved in a philosophical theory which must be rejected as false. Can we find an idea like this in one of the theories of the ultimate sources of our knowledge?

I believe we can; and I suggest that it is one of the two main ideas which underlie the doctrine that the source of all our knowledge is super-natural. The first of these ideas is false, I believe, while the second is true.

The first, the false idea, is that we must justify our knowledge, or our theories, by positive reasons, that is, by reasons capable of establishing them, or at least of making them highly probable; at any rate, by better reasons than that they have so far withstood criticism. This idea implies, I suggest, that we must appeal to some ultimate or authoritative source of true knowledge; which still leaves open the character of that authority - whether it is human, like observation or reason, or super-human (and therefore super-natural).

The second idea - whose vital importance has been stressed by Russell - is that no man's authority can establish truth by decree; that we should submit to truth; that truth is above human authority.

Taken together these two ideas almost immediately yield the conclusion that the sources from which our knowledge derives must be super-human; a conclusion which tends to encourage self-righteousness and the use of force against those who refuse to see the divine truth.

Some who rightly reject this conclusion do not, unhappily, reject the first idea - the belief in the existence of ultimate sources of knowledge. Instead they reject the second idea - the thesis that truth is above human authority. They thereby endanger the idea of the objectivity of knowledge, and of common standards of criticism or rationality.

What we should do, I suggest, is to give up the idea of ultimate sources of knowledge, and admit that all human knowledge is human: that it is mixed with our errors, our prejudices, our dreams, and our hopes: that all we can do is to grope for truth even though it be beyond our reach. We may admit that our groping is often inspired, but we must be on our guard against the belief, however deeply felt, that our inspiration carries any authority, divine or otherwise. If we thus admit that there is no authority beyond the reach of criticism to be found within the whole province of our knowledge, however far it may have penetrated into the unknown, then we can retain, without danger, the idea that truth is beyond human authority. And we must retain it. For without this idea there can be no objective standards of inquiry; no criticism of our conjectures; no groping for the unknown; no quest for knowledge.

Two Kinds of Definitions (1945)

The chief danger to our philosophy, apart from laziness and woolliness, is scholasticism, . . . which is treating what is vague as if it were precise....

F. P. RAMSEY

The problem of definitions and of the 'meaning of terms' is the most important source of Aristotle's regrettably still prevailing intellectual influence, of all that verbal and empty scholasticism that haunts not only the Middle Ages, but our own contemporary philosophy; for even a philosophy as recent as that of L. Wittgenstein suffers, as we shall see, from this influence. The development of thought since Aristotle could, I think, be summed up by saying that every discipline, as long as it used the Aristotelian method of definition, has remained arrested in a state of empty verbiage and barren scholasticism, and that the degree to which the various sciences have been able to make any progress depended on the degree to which they have been able to get rid of this essentialist method. (This is why so much of our 'social science' still belongs to the Middle Ages.) The discussion of this method will have to be a little abstract, owing to the fact that the problem has been so thoroughly muddled by Plato and Aristotle, whose influence has given rise to such deep-rooted prejudices that the prospect of dispelling them does not seem very bright. In spite of all that, it is perhaps not without interest to analyse the source of so much confusion and verbiage.

Aristotle followed Plato in distinguishing between knowledge and opinion. Knowledge, or science, according to Aristotle, may be of two kinds - either demonstrative or intuitive. Demonstrative knowledge is also a knowledge of 'causes'. It consists of statements that can be demonstrated - the conclusions - together with their syllogistic demonstrations (which exhibit the 'causes' in their 'middle terms'). Intuitive knowledge consists in grasping the 'indivisible form' or essence or essential nature of a thing (if it is 'immediate', i.e. if its 'cause' is identical with its essential nature); it is the originative source of all science since it grasps the original basic premisses of all demonstrations.

Undoubtedly, Aristotle was right when he insisted that we must not attempt to prove or demonstrate all our knowledge. Every proof must proceed from premisses; the proof as such, that is to say, the derivation from the premisses, can therefore never finally settle the truth of any conclusion, but only show that the conclusion must be true provided the premisses are true. If we were to demand that the premisses should be proved in their turn, the question of truth would only be shifted back by another step to a new set of premisses, and so on, to infinity. It was in order to avoid such an infinite regress (as the logicians say) that Aristotle taught that we must assume that there are premisses which are indubitably true, and which do not need any proof; and these he called 'basic premisses'. If we take for granted the methods by which we derive conclusions from these basic premisses, then we could say that, according to Aristotle, the whole of scientific knowledge is contained in the basic premisses, and that it would all be ours if only we could obtain an encyclopaedic list of the basic premisses. But how to obtain these basic premisses? Like Plato, Aristotle believed that we obtain all knowledge ultimately by an intuitive grasp of the essences of things. 'We

can know a thing only by knowing its essence', Aristotle writes, and 'to know a thing is to know its essence'. A 'basic premiss' is, according to him, nothing but a statement describing the essence of a thing. But such a statement is just what he calls 3 a definition. Thus all 'basic premisses of proofs' are definitions.

What does a definition look like? An example of a definition would be: 'A puppy is a young dog.' The subject of such a definition sentence, the term 'puppy', is called the term to be defined (or defined term); the words 'young dog' are called the defining formula. As a rule, the defining formula is longer and more complicated than the defined term, and sometimes very much so. Aristotle considers 4 the term to be defined as a name of the essence Of a thing, and the defining formula as the description of that essence. And he insists that the defining formula must give an exhaustive description of the essence or the essential properties of the thing in question; thus a statement like 'A puppy has four legs', although true, is not a satisfactory definition, since it does not exhaust what may be called the essence of puppiness, but holds true of a horse also; and similarly the statement 'A puppy is brown', although it may be true of some, is not true of all puppies; and it describes what is not an essential but merely an accidental property of the defined term.

But the most difficult question is how we can get hold of definitions or basic premisses, and make sure that they are correct - that we have not erred, not grasped the wrong essence. Although Aristotle is not very clear on this point, there can be little doubt that, in the main, he again follows Plato. Plato taught 5 that we can grasp the Ideas with the help of some kind of unerring intellectual intuition; that is to say, we visualise or look at them with our 'mental eye', a process which he conceived as analogous to seeing, but dependent purely upon our intellect, and excluding any element that depends upon our senses. Aristotle's view is less radical and less inspired than Plato's, but in the end it amounts to the same. 6 For although he teaches that we arrive at the definition only after we have made many observations, he admits that sense experience does not in itself grasp the universal essence, and that it cannot, therefore, fully determine a definition. Eventually he simply postulates that we possess an intellectual intuition, a mental or intellectual faculty which enables us unerringly to grasp the essences of things, and to know them. And he further assumes that if we know an essence intuitively, we must be capable of describing it and therefore of defining it. (His arguments in the Posterior Analytics in favour of this theory are surprisingly weak. They consist merely in pointing out that our knowledge of the basic premisses cannot be demonstrative, since this would lead to an infinite regress, and that the basic premisses must be at least as true and as certain as the conclusions based upon them. 'It follows from this', he writes, 'that there cannot be demonstrative knowledge of the primary premisses; and since nothing but intellectual intuition can be more true than demonstrative knowledge, it follows that it must be

intellectual intuition that grasps the basic premisses.' In the *De Anima*, and in the theological part of the *Metaphysics*, we find more of an argument; for here we have a theory of intellectual intuition - that it comes into contact with its object, the essence, and that it even becomes one with its object. 'Actual knowledge is identical with its object.')

Summing up this brief analysis, we can give, I believe, a fair description of the Aristotelian ideal of perfect and complete knowledge if we say that he saw the ultimate aim of all inquiry in the compilation of an encyclopaedia containing the intuitive definitions of all essences, that is to say, their names together with their defining formulae; and that he considered the progress of knowledge as consisting in the gradual accumulation of such an encyclopaedia, in expanding it as well as in filling up the gaps in it and, of course, in the syllogistic derivation from it of 'the whole body of facts' which constitute demonstrative knowledge.

Now there can be little doubt that all these essentialist views stand in the strongest possible contrast to the methods of modern science. (I have the empirical sciences in mind, not perhaps pure mathematics.) First, although in science we do our best to find the truth, we are conscious of the fact that we can never be sure whether we have got it. We have learnt in the past, from many disappointments, that we must not expect finality. And we have learnt not to be disappointed any longer if our scientific theories are overthrown; for we can, in most cases, determine with great confidence which of any two theories is the better one. We can therefore know that we are making progress; and it is this knowledge that to most of us atones for the loss of the illusion of finality and certainty. In other words, we know that our scientific theories must always remain hypotheses, but that, in many important cases, we can find out whether or not a new hypothesis is superior to an old one. For if they are different, then they will lead to different predictions, which can often be tested experimentally; and on the basis of such a crucial experiment, we can sometimes find out that the new theory leads to satisfactory results where the old one breaks down. Thus we can say that in our search for truth, we have replaced scientific certainty by scientific progress. And this view of scientific method is corroborated by the development of science. For science does not develop by a gradual encyclopaedic accumulation of essential information, as Aristotle thought) but by a much more revolutionary method; it progresses by bold ideas, by the advancement of new and very strange theories (such as the theory that the earth is not flat, or that 'metrical space' is not flat), and by the overthrow of the old ones.

But this view of scientific method [developed in selections 9-14 below] means that in science there is no 'knowledge', in the sense in which Plato and Aristotle understood the word, in the sense which implies finality; in science, we never have sufficient reason for the belief that we have attained the truth. What we usually call 'scientific

knowledge' is, as a rule, not knowledge in this sense, but rather information regarding the various competing hypotheses and the way in which they have stood up to various tests; it is, using the language of Plato and Aristotle, information concerning the latest, and the best tested, scientific 'opinion'. This view means, furthermore, that we have no proofs in science (excepting, of course, pure mathematics and logic). In the empirical sciences, which alone can furnish us with information about the world we live in, proofs do not occur, if we mean by 'proof' an argument which establishes once and for ever the truth of a theory. (What may occur, however, are refutations of scientific theories.) On the other hand, pure mathematics and logic, which permit of proofs, give us no information about the world, but only develop the means of describing it. Thus we could say (as I have pointed out elsewhere 7): 'In so far as a scientific statement speaks about reality, it must be falsifiable; and in so far as it is not falsifiable, it does not speak about reality.' But although proof does not play any part in the empirical sciences, argument still does; indeed, its part is at least as important as that played by observation and experiment.

The role of definitions in science, especially, is also very different from what Aristotle had in mind. Aristotle taught that in a definition we have first pointed to the essence - perhaps by naming it - and that we then describe it with the help of the defining formula; just as in an ordinary sentence like 'This puppy is brown', we first point to a certain thing by saying 'this puppy', and then describe it as 'brown'. And he taught that by thus describing the essence to which the term points which is to be defined, we determine or explain the meaning 8 of the term also.

Accordingly, the definition may at one time answer two very closely related questions. The one is 'What is it?', for example 'What is a puppy?'; it asks what the essence is which is denoted by the defined term. The other is 'What does it mean?', for example, 'What does "puppy" mean?'; it asks for the meaning of a term (namely, of the term that denotes the essence). In the present context, it is not necessary to distinguish between these two questions; rather, it is important to see what they have in common; and I wish, especially, to draw attention to the fact that both questions are raised by the term that stands, in the definition, on the left side and answered by the defining formula which stands on the right side. This fact characterizes the essentialist view, from which the scientific method of definition radically differs.

While we may say that the essentialist interpretation reads a definition 'normally', that is to say, from the left to the right, we can say that a definition, as it is normally used in modern science, must be read back to front, or from the right to the left; for it starts with the defining formula, and asks for a short label for it. Thus the scientific view of the definition 'A puppy is a young dog' would be that it is an answer to the question 'What shall we call a young dog?' rather than an answer to the question 'What is a puppy?' (Questions like 'What is life?' or 'What is gravity?' do not play any role in

science.) The scientific use of definitions, characterized by the approach 'from the right to the left', may be called its nominalist interpretation, as opposed to its Aristotelian or essentialist interpretation. 9 In modern science, only 10 nominalist definitions occur, that is to say, shorthand symbols or labels are introduced in order to cut a long story short. And we can at once see from this that definitions do not play any very important part in science. For shorthand symbols can always, of course, be replaced by the longer expressions, the defining formulae, for which they stand. In some cases this would make our scientific language very cumbersome; we should waste time and paper. But we should never lose the slightest piece of factual information. Our 'scientific knowledge', in the sense in which this term may be properly used, remains entirely unaffected if we eliminate all definitions; the only effect is upon our language, which would lose, not precision, but merely brevity. (This must not be taken to mean that in science there cannot be an urgent practical need for introducing definitions, for brevity's sake.) There could hardly be a greater contrast than that between this view of the part played by definitions, and Aristotle's view. For Aristotle's essentialist definitions are the principles from which all our knowledge is derived; they thus contain all our knowledge; and they serve to substitute a long formula for a short one. As opposed to this, the scientific or nominalist definitions do not contain any knowledge whatever, not even any 'opinion'; they do nothing but introduce new arbitrary shorthand labels; they cut a long story short.

In practice, these labels are of the greatest usefulness. In order to see this, we only need to consider the extreme difficulties that would arise if a bacteriologist, whenever he spoke of a certain strain of bacteria, had to repeat its whole description (including the methods of dyeing, etc., by which it is distinguished from a number of similar species). And we may also understand, by a similar consideration, why it has so often been forgotten, even by scientists, that scientific definitions must be read 'from the right to the left', as explained above. For most people, when first studying a science, say bacteriology, must try to find out the meanings of all these new technical terms with which they are faced. In this way, they really learn the definition 'from the left to the right', substituting, as if it were an essentialist definition, a very long story for a very short one. But this is merely a psychological accident, and a teacher or writer of a textbook may indeed proceed quite differently; that is to say, he may introduce a technical term only after the need for it has arisen.

So far I have tried to show that the scientific or nominalist use of definitions is entirely different from Aristotle's essentialist method of definitions. But it can also be shown that the essentialist view of definitions is simply untenable in itself. In order not to prolong this discussion unduly, I shall criticize two only of the essentialist doctrines; two doctrines which are of significance because some influential modern schools are still based upon them. One is the esoteric doctrine of intellectual intuition,

and the other the very popular doctrine that 'we must define our terms', if we wish to be precise.

Aristotle held with Plato that we possess a faculty, intellectual intuition) by which we can visualize essences and find out which definition is the correct one, and many modern essentialists have repeated this doctrine. Other philosophers, following Kant, maintain that we do not possess anything of the sort. My opinion is that we can readily admit that we possess something which may be described as 'intellectual intuition'; or more precisely, that certain of our intellectual experiences may be thus described. Everybody who 'understands' an idea, or a point of view, or an arithmetical method, for instance, multiplication, in the sense that he has 'got the feel of it', might be said to understand that thing intuitively; and there are countless intellectual experiences of that kind. But I would insist, on the other hand, that these experiences, important as they may be for our scientific endeavours, can never serve to establish the truth of any idea or theory, however strongly somebody may feel, intuitively, that it must be true, or that it is 'self-evident.'" Such intuitions cannot even serve as an argument, although they may encourage us to look for arguments. For somebody else may have just as strong an intuition that the same theory is false. The way of science is paved with discarded theories which were once declared self-evident; Francis Bacon, for example, sneered at those who denied the self-evident truth that the sun and the stars rotated round the earth, which was obviously at rest. Intuition undoubtedly plays a great part in the life of a scientist, just as it does in the life of a poet. It leads him to his discoveries. But it may also lead him to his failures. And it always remains his private affair, as it were. Science does not ask how he has got his ideas, it is only interested in arguments that can be tested by everybody. The great mathematician, Gauss, described this situation very neatly once when he exclaimed: 'I have got my result; but I do not know yet how to get it.' All this applies, of course, to Aristotle's doctrine of intellectual intuition of so-called essences, which was propagated by Hegel and in our own time by E. Husserl and his numerous pupils; and it indicates that the 'intellectual intuition of essences' or 'pure phenomenology', as Husserl calls it, is a method of neither science nor philosophy. (The much debated question whether it is a new invention, as the pure phenomenologists think, or perhaps a version of Cartesianism or Hegelianism, can be easily decided; it is a version of Aristotelianism.)

The second doctrine to be criticized has even more important connections with modern views; and it bears especially upon the problem of verbalism. Since Aristotle, it has become widely known that one cannot prove all statements, and that an attempt to do so would break down because it would lead only to an infinite regression of proofs. But neither he nor, apparently, a great many modern writers seems to realize that the analogous attempt to define the meaning of all our terms must, in the

same way, lead to an infinite regression of definitions. The following passage from Crossman's *Plato To-day* is characteristic of a view which by implication is held by many contemporary philosophers of repute, for example, by Wittgenstein: 13 ' . . . if we do not know precisely the meaning of the words we use, we cannot discuss anything profitably. Most of the futile arguments on which we all waste time are largely due to the fact that we each have our own vague meaning for the words we use and assume that our opponents are using them in the same sense. If we defined our terms to start with, we could have far more profitable discussions. Again, we have only to read the daily papers to observe that propaganda (the modern counterpart of rhetoric) depends largely for its success on confusing the meaning of the terms. If politicians were compelled by law to define any term they wished to use, they would lose most of their popular appeal, their speeches would be shorter, and many of their disagreements would be found to be purely verbal.' This passage is very characteristic of one of the prejudices which we owe to Aristotle, of the prejudice that language can be made more precise by the use of definitions. Let us consider whether this can really be done.

First, we can see clearly that if 'politicians' (or anybody else) 'were compelled by law to define any term they wished to use', their speeches would not be shorter, but infinitely long. For a definition cannot establish the meaning of a term any more than a logical derivation can establish the truth of a statement; both can only shift this problem back. The derivation shifts the problem of truth back to the premises, the definition shifts the problem of meaning back to the defining terms (i.e., the terms that make up the defining formula). 14 But these, for many reasons, are likely to be just as vague and confusing as the terms we started with; and in any case, we should have to go on to define them in turn; which leads to new terms which too must be defined. And so on, to infinity. One sees that the demand that all our terms should be defined is just as untenable as the demand that all our statements should be proved.

At first sight this criticism may seem unfair. It may be said that what people have in mind, if they demand definitions, is the elimination of the ambiguities so often connected with words such as 15 'democracy', 'liberty', 'duty', 'religion', etc.; that it is clearly impossible to define all our terms, but possible to define some of these more dangerous terms and to leave it at that; and that the defining terms have just to be accepted, i.e., that we must stop after a step or two in order to avoid an infinite regression. This defence, however, is untenable. Admittedly, the terms mentioned are much misused. But I deny that the attempt to define them can improve matters. It can only make matters worse. That by 'defining their terms' even once, and leaving the defining terms undefined, the politicians would not be able to make their speeches shorter, is clear; for any essentialist definition, i.e. one that 'defines our terms' (as opposed to the nominalist one which introduces new technical terms), means the

substitution of a long story for a short one, as we have seen. Besides, the attempt to define terms would only increase the vagueness and confusion. For since we cannot demand that all the defining terms should be defined in their turn, a clever politician or philosopher could easily satisfy the demand for definitions. If asked what he means by 'democracy', for example, he could say 'the rule of the general will' or 'the rule of the spirit of the people'; and since he has now given a definition, and so satisfied the highest standards of precision, nobody will dare to criticize him any longer. And, indeed, how could he be criticized, since the demand that 'rule' or 'people' or 'will' or 'spirit' should be defined in their turn, puts us well on the way to an infinite regression so that everybody would hesitate to raise it? But should it be raised in spite of all that, then it can be equally easily satisfied. On the other hand, a quarrel about the question whether the definition was correct, or true, can only lead to an empty controversy about words.

Thus the essentialist view of definition breaks down, even if it does not, with Aristotle, attempt to establish the 'principles' of our knowledge, but only makes the apparently more modest demand that we should 'define the meaning of our terms'.

But undoubtedly, the demand that we speak clearly and without ambiguity is very important, and must be satisfied. Can the nominalist view satisfy it? And can nominalism escape the infinite regression?

It can. For the nominalist position there is no difficulty which corresponds to the infinite regression. As we have seen, science does not use definitions in order to determine the meaning of its terms, but only in order to introduce handy shorthand labels. And it does not depend on definitions; all definitions can be omitted without loss to the information imparted. It follows from this that in science, all the terms that are really needed must be undefined terms. How then do the sciences make sure of the meanings of their terms? Various replies to this question have been suggested, 16 but I do not think that any of them is satisfactory. The situation seems to be this. Aristotelianism and related philosophies have told us for such a long time how important it is to get a precise knowledge of the meaning of our terms that we are all inclined to believe it. And we continue to cling to this creed in spite of the unquestionable fact that philosophy, which for twenty centuries has worried about the meaning of its terms, is not only full of verbalism but also appallingly vague and ambiguous, while a science like physics which worries hardly at all about terms and their meaning, but about facts instead, has achieved great precision. This, surely, should be taken as indicating that, under Aristotelian influence, the importance of the meaning of terms has been grossly exaggerated. But I think that it indicates even more. For not only does this concentration on the problem of meaning fail to establish precision; it is itself the main source of vagueness, ambiguity, and confusion.

In science, we take care that the statements we make should never depend upon the meaning of our terms. Even where the terms are defined, we never try to derive any information from the definition, or to base any argument upon it. This is why our terms make so little trouble. We do not overburden them. We try to attach to them as little weight as possible. We do not take their 'meaning' too seriously. We are always conscious that our terms are a little vague (since we have learnt to use them only in practical applications) and we reach precision not by reducing their penumbra of vagueness, but rather by keeping well within it, by carefully phrasing our sentences in such a way that the possible shades of meaning of our terms do not matter. This is how we avoid quarrelling about words.

The view that the precision of science and of scientific language depends upon the precision of its terms is certainly very plausible, but it is none the less, I believe, a mere prejudice. The precision of a language depends, rather, just upon the fact that it takes care not to burden its terms with the task of being precise. A term like 'sand-dune' or 'wind' is certainly very vague. (How many inches high must a little sand-hill be in order to be called 'sand-dune'? How quickly must the air move in order to be called 'wind'?) However, for many of the geologist's purposes, these terms are quite sufficiently precise; and for other purposes, when a higher degree of differentiation is needed, he can always say 'dunes between 4 and 30 feet high' or 'wind of a velocity of between 20 and 40 miles an hour'. And the position in the more exact sciences is analogous. In physical measurements, for instance, we always take care to consider the range within which there may be an error; and precision does not consist in trying to reduce this range to nothing, or in pretending that there is no such range, but rather in its explicit recognition.

Even where a term has made trouble, as for instance the term 'simultaneity' in physics, it was not because its meaning was imprecise or ambiguous, but rather because of some intuitive theory which induced us to burden the term with too much meaning, or with too 'precise' a meaning, rather than with too little. What Einstein found in his analysis of simultaneity was that, when speaking of simultaneous events, physicists made a false assumption which would have been unchallengeable were there signals of infinite velocity. The fault was not that they did not mean anything, or that their meaning was ambiguous, or the term not precise enough; what Einstein found was, rather, that the elimination of a theoretical assumption, unnoticed so far because of its intuitive self-evidence, was able to remove a difficulty which had arisen in science. Accordingly, he was not really concerned with a question of the meaning of a term, but rather with the truth of a theory. It is very unlikely that it would have led to much if someone had started, apart from a definite physical problem, to improve the concept of simultaneity by analysing its 'essential] meaning', or even by analysing what physicists 'really mean' when they speak of simultaneity.

I think we can learn from this example that we should not attempt to cross our bridges before we come to them. And I also think that the preoccupation with questions concerning the meaning of terms, such as their vagueness or their ambiguity, can certainly not be justified by an appeal to Einstein's example. Such a preoccupation rests, rather, on the assumption that much depends upon the meaning of our terms, and that we operate with this meaning; and therefore it must lead to verbalism and scholasticism. From this point of view, we may criticize a doctrine like that of Wittgenstein, 17 who holds that while science investigates matters of fact, it is the business of philosophy to clarify the meaning of terms, thereby purging our language, and eliminating linguistic puzzles. It is characteristic of the views of this school that they do not lead to any chain of argument that could be rationally criticized; the school therefore addresses its subtle analyses 18 exclusively to the small esoteric circle of the initiated. This seems to suggest that any preoccupation with meaning tends to lead to that result which is so typical of Aristotelianism: scholasticism and mysticism.

Let us consider briefly how these two typical results of Aristotelianism have arisen. Aristotle insisted that demonstration or proof, and definition, are the two fundamental methods of obtaining knowledge. Considering the doctrine of proof first, it cannot be denied that it has led to countless attempts to prove more than can be proved; medieval philosophy is full of this scholasticism and the same tendency can be observed, on the Continent, down to Kant. It was Kant's criticism of all attempts to prove the existence of God which led to the romantic reaction of Fichte, Schelling, and Hegel. The new tendency is to discard proofs, and with them, any kind of rational argument. With the romantics, a new kind of dogmatism becomes fashionable, in philosophy as well as in the social sciences. It confronts us with its dictum. And we can take it or leave it. This romantic period of an oracular philosophy, called by Schopenhauer the 'age of dishonesty', is described by him as follows: 19 'The character of honesty that spirit of undertaking an inquiry together with the reader, which permeates the works of all Previous philosophers, disappears here completely. Every page witnesses that these so-called philosophers do not attempt to teach, but to bewitch the reader.'

A similar result was produced by Aristotle's doctrine of definition. First it led to a good deal of hairsplitting. But later, philosophers began to feel that one cannot argue about definitions. In this way, essentialism not only encouraged verbalism, but it also led to the disillusionment with argument, that is, with reason. Scholasticism and mysticism and despair in reason, these are the unavoidable results of the essentialism of Plato and Aristotle. And Plato's open revolt against freedom becomes, with Aristotle, a secret revolt against reason.

As we know from Aristotle himself, essentialism and the theory of definition met with strong opposition when they were first proposed, especially from Socrates's old companion Antisthenes, whose criticism seems to have been most sensible. 20 But this opposition was unfortunately defeated. The consequences of this defeat for the intellectual development of mankind can hardly be overrated.