

The Analysis of Knowledge (October, 1979)

Peter J. R. Millican

Lincoln College, Oxford

If anyone were to require proof that philosophy is more akin to madness than to sense, the recent controversy concerning the analysis of knowledge might well provide it. Theories have been supported against their rivals with counter-examples which would not seem out of place in Alice's world through the looking-glass, whilst the only point on which agreement has been reached is the somewhat less than stupendous result that in order to know something, it must be true. Indeed, were it not for the constant appeals to 'ordinary language', and some degree of consensus over particular puzzle-cases, one would be tempted to suppose that the theory of meaning being tacitly employed is that proposed by Humpty Dumpty: 'When I use a word, it means just what I choose it to mean – neither more nor less'.

The debate was precipitated by Edmund Gettier's famous refutation of the traditional view, which had involved an analysis of the statement 'S knows that P' into the three conditions: 'P is true', 'S believes that P' and 'S is justified in believing that P'. He supplied counter-examples which satisfied these criteria without being instances of knowledge, but since they were rather artificial, I would like to start my catalogue of fables with a more realistic example along the same lines. Suppose that a man in the desert sees the image of an oasis and, pointing, says 'There's an oasis over there'. In fact, he is looking at a mirage, but by a lucky coincidence there is an oasis nearby in the specified direction. It would appear that the man's true belief is justified, though we would certainly not call this a case of knowledge, as the belief is founded upon a false assumption, namely, that the man can see the oasis.

In common with most of the recent literature, this example is designed to illustrate a difficulty within the realm of belief justification. This has been the principal area of dispute, but since some writers have even cast doubt on the place of belief in knowledge, it is to the latter issue that I must first turn. Having examined whether knowledge implies belief, I shall come back to the question of justification, with the aim of investigating what more is required if a true belief is to count as knowledge. Having dealt with several unsuccessful attempts to solve this problem, I intend finally to show that the quest for a precise analysis of knowledge is misguided, arising from a failure to observe the concept's role when used in language.

I

The necessity of belief for knowledge has been contested principally on two grounds. The first of these depends on certain aspects of linguistic usage which suggest that knowledge and belief are mutually exclusive, whilst the second relies on counter-examples purporting to demonstrate specific instances of knowledge without belief. Let us look at these arguments in turn.

According to Carolyn Black, 'believing that P is characteristically weaker epistemically than believing that one knows that P. Knowledge that P does not entail belief that P in cases like this.' She gives the example:

‘I say that my books are in my office. You ask: “Do you believe that your books are in your office?” I say, “No. I know that my books are in my office.”

Miss Black undoubtedly has a point here, for it is, as she says, ‘epistemically unconventional to ascribe belief to someone to whom we are prepared to ascribe knowledge’. On the other hand, her example seems to parallel the situation when a child, asked whether his grandfather was old, replied: ‘No, he’s ancient!’. Conclusions about the logic of belief require more than such mere conversational idioms for support, since a natural interpretation of ‘I don’t believe, I know’ is to assume, as implicit, ‘I don’t *only* believe, I know’. It is, after all, quite possible to say of someone: ‘I am not saying whether he knows, but he does believe’, and even were this not so, ‘epistemically unconventional’ is not the same as ‘false’.¹

In philosophical controversies, it is often illuminating to observe exactly how much hangs on the result. In this case, as in many others, the method reveals the only point at issue to be a definition. No great epistemological problems are going to be solved by an agreement over the use of the word and if the question is simply about ‘ordinary language’, then it must be admitted that this is ambiguous. The somewhat trivial dispute can, it seems, be settled by a stipulation in favour of either the narrow or the broad sense of ‘believe’, though there are good reasons for preferring the latter. If I think that I know something, and I am mistaken, then it is certainly true that I believe it. But ‘what goes on in my mind’ can, in such a case, be precisely the same as when I do know. It follows that, if we take ‘belief’ to exclude knowledge, we are not using the term simply to describe a mental state. Since there is apparently no other word to do the job either, we must resort to such an expression as: ‘the state of mind of one who believes’. It seems to me to be far more convenient to define ‘belief’ itself in this way, and for the purposes of this essay that is what I shall do. A problem still remains, however.

Radford and Miss Black have put forward examples which appear to demonstrate that, even if we accept the broader sense of ‘belief’, such a belief is not necessary for knowledge. Radford’s librarian, who is always right about what has happened to the books, even when neurotically uncertain, shows that one need not be sure of something to know it. Jean, the French-Canadian, who correctly ‘guesses’ several dates from English history, of which he sincerely claims total ignorance, seems to be an example of knowledge that P without belief that P. Finally, Miss Black’s Greek student, who does not believe that he knows the Greek alphabet, and yet recites it perfectly in an examination, appears to be a similar case, although it is not clear exactly which proposition the student is supposed to know.

It has been argued by Armstrong that, whilst one can know something without consciously believing it, one must in these circumstances believe it *unconsciously*. Radford has countered this by claiming that Armstrong’s inference from knowledge to belief is ‘immediate and vacuous. For there is nothing else and nothing less than the subject’s answers, which, when correct, make up his knowledge, to constitute his belief.’ Radford has, however, overlooked an important point: if we characterise ‘unconscious belief’ as the phenomenon manifested when a man, though sincerely denying any opinion on a given question, invariably produces the same ‘guess’, then Armstrong’s inference is by no means vacuous. As Jones argues, faced with Jean’s quiz we would certainly want to inquire further. If it turned out that he gave different answers on different occasions, and had no inclination to stick to his original ‘guess’ in preference to alternative suggestions, I am sure we

¹ The convention according to which knowing excludes saying ‘I believe’ can, furthermore, be explained by an appeal to Grice’s ‘conversational implicature’. His ‘Co-operative Principle’ demands that one make the strongest statement one can, and thus claim knowledge, rather than mere belief, when in a position to do so.

would not say that he knew the corresponding history at all. If, on the other hand, his answers were consistent, then he would indeed possess the ‘unconscious belief’ defined above.

Having come so far, we can go no further. Although we may conclude that, in order to know that P, P must at least be one’s consistently preferred choice when faced with a relevant question, I do not see how the relation of knowledge to conscious belief is to be established. If philosophers disagree over the status of a particular example, and construct their theories accordingly, then presumably neither will be shaken by the other’s arguments. All that can be said is that, depending on which side one takes, one will have different meanings of the word ‘knowledge’ which, whilst in accord when faced with straightforward examples, nevertheless disagree when they come up against cases in which standard linguistic usage gives no clear judgement.

II

Let us now turn to the difficulty surrounding the notion of justification. Most of the positive accounts have concentrated on this, with examples and counter-examples in abundance. The principal feature of the latter has been belief based on false assumptions, or on arguments which, though reasonable, contain stages which are actually incorrect. It will not do, however, simply to add a condition to the effect that knowledge must be based on true propositions alone, as this would also exclude cases where the unsuccessful reasoning is superfluous, the conclusion being well supported by other, correct arguments. Lehrer’s sophisticated analysis avoids this pitfall, but perhaps runs into a complementary difficulty. To the traditional formulation, he attaches the following stipulation:

‘If S is completely justified in believing any false statement F which entails (but is not entailed by) P, then S would be completely justified in believing P even if S were to suppose that F is false.’²

In other words, take away all the false propositions on which S’s belief is based; if S is still justified in believing that P, then S knows that P. If he is not, then he doesn’t.

It might be that Lehrer has gone too far here, for although one correct path of reasoning is sufficient in many cases, this could depend on the emphasis given to it by the subject. If a politician supports a true assertion with ten arguments, of which only his last, which he thinks to be the least effective, is sufficient to establish the conclusion, it is not immediately obvious whether he should be credited with knowledge or not. Experience does not give a decisive verdict here, but let us allow Lehrer the benefit of the doubt until we have considered some rather more challenging objections.

One kind of situation which is likely to present difficulties for any account in terms of truth-conditions is that in which the subject’s belief is founded upon a statement which is inexact. To illustrate the point, let us take an example. Tom tells Dick: ‘The death of Harry was reported in every newspaper on my news-stand today’. Dick believes this (which is false: the Little Puddleton Gazette did not mention Harry), and infers that Harry is dead. Surely we would say that Dick knows of Harry’s death and yet, since this knowledge is based just on one incorrect proposition, it appears to violate Lehrer’s condition.³ He has foreseen the danger, however, for he adds cautiously:

² Lehrer uses the technical term ‘complete justification’ to avoid ambiguity, but its sense is the same as that of ‘justification’.

³ More versatile is the following example, in which the quoted figures can both be varied. It is thus similar in principle to the speedometer example mentioned later. Bill conducts an exhaustive poll, and on the basis of this he tells Fred: ‘82 per cent of Oxford’s electorate will vote against the government in the election’. In conversation with Bert, Fred remarks:

'We can stipulate that the expression "if S were to suppose that F is false" is short for "if S were to suppose that F is false but neither suppose anything else (except what is entailed by the supposition that F is false) which is adverse to justifying his belief that P is true, nor alter his beliefs in any way which is adverse to such justification".'

Lehrer's intention in this note is not entirely clear. If he means that, whilst supposing F to be false, S is nonetheless still *justified* in believing conclusions previously obtained when assuming F to be true (except, of course, those entailed false by the supposition that F is false), then he has failed to rule out Gettier's examples. I imagine, therefore, he is simply stating that one must not allow the supposition that F is false to influence other, independently established beliefs. Also, that this supposition does not jeopardise the justification of these other beliefs.

Returning to our example, we must show that Dick's belief would still be justified even were he to suppose that Tom's report was inaccurate (though not necessarily grossly inaccurate). An analysis of justification is called for, and Lehrer, after stating some preliminary definitions, provides this in his book. First, the *doxastic system* of a person S is defined as 'a set of statements of the form, S believes that P, S believes that Q, and so forth, which describes what S believes'. The 'corrected doxastic system' of S is 'that subset of the doxastic system resulting when every statement is deleted which describes S as believing something he would cease to believe as an impartial and disinterested truth-seeker'. It is coherence with the corrected doxastic system that constitutes justification. Formally:

'S is completely justified in believing that P if and only if, within the corrected doxastic system of S, P is believed to have a better chance of being true than the denial of P or any other statement that competes with P'.⁴

The problem, then, is to show that, were Dick to suppose that Tom's statement was false, he would still believe the proposition 'Harry is dead' to be more likely than any of its competitors. Vicious circles entangle any argument starting from the claim that he would be *justified* in believing this, so all that we can do is to seek evidence for the desired conclusion amongst Dick's present beliefs. There is no reason to suppose that his consciously considered beliefs should necessarily furnish this, and so it is to his unconsidered assumptions that we must turn. The obvious move is to credit him with a belief something like: 'Tom would not have said it unless it were substantially correct', a true belief which would be sufficient, when combined with Tom's statement, to demonstrate deductively that Harry is dead. Anything less will leave totally unsolved the difficulty of establishing what Dick's beliefs would be, and there are, besides, strong independent reasons for including such unconscious assumptions in the analysis, which might otherwise be too weak. Clark provides the example of Mr Brown, who is known by Mr. Smith to be generally reliable and honest, and yet, quite out of character, makes a wild guess when talking to him which happens to be correct. Clearly Smith's resulting belief is not knowledge, but to disallow it Lehrer must include the proposition 'Brown was not guessing' amongst Smith's beliefs.

It could be argued, in fact, that if Lehrer is right, these tacitly assumed 'principles' must be present in any situation of empirical knowledge whatsoever. Consider a subject S who knows that P, and suppose that something then happens of which S is unaware and has taken no account, but that owing to some lucky chance the truth of P is unaffected. It might be that the unknown occurrence

'According to the results of Bill's poll, 92 per cent of voters in Oxford will vote against the government'. Bert believes this, and knows that the government will lose the seat.

⁴ 'One statement is negatively relevant to a second if and only if the second statement has a lower chance of being true on the assumption that the first is true than otherwise A statement competes with all those statements believed to be negatively relevant to it within the corrected doxastic system of the man in question.'

makes no difference at all to the truth-values of S's relevant conscious beliefs, and if these alone are involved S will still be justified in believing that P in exactly the same way as he was previously. It is only by attributing to S a now false belief of the form: 'There were no other significant factors in the situation' that Lehrer can rule it out as a case of knowledge, and it is quite possible to know something without having this explicitly in mind.

It seems, then, that we must supplement Lehrer's analysis by counting amongst S's beliefs suppositions which might never have been distinctly thought at all. This will, of course, let such 'beliefs' in as possible candidates for knowledge themselves, and those with a wary eye for the looming spectre of innate knowledge may require a stipulation preventing this. Let us for the moment, though, consider them purely for the purpose of deciding whether or not S is 'conditionally justified' (that is, 'justified in believing P even if S were to suppose that F is false'). To affirm this, we must produce 'principles' which, when combined with S's conscious true beliefs, are sufficient to yield P deductively. They will only be brought in to deal with awkward puzzle-cases, and it is clear, for example, that Dick's belief can be vindicated by this method. There are instances, however, where a lot of juggling is needed to obtain adequate assumptions, and this can result in extremely cumbersome constructions.

Take Cynthia Sighs, the electronic pop star. She believes that the new magnetic tape which she has developed possesses certain unusual qualities. Firstly, that it retains its recordings perfectly for several years; secondly, that when they start to deteriorate they are distorted instantly, owing to a special chemical which is added to enable one to know that the time has come for replacement. She gives such a tape to her friend Nick Nasty, the notorious punk rocker. Unfortunately he puts a safety pin through it during testing but, concerned lest he offend her, transfers the label to a tape which he has purchased, carefully avoiding his own brand of 'inverting' tape, which distorts anything put onto it. Having recorded his latest 'song' on the doctored tape, he gives it to Cynthia to hear.

Now, Cynthia is a delicate girl, and rather than subject her sensitive ears to Nick's discordant 'punk' without warning, she runs the tape through her 'system' to check the volume levels. This gives her the correct information that the 'music' increases in intensity throughout but, unbeknown to her, the tape has in fact faded, and only reproduces about one half of the original signal, and that a little inconsistently.⁵

Does Cynthia know that Nick's composition is one uninterrupted crescendo? I would say that she does, although Lehrer, even with our allowance for unconscious ideas, would have considerable difficulty accounting for this. Her belief is based on the assumption that the tape reproduces faithfully, but she cannot fall back on the principle: 'Even if it had deteriorated, it would still yield the approximate form of the sound', for she does not believe this to be the case. Nor will it do to try 'every tape reproduces a crescendo', in view of Nick's inverting tape. Perhaps the only hope is something like: 'Nick would not have given me the tape unless, faded or not, it would yield the approximate form of the sound'. All that is shown by this clumsy attempt, however, is the total inadequacy of Lehrer's formulation when faced with such a situation.

The production of deductive principles to support the ascription of conditional justification, then, involves considerable artificiality and seems, moreover, quite implausible. There are, though, more fundamental reasons for excluding, or at least limiting, the application of any unconscious data in an analysis of knowledge.

⁵ This example can be made even more complicated. What verdict would we give if Cynthia's tape had itself faded (and therefore distorted)? Or if its distortion had been caused by Nick's safety pin?

Lehrer himself tells us that ‘if a person has evidence adequate to completely justify his belief, he may still fail to be completely justified in believing what he does because his belief is not based on that evidence’. I might know well that even numbers have even cubes and yet, asked in a test to select the cube of six from a set of four numbers, three of which are odd, guess the correct answer. My awareness of that ‘principle’ does not give me knowledge, because I do not employ it in my reasoning. Why, then, should we accord any unconscious propositions a special status in our arguments, unless it is clear that they are indeed used in the particular case under study? Evidence for this might be provided by the subject’s behaviour, so we need not necessarily rule them out completely. But to introduce them in as arbitrary a fashion as in the instances which we have been examining obviously will not do.⁶

An appeal to assumed principles cannot, it seems, give substance to Lehrer’s conditional justification. One of his own examples, however, shows that, even were it possible to do this in some other way, his formulation would still be unsatisfactory:

‘If a person has evidence adequate to completely justify his belief and his belief is based on that evidence, he may still fail to be completely justified in believing what he does. For he may be unable to provide any plausible line of reasoning to show how one could reach the conclusion he believes from the evidence that he has. For example, a detective who has a complicated mass of evidence that is conclusive evidence for the conclusion that Little Nelson is the leader of the gang might reach that conclusion from his evidence by what is nothing more than a lucky guess. Imagine that the only line of reasoning he can supply to show how he reached his conclusion is entirely fallacious or that he can supply none. In that case the detective would not be completely justified in believing what he does.’

If a subject is to be justified in his belief, not only must he use adequate evidence to establish the conclusion, he must also use it correctly. But Lehrer is hoist with his own petard, for his analysis of complete justification makes no mention of this. Consequently, if he is to disallow his detective as a possessor of knowledge, he must claim that the detective’s ‘guess’ is founded on a false belief. This will involve counting as beliefs the implications which the subject assumes to hold between the various items of evidence. And this is precisely the same as demanding deductive principles for his reasoning.

In conclusion, Lehrer’s definition of knowledge is quite unserviceable. It relies heavily on the notion of conditional justification, but his own account of justification leaves no room for a satisfactory interpretation of this. It is significant, I think, that in his treatment of puzzle-cases he totally ignores the very formulations which he has taken such pains to contrive. We would be well advised to do the same.

III

It is now time to deal with a problem raised by Sosa’s article, which threatens to complicate things considerably. He gives the following example:

‘Suppose that A and B are outside room R and seven persons P1-P7 come out of the room. A takes four of these, P1-P4 into his office and B takes the remaining three, P5-P7, into his office. Each of the seven subjects is asked the same question: “Is there a chair in room R?” P1-P4 answer: “Yes”. P5-P7 answer: “No”. When P1-P4 emerge from A’s office they are interviewed by B. A, however, does not think it necessary to interview P5-P7.’

⁶ In particular, it is circular to characterise unconscious belief in terms of potential changes in belief caused by corresponding possible suppositions.

Sosa goes on to say that, unless a suitable stipulation is included in our analysis,⁷ ‘A could know that there is a chair in room R on the basis of the testimony he heard, whereas B couldn’t on the basis of the testimony he heard. But B has just as much evidence for the proposition that there is a chair in the room as A does, i.e. four affirmative answers. And, furthermore, he has a wider range of data bearing on the issue. It would thus be preposterous to deny him entitlement to knowledge while granting it to A.’

There are several points to be made here. First, it is quite irrelevant whether A could ‘reasonably be expected to have found out’ the conflicting evidence, for if B had sent P5-P7 away the situation would still be substantially the same. Secondly, it is implausible to say that in claiming knowledge that P, one thereby rules out the existence of any evidence which would be sufficient to make one doubt that P. If I telephone five people at random, and decide on the basis of their (consistent) testimony that the Queen referred to Europe in her Christmas speech on television, am I really committing myself to the view that there are not ten people in Britain who would disagree with them? Surely, if I make a claim to knowledge my statement essentially concerns some fact and the evidence which I have for it, not merely possible evidence. But if this is true then Sosa’s argument indeed appears to raise a difficulty, so let us look again at his example, and see what we can make of it.

Sosa presents the case as a paradox, one that cannot be dealt with save by the sort of stipulation which he introduces. For this to be so, we must assume that P1-P4’s reports, in the absence of P5-P7, would yield sufficient justification for knowledge. If not, there would be no question of knowledge, and hence no paradox. Having established this,⁸ we must investigate whether or not P5-P7’s testimony is good reason to doubt that the room contains a chair. If it were badly lit and full of pillars,⁹ then presumably this would not be the case, and B, if he knew of it, would have adequate justification for knowledge himself. Obviously, either A or B might be denied knowledge if unaware of the room’s condition, but this presents no more of a problem than does any other situation in which a subject’s ignorance undermines his evaluation of conclusive evidence.

A puzzle remains only if both A and B know that the circumstances are such as to make the reports of P5-P7 a good reason to doubt that room P contains a chair. If this be so, however, then it must be due to some fluke that P5-P7 are mistaken. If we nevertheless refuse A knowledge on the grounds of their testimony to B, then it seems to follow that, given any case of empirical knowledge whatever, a miraculous occurrence supplying contrary evidence to someone other than the subject could render his ‘knowledge’ illegitimate. It is not easy to find examples to illustrate the absurdity of this result, since miracles are unfortunately none too common. I shall therefore take one from the Bible, where the account of Paul’s conversion suits our purpose admirably.

In the book of Acts, we first come across Paul as a ravager of the church, dragging off Christians and committing them to prison: ‘... breathing threats and murder against the disciples of the Lord, he went to the high priest and asked him for letters to the synagogues at Damascus, so that if he found any Christians, men or women, he might bring them bound to Jerusalem’. On the road to Damascus, Paul was converted by a voice from heaven, and when he reached his destination

⁷ Sosa includes the conditions: ‘S is justified in not believing that there is any set of [statements] with true members, which casts sufficient doubt on P to make it false that [S has] strong enough evidence for F’ and ‘There is no set of [such statements] such that S could reasonably be expected to have found out or otherwise know their truth’.

⁸ This would not be the case if, for instance, room R were full of actual-size photographs of chairs.

⁹ Sosa gives these as conditions designed to make the example more convincing.

‘proclaimed Jesus in the synagogues’. Now, let us imagine that rumours of his preaching have filtered back to Jerusalem, and have generated severe scepticism regarding the motive behind his journey. The high priest has yet to be told of Paul’s activities, and therefore still believes that he undertook the enterprise with the sole intention of persecuting the church. It is, I think, quite clear that this belief constitutes knowledge, and yet the situation here is entirely analogous to that of Sosa’s room. Applying these considerations to the latter, if P1-P4’s reports by themselves can yield knowledge, then the existence of P5-P7 makes no difference at all to A’s condition. B, of course, may be affected by their testimony, and denied knowledge on the basis of his uncertainty, but this was never in dispute.

Well, where has our discussion taken us? Despite the elimination of Sosa’s paradox, it appears that a complete analysis of knowledge is going to be a lot more involved than first we thought. We have not even touched on many problems, such as those concerning ‘basic’ beliefs (which require no justification), or those relating particularly to the status of contrary evidence. I would like to leave aside these questions of detail, however, and go on to consider difficulties of a rather more fundamental nature.

IV

Hitherto in our investigation, we have been taking it for granted that the idea of justification is relatively unproblematic, and that in order to render the traditional analysis watertight, we need but add a fourth condition to stipulate what kind is required for knowledge. Some, though, have preferred to take the ‘heroic’ step of applying the title ‘justification’ only where knowledge is attained. A fourth condition is then unnecessary, but the task remains to expand the third, to spell out this concept of sufficient justification.

The danger in following this line is that of reducing the account to triviality. Ayer, for example, apparently does just this when giving as his criterion for knowledge, possession of ‘the right to be sure’. Now, if we came across this term in a legal context (perhaps a killer’s ‘right to be sure’ that his victim was armed), we would expect it to refer to the reasonableness of a subject’s belief, and thus to involve a ‘right’ very similar to Gettier’s ‘justification’.¹⁰ It is clear, however, that Ayer has something quite different in mind. Consider his treatment of the man who consistently predicts the outcome of a lottery:

‘Now does our man who knows what the results of the lottery will be differ from one who only makes a series of lucky guesses? The answer is that, so far as the man himself is concerned, there need not be any difference. His procedure and his state of mind, when he is said to know what will happen, may be exactly the same as when it is said that he is only guessing. The difference is that to say that he knows is to concede to him the right to be sure, while to say that he is only guessing is to withhold it. ... If he were repeatedly successful ..., we might very well come to say that he knew ..., even though we could not explain how he knew We should grant him the right to be sure, simply on the basis of his success.’

Ayer’s ‘right to be sure’, then, is not dependent merely on subjective conditions, and this for very good reason.¹¹ Too weak because of the oasis example already considered, his analysis would also, it seems, be too strong, excluding our man with the fortunate gift of prophecy. But Ayer leaves

¹⁰ Gettier is certainly thinking of a ‘subjective’ justification, for he consistently uses the locution ‘S is justified in believing that P’.

¹¹ I use the term ‘subjective’ to refer to the condition of a subject’s mind, including his past or present awareness of external (‘objective’) things.

us totally in the dark as to how the possession of his 'right' is to be established. Obviously it requires a certain connection between objective and subjective factors, but we are given no clue as to the nature of this connection. I suspect that he would concede the 'right to be sure' in precisely those cases, and those alone, which he would describe as instances of knowledge. Since, however, the notion is apparently more obscure than the concept of knowledge which it is brought in to explain, I am tempted to think that, in order to discover whether somebody satisfied his criterion, Ayer would first have to decide whether or not that person knew the relevant fact. This may render his analysis sound, but hardly satisfactory.

Before leaving Ayer, let us look again at his lucky prophet. If he were sufficiently consistent, we probably would credit him with knowledge, but one of the reasons for this is that his performance is quite unaccountable. It is vital not to confuse those cases where, having all the facts, we choose to describe belief as knowledge, with those where our ignorance allows us to do no more than suppose that knowledge is achieved. Inexplicable prediction is always of the latter kind and, since there is no way of assessing the justification behind it, we can go only on the basis of its reliability. This is not to say that we justify the belief as though it were inductive: if our man foretold the outcome of many lotteries without error, we would just as readily grant him knowledge of the first as of the last. Rather, the point is that the only means by which we can examine the credentials of his prophecy is to study its results. Without a convincing run of success, we should have no reason to postulate any psychic power, and if he were even once mistaken we might deny him knowledge entirely. Something leads us to expect supernatural agencies to act infallibly (though this may be nothing more than our scepticism in cases where they do not), whereas less bizarre 'intuitions' would not have to meet such stringent demands. A farmer could be said to know what the weather will be despite occasional lapses, because of an assumption that there are fairly reliable causal connections between the state of the atmosphere and his feelings. The fact that he himself may be unable to explain them need not matter, but this can be put down to the influence of induction: the forecasts are not mysterious, for they are conditioned by past regularities.

The contrast between the epistemic status of prophet and farmer, which we have attributed to a difference of situation, has been taken by Mackie as indicating a variation in the meaning of the word 'knowledge'. Although it can be understood, he says, in a way which would apply only to epistemically reasonable belief, there is an important alternative sense, according to which it is sufficient that a belief be 'non-accidentally true'. These are two possibilities, but even finer distinctions can be drawn within them. Dissecting the class of 'non-accidental' beliefs, for example, 'we might require that there should be some actual connexion, in principle at least capable of being discovered or further explained, between the believing and what makes the belief true. Or we might be satisfied with the bare fact that if what is believed had been otherwise, the believing would have been at least likely to have been otherwise too, even if we thought that this was unexplainable even in principle.'¹² The prophet provides a case in which these narrow senses conflict, and this, claims

¹² From subsequent discussion with Mr. Mackie, I gather that he would now distinguish between the prophet and the farmer not on the basis of the unpredictability of lotteries, but rather because he would attribute the farmer's success to a history of psychological 'reinforcement'. Thus the confirmation of his previous 'feelings' about the weather has a cumulative subconscious effect which would account for his predictions, whereas the prophet's position, having no such explanation, would remain essentially unchanged even were he to obtain no 'feedback' to inform him of his past results. In this case, a subjective difference would be responsible for their different epistemic status, and not a non-subjective causal factor.

Though such an interpretation would accord well with my argument, however, I am not convinced that it can be supported by a reading of Mackie's text. The passage quoted continues:

Mackie, accounts for any uncertainty over his position. The farmer, on the other hand, passes both tests, and so could be granted knowledge though the prophet be refused. Such a ruling, however, would do no more than reveal the particular criterion being employed by the commentator, for there are many available to him, some fulfilled by both men and some fulfilled by neither. The existence of this 'spectrum of fine variations' renders it quite pointless to ask which is the true criterion of knowledge, and so his choice is more or less arbitrary. Mackie concludes that an investigation into the concept of knowledge should aim, not to decide what is 'properly' to be included, but rather to see 'just what is interesting and important about the cognitive situation that each sense of "know" marks off'.

If it can be shown that the categories of 'knowledge' which Mackie has isolated differ significantly, then his argument is certainly decisive. It is not yet clear what is to count as significant in this context, but a close scrutiny of the distinctions involved might nevertheless provide us with a fruitful line of inquiry. The most obvious feature of Mackie's approach is his concern with the question of innate knowledge, and it is hardly surprising that the principal variable amongst his criteria is the subject's awareness of the causal link between belief and fact believed. Thus:

'We might distinguish, within the sense of "know" which requires that the knower should be epistemically justified, a sense in which someone knows only if he can defend his belief from a sense in which someone knows if he has, say, made sufficient relevant observations but has now forgotten them.'

We have now been given four fairly precise senses of 'knowledge', two in which the subject is epistemically justified to some degree, and two in which his belief is merely non-accidentally true. The latter, though, seem to require no justification at all on the part of the subject, and the difference between them lies exclusively in the demand for an explainable connexion between belief and fact, which is made by the one and not by the other. Mackie supports his case in favour of these senses by citing examples resembling those of our farmer and our prophet respectively, but we have already registered a disagreement on this score, so a further examination may not be amiss.

Before proceeding, it will be helpful to clarify the point at issue. It appears that there are variations in the meaning of the word 'knowledge', and that according to at least one of these, the farmer could be granted knowledge and the prophet refused. But if the only relevant difference between them is the explicability of their belief (a non-subjective matter) then we have, with the minimal sense which covers both men, two senses of 'know' which are identical in their subjective requirements. This implies that even the non-subjective conditions for knowledge can vary from one sense to another, and therefore leaves us with little hope of ever achieving a comprehensive analysis.

'For example, if someone consistently predicts the result of some random process, such as the tossing of a coin, with, say, 90 per cent accuracy, can we admit, in the end, that his performance is in principle unexplainable but still say that he knew it would be heads on some occasion when he has correctly predicted this, on the ground that if it had not been going to be heads he would probably not have said so? Or should we decline, in this case, to say that he knew, but allow that the countryman who predicts the next day's weather, again, say, with 90 per cent accuracy knows, on the occasions when he gets it right, what the next day's weather will be, because we think that, although the countryman may be unable to give any account of how he knows it, there are fairly reliable connections between the recent and present state of the weather and, on the one hand, tomorrow's weather, and, on the other hand, what the countryman feels — that is, that the earlier weather conditions are a partial cause both of the prediction and of what is predicted? Surely we simply have two senses of "know" here, one sense in which the coin predictor and the weather forecaster both, on the occasions when they get it right, know what will happen, and another sense in which the weather forecaster knows but the coin predictor does not.'

Mackie gives no reason here to suppose that the prophet's predictions could not be due to a similar reinforcement of past feelings, and provides no basis for distinguishing him from the farmer beside the intrinsic unpredictability of lotteries.

I would like to challenge Mackie's assertion that it is essentially the mysterious nature of the prophet's belief which denies him knowledge. In the first place, this is not the only factor distinguishing him from the farmer, for if the farmer's belief is not to be equally mysterious, it must be justified by some form of induction. In the second place, there is an important consideration which has not yet been mentioned, namely, the reliability of the process responsible for the subject's belief. This is certainly related to its explicability, in that a belief which can be fully accounted for by a dependable causal sequence can thereby be shown to be trustworthy. On the other hand, there is no reason in principle why a belief should not drive a wedge between the two, and when this happens, knowledge and reliability will be found constant companions.

To take an example, let us imagine that an astrologer has discovered that every man born under a full moon possesses an infallible intuition for telling the time. Researched and thoroughly tested, this fact soon becomes established as common knowledge, and when the advantages of an internal clock are generally appreciated, scientists, who are totally mystified by the phenomenon, develop a method of reproducing it artificially by means of electronic timepieces (but never, perhaps, with quite the perfection of the original). In this situation, knowledge of the time could not be granted to those who are operated upon in preference to those with the supernatural faculty, for their beliefs, though explicable, are not any more reliable. Members of both groups could be denied knowledge if unaware of their ability, but even this strict demand for epistemic justification would be satisfied eventually by an appeal to induction. This appeal might take the form of an 'explanation' in terms of scientific 'laws', or it might refer simply to the 'brute facts' of experience. Since, however, scientific laws are ultimately brute facts themselves, the difference is essentially one of degree.

In the minimal sense of knowledge, the sense which requires no subjective justification, explicability is relevant only as an indication of reliability. Indeed, the production of a belief by a reliable process, which is obviously necessary for any kind of knowledge at all, is also in this case quite sufficient. Mackie's distinction, drawn within this single unambiguous criterion, is but a confusion of truth with demonstration.¹³

Having dealt with explicability, the only variable now left amongst our possible requirements for knowledge is the subject's awareness of the causes of his belief. Beyond this, it is simply necessary that the belief be non-accidentally true. Yet here there is some obscurity, for as Harman says, 'what counts as an accident depends on what considerations one has in mind as well as upon how one describes what is to count as the accident'. Sosa gives the case of a palm reader S who has been indoctrinated to believe that palm pattern A is a symptom of a forthcoming disease D, where the indoctrination has no observational foundation but where, unknown to S and his mentor, A *is* a legitimate symptom of D. S could not, on the basis of A, know that someone will contract D, though his belief would be, on at least one interpretation, non-accidentally true. It is clearly the random nature of his mentor's teaching which provides the 'accident' in this instance, but other examples can be produced which are far less straightforward,

Suppose that a man S regularly drives to work at what he believes to be a steady speed of 40 m.p.h., and that his journey takes him along a road with a speed limit of 50 m.p.h. His speedometer, which is initially good, gradually corrodes, with the result that after one week it is only accurate to

¹³ Of course, I do not deny that Mackie could frame a definition of 'knowledge' in terms of explicability, only that such a definition would correspond to any sense of the word as it is actually used.

plus or minus 1 m.p.h. Deterioration continues at the same rate, and within a few months S is arrested for speeding.

At what point does S cease to know that he is complying with the law? A precise answer to this question cannot be given, for the decay of S's speedometer is continuous. True, there is no risk of his breaking the speed limit until the tenth week has elapsed, but it is implausible to say that he is epistemically justified up to that time. Neither is it any help to ask when his belief is non-accidentally true: is it an accident that after nine weeks his speedometer is sufficiently accurate to ensure his safety?

It is impossible to specify exactly the conditions which must be fulfilled if a belief is to be classed as non-accidentally true. Similarly, it is futile to attempt a precise definition of knowledge, for unless it takes into account the vagueness of language it will be forced to make a decision in cases which would correctly be described as undecidable.

But this is not all. We are still left with the variable subjective criteria for knowledge indicated by Mackie, and any choice between these can only be arbitrary. 'It is highly probable that competent speakers of English use the word "know" sometimes in one sense, sometimes in another, and perhaps sometimes with a sense that is indeterminate between several precise senses that might be distinguished. And even if they did not, even if there were near unanimity in favour of just one use, what would it matter? The other senses that can be distinguished would still be possible senses of the word, they would still point to potentially interesting differences in the field of cognitive attitudes and situations.' Briefly, there are some cases (such as Ayer's prophet) in which it is equally plausible both to grant knowledge and to withhold it. Here the decision will depend on the commentator's chosen criterion, and will therefore add nothing to the situation, being rather an expression of his own epistemic attitudes.

Why, then, do we rarely find philosophers disagreeing when faced with unusual examples? The reasons are twofold. First, most of the proposed analyses have concentrated on the non-subjective requirements for knowledge and these, though vague, are not ambiguous. Causal theories, for instance, try to give a rigorous account of the intuitive notion of non-accidental truth, and it is precisely because our intuitions generally coincide in this respect that puzzle-cases provide a useful tool of analysis. Secondly, on the subjective side, there is 'the purely contingent fact that men have few, if any, beliefs that would drive a deep wedge between the various possible concepts of knowledge'. Nearly all our non-accidentally true beliefs are based on our own reasonings and observations, though frequently these would be invoked but indirectly, to establish the reliability of others' testimony. Even intuitions such as the internal clock of our example would eventually assume the mantle of authoritative knowledge when guaranteed by induction, and it is thus only inconsistent intuitions, or those not yet guaranteed, which are left to pose any problem.

But now we can see a motive for disallowing such inauthoritative 'knowledge'. Men unable to justify their non-accidentally true beliefs would be in a hopelessly weak position when assailed by sceptical doubts, and would have no means of determining which of their beliefs were trustworthy. Although it might be gratuitous to deny them 'real' knowledge, we can get beyond mere words and legitimately deny them 'durable' knowledge on the very practical criteria of stability and usefulness. 'True opinions', says Socrates, 'are a fine thing and do all sorts of good so long as they stay in their place, but they will not stay long. They run away from a man's mind, so they are not worth much until you tether them by working out the reason. Once they are tied down, they become knowledge, and are stable. That is why knowledge is something more valuable than right opinion. What distinguishes one from the other is the tether.' Knowledge that can be defended against doubt can

furthermore be defended against dissent and, through an explanation of one's reasons for belief, be imparted to others. Justification may not be logically necessary for knowledge, but practically it is indispensable.

Utility can account for the requirement that knowledge be justified, but also for the requirement that knowledge be believed. Without belief it will be useless and incommunicable, and only in the artificial situation of an examination does the issue arise. Indeed, belief is sometimes our chief concern when speaking of knowledge': where the relevant fact is not itself in dispute, a question as to a person's knowledge is often no more than an inquiry regarding his assent to that fact.

V

The attractions of the traditional analysis can be seen easily in the light of linguistic practice. It will be interesting, therefore, to examine in the same light the points raised by the counter-examples which refute that analysis.

The principal function of the concept of knowledge is to convey information, but it is not merely sufficient that the information be true: justification of the subject's belief is necessary to establish its truth. This involves showing that the belief implies the fact believed; that there is a reliable causal connection from the one to the other. Now, the fact will not be established successfully if any step in the justifying reasoning fails; practically, if somebody challenges a stage in the argument, he will not be convinced of its conclusion. Thus a belief, to be capable of resisting any genuine, factually-based challenge, must be justified in such a way that the justification would remain valid even were that challenge to be taken into account by the subject. And this seems virtually identical to Lehrer's formula for knowledge in terms of 'conditional justification'! There is, though, a subtle difference, which can best be brought out by an example.

Let us imagine Professor Gettier sitting in his study, hard at work on his famous article. Having produced counter-examples to the 'justified true belief' analysis of knowledge, he concludes with the remark that they are equally effective against Ayer. But here he is mistaken, for Ayer's 'right to be sure' is not simply subjective justification. It avoids subjectivity, as we have seen, only at the cost of triviality, and it might be that Gettier would realise this immediately were he to consider the matter. Until he had done so, however, we could not say that he knew Ayer's analysis to be inadequate. His belief in its inadequacy would still be justified if he were to suppose that 'the right to be sure' is not subjective (so Lehrer's condition is fulfilled), but he does not know, because it would then be justified by an argument which has not yet occurred to him.

If a belief is to be knowledge, the reasoning necessary to support it when challenged must already lie within the intellectual resources of the subject. Such a condition only applies, of course, to 'authoritative' knowledge, but we would do well to confine our present attention to this kind anyway, since it is undoubtedly the most interesting from a practical point of view. The knowledge with which we are generally concerned is that supplied by sources of information which we regard as reliable, and we would not regard as reliable any process (mental or physical) whose validity was rendered questionable by counter-evidence. Thus reasoning which cannot be defended when faced with conflicting facts will not yield knowledge, and, putting this formally, we reach the following stipulation: 'S knows that P only if, for any false statement F, if S were to suppose that F is false S would still be justified in believing that P by essentially the same reasoning which now justifies him'. This condition, when combined with the three of the traditional account, appears to be both necessary and sufficient for knowledge, but rejoicing would be premature until two further problems

have been dealt with. The first problem is that of explaining cases such as the news-stand and Cynthia's tape. In the latter, particularly, we would not expect the subject's reasoning to remain *exactly* the same upon supposing the relevant falsehood to be false, simply because such a supposition would rob the subject of an appropriate premise ('All of Tom's newspapers reported Harry's death' and 'This is my tape' respectively). Intuitively the argument might be 'essentially' the same, but it is difficult to give substance to this notion without begging the question. Perhaps a provision could be framed to allow for cases where an incorrect proposition, which approximates to a correct one, is used to obtain conclusions which that correct proposition would yield, but such a provision would have to be vague, and its application to the Cynthia example would be far from clear. Besides, many beliefs are not justified by a straightforward argument.

The second problem arises when, as in the case of Sosa's room, there is misleading evidence which is not available to the subject. This evidence might be consistent with every stage of the subject's reasoning, and yet be such that, had he known of it, his belief would not be justified. Our stipulation would rule out all cases of this type, but since in the Paul example we were inclined to grant the high priest knowledge, it seems that the identification of knowledge with 'belief resistant to factual challenge' is incorrect. We could, of course, modify the stipulation to 'S knows that P if and only if, for any false statement F, if S were to suppose that F is false S would still be able to reach the conclusion that P, without inconsistency and by essentially the same reasoning which now justifies him'. But Harman argues that this would be too weak, supporting his position with the following example.

Suppose a man reads in a newspaper that a political leader has been assassinated. The story is written by a dependable reporter who in fact witnessed and accurately reported the event. The leader's associates, however, fearing a coup, decide to pretend that the bullet hit somebody else, and make announcements on radio and television to that effect. Could we say that our man who, by accident, has not heard the denials, knows that the leader has been assassinated? Both Harman and Lehrer claim that we could not. Suppose, though, that as the political leader's associates are about to make their announcement, a saboteur cuts the wire leading to the transmitter. The announcement is therefore heard only by those in the studio, all of whom are parties to the deception. Now, we are told, our man does come to know that the leader has been assassinated. Yet as before, if he had known what was to have been announced, he would not have been justified in his belief simply on the basis of the newspaper story. 'Here', says Harman, 'a cut wire makes the difference between evidence that undermines knowledge and evidence that does not undermine knowledge'. And he follows this surprising result with a less surprising confession: 'I am unable to formulate criteria that would distinguish among these cases'.

If Harman is right, then any stipulation of the kind which we have been considering will be unsatisfactory, for it will fail to draw a line between the two parts of his example. But his evaluation is certainly open to question, and it is noticeable that Lehrer, who reaches the same conclusion in this instance, nevertheless disagrees with him over the assessment of a parallel case. Intuitions are not unanimous in such examples, and even though many of them can be analysed in the manner of Sosa's room, there will always remain a decision to be made when the analysis is complete.¹⁴ This decision cannot be made by a prior theory of knowledge, for it itself must supply a premise of that theory. Thus, once again, the choice is not between a correct and an incorrect account of knowledge, but rather between (at least) two possible senses of the word, senses which will be selected according to their correspondence with the commentator's view of particular puzzle-cases.

¹⁴ As over the status of A's belief when it is established that a 'miracle' is responsible for B's misleading evidence.

Before attempting to summarise our findings, it will be interesting to ask why the contentious and somewhat odd claim that misleading evidence can undermine knowledge should seem plausible. The answer is that the concept of knowledge is not exclusively subject-centred; that although the commentator can put himself in the place of the subject to examine a belief's justification, he can equally well adopt what he sees as the position of an 'ordinary man in the street'. In this guise, he will not be convinced of a fact by anyone who is unaware of some generally acknowledged contrary evidence (and so might deny that person knowledge). Where such evidence is little known, however, it will present no epistemic obstacle to the ignorant 'common man'.¹⁵

VI

Knowledge can be centred on the subject, but need not be. The subject must believe it, though perhaps not consciously. His belief should be justified, but not necessarily. That justification must be valid, if only inexactly. With every new consideration our analysis seems to become more vague, and more ambiguous.

Knowledge is tied to the pursuit of truth; it is the relation between a man and a fact, when the man can furnish evidence sufficient to establish that fact. But what form should this evidence take: memories, beliefs, or just answers to questions? And must it be 'sufficient' for the subject, for an observer, or just causally? Replies will vary, and with them the meaning of 'knowledge'.

Knowledge requires that a fact be implied by a situation; and any situation may be specified from a number of different points of view. In particular, the commentator himself might infer some fact from circumstances which include the subject's belief, and grant him knowledge on those grounds alone. Justification, even for the commentator, is sometimes important only as a means of reaching the truth.

Knowledge is spoken of in conveying information, and this information will generally concern either someone's belief or its content. If we are aware of a fact, we might wish to know whether our subject believes that fact. If we know of his belief, our interest will lie in its truth. Philosophers alone, faced with a belief already established as true, ask the artificial: 'Is it knowledge?'

To conclude, the word 'knowledge' can be used in many senses which are subtly different from each other, and whose employment in language can frequently be explained by reference to the situation in which they occur. The quest for a precise analysis of knowledge, if taken as a fundamental issue in epistemology, would therefore better be replaced by an inquiry into the very possibility of validly inferring statements about the world from our given data. Knowledge concerns facts, and an investigation as to what are the facts will be more helpful than a futile dispute over the true 'essence' of knowledge. Scepticism presents us with a real challenge in this respect, and with a problem that is not, like so many, mysteriously born in the confused looking—glass world of language. To those which are, Humpty Dumpty has the answer.¹⁶

¹⁵ Some might also deny a subject knowledge on the grounds of little-known misleading evidence, where that evidence is likely to become available to the subject in the near future. Perhaps the motive behind such a denial would be the requirement that knowledge be durable.

¹⁶ I am very grateful to my mother, who kindly typed the script, and to Mr. J. M. Shorter, for his helpful comments on the question of unwitting knowledge.

REFERENCES

The following articles and books are referred to in the text. They are listed in the order in which they are mentioned.

- E. L. Gettier, 'Is Justified True Belief Knowledge?', *Analysis* 23.6 (June 1963).
- Carolyn Black, 'Knowledge Without Belief', *Analysis* 31.5 (April 1971).
- C. Radford, 'Knowledge – By Examples', *Analysis* 27.1. (October 1966).
- D. M. Armstrong, 'Does Knowledge Entail Belief?', *Proceedings of the Aristotelian Society* (1969-1970).
- C. Radford, 'Does Unwitting Knowledge Entail Unconscious Belief?', *Analysis* 30.3 (January 1970)
- O. R. Jones, 'Knowing and Guessing – By Examples', *Analysis* 32.1 (October 1971)
- K. Lehrer, 'Knowledge, Truth and Evidence', *Analysis* 25.5 (April 1965)
- K. Lehrer, *Knowledge* (Oxford: Clarendon Press, 1974), 189-226.
- M. Clark, 'Knowledge and Grounds: A Comment on Mr. Gettier's Paper', *Analysis* 24.2 (December 1963).
- E. Sosa, 'The Analysis of "Knowledge That P"', *Analysis* 25.1 (October 1964).
- A. J. Ayer, *The Problem of Knowledge* (Harmondsworth, Middlesex: Penguin, 1956), 31-35.
- J. L. Mackie, 'The Possibility of Innate Knowledge', *Proceedings of the Aristotelian Society* (1969-1970).
- G. Harman, 'Unger on Knowledge', *Journal of Philosophy* 64.12 (June 1967).
- E. Sosa, 'On Our Knowledge of Matters of Fact', *Mind* 83 (1974).
- Plato, *Meno*, trans. W.K.C. Guthrie (Harmondsworth, Middlesex: Penguin, 1956), 98a.
- G. Harman, *Thought* (Princeton University Press, 1973), 142-149.