PHILOSOPHICAL LECTURE

KNOWLEDGE, TRUTH, AND RELIABILITY

By SIMON BLACKBURN

Read 10 May 1984

THE philosophy of knowledge and truth is dominated by two metaphors: that of a system of elements which correspond with particular facts in the world, on the one hand, and that of a raft or boat of interconnected judgements, where no element corresponds to anything external, on the other hand. The choice is between realism, and a leaning towards a correspondence theory of truth, and holism, verging towards idealism, with a coherence theory of truth: the pyramid and the raft.[1] These images have complementary attractions, but neither provides a solid, stable metaphor whereby we can understand how truth connects us to the world and its facts. The first prompts the charge of foundationalism, or of the myth that we can step outside our best beliefs to estimate how well they correspond with the facts. And the second seems to disconnect the web of belief from proper control by the world, so falling into idealism, or relativism. In spite of rearguard actions it is the second image which dominates philosophy today.[2] I shall have little to say to oppose that movement of opinion. But if we follow it, does it leave our concept of knowledge where we would like it to be, or does it demand, as some have maintained, abandoning the concept as the remnant of a classical, but outmoded, self image?

Suppose we come to sympathize with the image of the raft through mistrust of 'the given', or through a Humean, or Wittgensteinian naturalism. Then it is possible to argue that we must retreat to a coherence theory of truth, and that this in turn gives a ready answer to scepticism. It brings truth down to the natural earth. We learn, on this account, that knowledge and truth are concepts which *we* use, in *our* world. Philosophies may have falsely

[1] Ernest Sosa, 'The Raft and The Pyramid', *Midwest Studies in Philosophy*, vol. v, ed. P. French *et al.*, Minnesota University Press, 1980.

[2] The image seems to me common to Quine, Goodman, Putnam, Rorty, and many others.

promised us a real correspondence, in which facts impinged upon us unfiltered by our own concepts, ways of classifying, or perception of similarities. But by seeing how false that promise is, we learn not to fall into scepticism when it is unfulfilled. We learn to feel comfortable claiming truth and knowledge within our own terms, and not to respect an alleged demand—which could be met only by an a priori argument—that those terms be given any foundation beyond the fact that we find them natural. I hesitate to ascribe this position to any one writer. Its characteristic combination of empirical (or internal) realism and transcendental idealism (or conventionalism) is, I should have thought, almost orthodox, and writers who oppose it are self-consciously fighting not just one author, but a whole tide of thought.

But somehow it doesn't seem to work as it should. For it is at least equally easy to feel that the combination destroys any right to regard ourselves as knowing—really knowing—what the world is like. This is obvious if the transcendental part—the part gestured at by thought experiments involving bent classifications, Goodman's predicates, or what are generally called the 'rule-following considerations'—issues in a kind of conventionalism.[1] But it is nearly as bad if it only issues in a kind of naturalism. The fear that nature—whatever it may be—has grown as not so much with a mirror as with a veil, or a distorting lens, is not easy to exorcize. So who is right: those who find a comfortable answer to scepticism in the combination, or those who fear that it plays into the hands of the sceptic? Or are we to suppose that part of the package is a new, appropriate, concept of knowledge, which itself supersedes any which permits scepticism to remain a real challenge?

My object in this paper is to approach these problems using a natural, everyday, requirement of reliability. I shall start by placing that in the relatively pedestrian context of the problems with the analysis of knowledge, on which there have been so many recent assaults. The position I arrive at is a version of 'reliabilism', but one which ought not to be opposed by theories of knowledge which insist upon justification. The reason for this combination emerges in due course. The position affords an argument against scepticism, but it would be idle to pretend that it 'refutes' it: indeed it offers a diagnosis of the permanent appeal of sceptical thoughts. This diagnosis does not depend upon

[1] This threat is discussed in my *Spreading the Word*, Oxford University Press, 1984, ch. 7.7.

a profound internal versus external reading of the sceptical concept of knowledge.[1] It sees the sceptic, certainly, as introducing a new context of inquiry, but it offers no straightforward way of dismissing that context either as illegitimate, or as involving new and different concepts to any in everyday use.

So I start with some observations about the concept—our concept—of knowledge. These ought to help us to see just what that concept involves, and therefore to understand how much survives the drift towards idealism.

I

The classical problem is to find the condition which adds to

p is true and
x believes p

to give sufficient conditions for: x knows p. The standard suggestions include refinements of the requirement that x be justified, refinements of the requirement that x be situated reliably with respect to the fact that p, and versions of the requirement that x's belief be not defeasible, meaning that further evidence ought merely to confirm his belief that p.[2] The chase for more accurate versions of these conditions, and the rivalry that can develop between them, has been called Gettier's salt-mine, and it can enslave us even against our will. So we might start by asking why we should need an extra condition in the first place.

If the epistemic concepts earn an honest living, they must form a natural intellectual kind. Even if some multi-part analysis accurately matched our judgements in difficult cases, it would still need asking why we are interested in just *that* set of conditions (a similar question arises when we propose complex psychological conditions for meaning, and in many other areas). Seeing ourselves and each other as knowing things is to be important. But how can it be important to organize our lives around one complex of conditions rather than another? We need a role for the epistemic concepts, and the role which seems most natural is that of ranking and selecting titles to respect. We have to pick up our beliefs about the world from our senses and from each other. So we need a

[1] Barry Stroud, *The Philosophical Significance of Scepticism*, Oxford University Press, 1984, especially chs. 3, 4, and 5. The distinction of course derives from Kant and Carnap.

[2] K. Lehrer and T. Paxson, 'Knowledge: Undefeated Justified True Belief', in *Essays on Knowledge and Justification*, ed. G. Pappas and M. Swain, Cornell University Press, 1978.

vocabulary to settle whether our sources are ones which themselves properly indicate the truth. This is a natural need, and it gives us the natural intellectual kind in which to place our epistemic verdicts.

So consider a subject who believes correctly that something is so. His being right gives him one title to correctness. Why isn't this enough? Because his position may not deserve respect as the *kind* of position from which one may safely accept information. A subject may believe truly by exercising defective propensities to form belief on occasions on which, by luck, he is right, or by exercising proper propensities, but when it is not they which are responsible for his being right, but an admixture of luck. Given this reply, the two concepts which are anathema to each other are *knowing* and being in an unreliable, defective state, or using an unreliable propensity to form belief (the close analogy, of course, is with the agent who does the best thing by accident, but has not exercised virtue in doing so). It is natural to detect two components in a subject's epistemic virtue on an occasion. There is the amount of information at his disposal, which may be more or less adequate, and there is what he makes of it, which may involve more or less rationality, or more or less reliable propensities to use information to deliver belief. These two components need not march in step, of course. But for the moment the difficulties this could cause, and indeed the difficulties of effecting the division in any accurate way, need not concern us.

Let us say that someone in a certain state of information, and exhibiting some disposition to form belief, is also showing a degree of soundness, or solidity as a source of information. We can call this a degree of value to a would be information-receiver, or IRV (information-receiver value). If this is the normative dimension, as it were, in which to place knowledge, then we would expect to be able to put the following principle down:

The Mirv/Pirv principle: If two subjects each believe truly that p, the one cannot know, when the other does not, unless the former is in a position with at least as much IRV as the latter.

Since the role of the epistemic concepts is to rank sources of information, then if one source knows, when another does not, it cannot be that the belief of the knowing subject is unsafe in ways which give him less IRV than the subject who does know. This is a principle concerning belief. So there is a caveat to enter in the use of this principle: we might call it the Matilda caveat. Matilda told such dreadful lies that when she eventually shouted that the

house was on fire, nobody believed her.[1] Her *effective* IRV had disappeared with her credibility, but for all that, she knew that the house was on fire. To use the Mirv/Pirv principle properly, we must ignore differences between safety of report and safety of original belief: we should say that Matilda's report actually *had* IRV, because the belief to which it gave voice was solid, even if rational hearers might have doubted it. Ultimately, of course, we are to be concerned with problems of our own reliability, and problems of insincerity in report, or of difficulties of interpretation of our own language, do not arise.

The Mirv/Pirv principle comes initially as a constraint upon the missing clause in the proposals for defining knowledge, and I suggest that it guides many verdicts in contested cases. To give a simple example, consider the subject who forms a true belief well enough, but who should have done something else as well, albeit that the extra thing would in fact have misled him (he believes, rightly and reliably, that the president has been assassinated, but others who did believe this have by now read the usually reliable morning papers, which deny it . . .). If we are reluctant to describe him as the only person who knows that the president has been assassinated, this is because someone who has done the extra has done the kind of thing which makes them a better source of information on this kind of issue.

The principle serves to rule out even powerful and plausible attempts to analyse knowledge. More importantly, I suggest that it explains our unease with these attempts: our sense that somewhere things are going to go wrong for them. Consider, for example, the conditional analysis of Dretske and Nozick.[2] This finds the missing clause in the two conditionals:

> If p then x believes that p
> If −p, then it is not the case that x would believe that p.

The idea is that x's believing should be sensitive to the truth, so that x should be what Nozick felicitously calls 'tracking' the truth. This idea is a good one: sensitivity to truth is indeed the kind of solidity we are looking for. But its realization in the two conditionals is not so good. For a little thought will show that a person could satisfy them through possession of a defect, compared with someone else who does not satisfy them, and that, for some audiences, this defect could make him a worse informant on the

[1] Hilaire Belloc, *Cautionary Tales for Children*, Puffin, 1950.

[2] F. Dretske, 'Conclusive Reason', *Australasian Journal of Philosophy*, 1971; R. Nozick, *Philosophical Explanations*, Oxford University Press, 1981, ch. 3.

kind of case in point. I shall illustrate this by a case, but it is the principle that matters.

Two freshmen, Mirv and Pirv, see the Professor in a car. They each believe, truly, that the Professor is in his own car, and they are each good at telling, in general, when propositions like this are true. Usually, for instance, when the Professor is not in his own car he drives very insecurely, and each freshman would judge that he was in an unfamiliar car. On this occasion, however, the Professor might not have been in his car, which was due for a service, and had he not been, the garage would have lent him a model of the same type, which would have been familiar to him. The only difference is that the garage model has a sticker of Mickey Mouse on the back, but the Professor wouldn't have minded that—indeed, he used to have such a sticker himself, let us say. However, Pirv comes from a puritanical and benighted part of the country, and could not bring himself to believe that anyone as distinguished as a Professor would ever own a vehicle with such a sticker. Mirv knows more about the world. But on this occasion his knowledge stands him in bad stead, by Nozick's lights. For through it he fails to satisfy the fourth condition: if the Professor had not been in his own car, Mirv would have continued to believe that he was. Whereas Pirv, through ignorance and misinformation, ends up satisfying the fourth condition: had the Professor not been in his own car, Pirv would not have believed that he was. So Dretske–Nozick would have us saying that Pirv knew the Professor to be in his own car, whereas Mirv did not. This flouts the principle, for Mirv is a better source of information about such things than Pirv. He is a better tuned car ownership detector, using the right parts of a better system of belief about such matters.

To say that Mirv is more solid on this *kind* of issue raises the question, noticed by Goldman,[1] of how we should classify 'kinds' of proposition, in order to evaluate the reliability in informants (for reliability is inevitably reliability in a *kind* of circumstance). And it raises the question of the antecedent position of the receiver of information. Someone who knows much more about a situation may rightly take information from a source who is generally worse, or who on an occasion is behaving quite irrationally, just because he knows that for particular reasons obtaining in this case, the irrationality is not involved in the informant's situation.

Compare acting as a second to a careless rock-climber, who ties the belaying knots in such a way that they might be safe, or they

[1] Alvin Goldman, 'What is Justified Belief', in *Justification and Knowledge*, ed. G. Pappas, Reidel, 1979.

might not, and does not check the difference. Was the weekend safe? God might have said so: perhaps all the knots the leader tied that weekend were luckily secure. His defect made no difference to your actual security. Swayed by this, we could say that, ontologically as it were, the weekend was really safe. But you mightn't think so: there is a good sense in which you cannot ever be safe behind such a person: you oughtn't to feel safe just because you don't know, any more than the leader does, what he is actually doing (suppose his defect only comes to light after the weekend is over: breaking out in a cold sweat, you correctly say 'what a dreadful risk I took!').

The issues here are close to those that arise in any application of statistical or dispositional facts to the singular case. Is it safe to bet on Fred surviving to the age of eighty? He is a sedentary, bran-eating, slender academic . . . We rapidly come to the narrowest class with weighty statistics, yet there is no end to Fred's peculiar combination of properties. Suppose he does survive: it does not follow that it was safe to bet on it. It may have been safe for God to bet on it, just as it is safe for him to follow the unsafe leader, or safe for him to ignore an exercise of irrationality on occasions when it is not in fact affecting the truth of his belief. But it would not have been safe for us.

In the case of chance, we suppose that the weightier the reference class, the better: we say that when we know more, so narrowing the kind in which to put the single case, we have a better estimate, or are nearer to the 'true' probability. This is easy to explain in pragmatic terms: someone using the fuller information wins when betting with someone who can use only the lesser.[1] But because the standard epistemic position is not one in which the receiver is the more knowledgeable party, we do not tailor the epistemic verdict to cases in the same way. We are not, as it were, concerned with how God might pick up information from a source: we are concerned with how we might. Thus we take into account causal factors which render a source more or less sound. But if flaws are involved it is *not* the weightier position, which happens to know that they are not responsible for the truth of the informant's belief, which counts. There are cases in the literature which, in effect, trade on this problem. Suppose, for instance, that Pirv is told by the President of the Royal Society that the dark room he is about to enter contains a perfect holographic illusion of a vase. Suppose that, irrationally, he takes no notice and

[1] I detail this in 'Opinions and Chances', in *Prospects for Pragmatism*, ed. D. H. Mellor, Cambridge University Press, 1980.

believes because of a cursory glance that there is a vase there anyhow. He doesn't know that there is a vase there, even if there was (for the President was lying, deceived, or just failing to remember that the machine was off). God, or anyone knowing this much more, could safely accept Pirv's word that there is a vase there, because they know enough to discount the exercise of irrationality in the way he came by his belief on this occasion. But someone knowing no more and no less than Pirv does could not accept his word. After the operations described, Pirv is not a solid source for him on the matter. Since the normal epistemic circumstance—the one which makes channels of information important—is that of wanting to know whether to accept information from a better placed source, it is not these superior positions which count. It is not so that he knew there was a vase on this occasion, because there is a kind of thing he is doing—forming beliefs irrationally—and it is dangerous to accept beliefs when this kind of thing is done. Someone knowing more can say that there is a narrower kind of thing he is doing—forming beliefs irrationally when the facts are such that the irrationality does not matter—and that it *is* safe to accept beliefs when this is true. But this superior epistemic position does not dictate our verdict: the concept of knowledge would lose its utility if it did (although there is a telling temptation to go soft on this: if someone's stoutly maintained, but irrational, belief turns out to be true we sometimes let this success alone dictate the epistemic verdict: 'funny how Beryl knew all along that Fred was . . .' The success makes us think that there must have been a reliable kind of belief formation involved, even if it would have taken a superior being to know what it was. Compare: 'so it *was* a safe bet after all', which is usually said when it wasn't.)

Puzzle cases and disputes arise because there are different ways of classing the kinds of case in which someone is reliable: there is the question of whether the informant is in some normal causal relationship to the facts; whether he is reliable over similar kinds of case; whether he would be justified or rational in believing himself to be reliable, and finally whether his background beliefs (which in turn may be rational or not) affect his standing as an informant. And all these can come apart. If any of them fails, then there will be a way of regarding the informant which makes him into a dangerous source of belief: there will be a kind of case over which he does badly. But it would be optimistic to expect principles to settle verdicts in such cases. Because of the 'holism' of belief, there is no principled limit to the flaws which may result in our being in kinds of state which are unreliable, and disqualify us

from acceptability as a source. In particular we will always be vulnerable, if we try to isolate some natural relation which a subject has to the fact that p, to cases in which he is nevertheless playing Pirv to someone else's Mirv, although the other person does not bear this particular relation to the fact. This explains the progress found constantly in the literature: someone proposes a natural relation to the facts sufficient for a subject to know that p, and someone comes along in a generally better state (he has read the newspapers, etc.) but who through the extra virtue, making him more solid on some *kind* of case, misses the title.

II

Solidity as an authority is a matter of degree. But knowledge, on the face of it is not. So how much solidity do we want: can knowledge tolerate chances of being wrong, or even the bare possibility of being wrong? To put the question in a closely related way, if a situation leaves it as much as barely or logically possible that one is playing Pirv to a non-knowing Mirv, does that destroy one's title? The most important initial division in the continuum of possible improvements comes where a subject is sufficiently solid to be an authority, and where any improvement in his state, or dispositions, would simply serve to sustain his belief. We could relativize this, if the possibility of different recipients with different standpoints is worrying, and say that a recipient should allow someone to know something just when anything which from his standpoint counts as an improvement, merely tends to confirm the original belief. This suggestion is of course close to the familiar non-defeasibility condition on knowledge. It differs only in that I put the notion of an improvement to the fore: it does not go without saying that increases in true belief, even when reasonably used, count as improvements on a particular kind of case. There is a caution too implied in putting the question as one of whether improvements *would* sustain the verdict. It will not be to the point to go in for thought experiments where improvements which could overthrow the belief, could happen, but in the actual world wouldn't. This kind of stability is sometimes easily achieved. Suppose I recognize my friend by a glance at his face. I know who he is, not because weightier investigations could not be made, but because they would simply confirm what I already know. Of course, anyone whose position is as solid as this cannot play the role of Pirv to someone else's Mirv: anyone in a genuinely improved position will also know.

If we used this as a cut-off point beyond which there is knowledge, we would be importing what Armstrong called an 'external' element into the notion of knowledge.[1] I could be in an informational state, and using dispositions sufficiently well, yet not know something because, as a matter of fact, the world does afford further evidence which would undermine proper confidence in the belief. And I could be in the same state and using the same intellectual dispositions, when on the contrary, any improvement would confirm my belief, and in that case, on this proposal, I know. People are uneasy with this for several reasons: notably, it seems to cut the concept off from any problems of objective justification, and it affords altogether too cheap a victory over scepticism (provided the way we are plugged into the world *is* alright, then we know, and the sceptic cannot show that it is *not* alright).

These worries may lead us to divide the continuum higher up. At the highest point, it is logically impossible that the subject should be playing Pirv to another's Mirv. The gap between the subject's informational state and the fact believed to obtain is to be closed altogether: it is to be logically impossible that the state should exist, yet the fact not obtain. This exorcizes all external elements with a vengeance: it tries to ensure that there is no element of luck, or even contingency, in the true believer's title to knowledge. Traditionally it requires that we shrink the area of fact known, potentially down to an entirely subjective realm, just as the parallel motivation in the theory of ethics shrinks the exercise of real virtue down from the chancy, external world where good intentions can go wrong, to the safe realm of acts of will. Alternatively, we might close the gap by expanding our conception of the state the believer has got himself into. The states we get into, and because of which we form beliefs, would be ones which we could not (logically) have been in had there not been a spatially extended, temporally ordered world, containing the other minds, numbers, possibilities, values, etc. in which we all believe.

What then is our best conception of the informational states whereby we come to believe things, or to know them? Let us say that informational states, in virtue of which we form beliefs, divide into two. There are those which, as a matter of necessity, could not have existed had not the beliefs formed in the light of them been true. We can call these guaranteeing states. And there are those which do not meet this strong condition. Call these indicative

[1] D. M. Armstrong, *Belief, Truth and Knowledge*, Cambridge University Press, 1973, p. 157.

states. The question in front of us is whether only guaranteeing states sustain knowledge, or whether indicative states, provided the external circumstances are right, can also do so. If we can happily see ourselves as largely possessed of guaranteeing states, the looming problems of scepticism might be thought to disappear. But can we? The 'informational state' in virtue of which a system is disposed to absorb something new can include any part of the deposit of previous times, as well as anything which could at all be thought of in terms of the impact of the immediate environment. We think of ourselves, of course, as getting into such states as a result of our physical positions and surroundings, the operations of our senses, and the use of concepts, beliefs and expectations which, be it because of reason (unlikely) or nature (most likely) or convention (let us hope not), we find ourselves forming. These banalities do nothing to support a 'guaranteeing' conception of informational states. On the contrary, they conjure up painful images of the ways in which the world responsible for our states might not conform to the way we end up taking it to be. (To adapt a metaphor of Kant's: a system of knowledge is a slow growth, like that of a crystal in a liquid. Its structure and strength are its own, and even if its composition is entirely determined by the matrix, nevertheless it need not reflect it.)

A guaranteeing conception of our epistemic positions is given spurious support by a spatial metaphor (Kant charged that it was a mistake of Locke to sensualize the understanding.[1] I think it is at least as important a mistake, and symptomatic of the same error, to spatialize it.) Thus we are often asked to pronounce upon what is manifest, disclosed, given, embraced, internal to our subjectivity, accessible *in* our experience, or to settle issues of what we really confront, or access, or what we can penetrate to, or what is transparent or open to us. The glassy blob of the mind reaches out to encompass (embrace, contain) facts, and knowledge stops at its boundaries. But then the blob cannot stop short of embracing all kinds of states of the world, for if it did it would be confined to embracing mental proxies of them, and these would so intervene that it could never know the world, nor even understand a vocabulary purporting to describe it. Whole issues, such as the realist/anti-realist confrontation as it is framed by Dummett and his commentators, are importantly distorted (or sustained) by this spatial metaphor.[2]

[1] I. Kant, *The Critique of Pure Reason*, A271 = B327.

[2] J. McDowell, 'Criteria, Defeasibility and Knowledge', *Proceedings of the British Academy*, lxviii, 1982, 455–79. The damage of the spatial metaphor is seen

To escape this error we might query whether the very notion of a 'state' plays us false here. Because of the interpenetration of theory and experience, and because of the temporal growth of the system of belief, it can seem artificial to analyse a response to new experience by thinking of informational states at all. Certainly, if we do, there is little better to say about them than that our state is one of being in an external world, surrounded by other minds, possessing a long past, and so on. The promise of a quick victory over the sceptic appears again, for our basic characterization of ourselves still entails that we live in the kind of world which he finds it possible to doubt. Unfortunately, we cannot retreat into dogmatism so comfortably. It will always seem a fragile response to scepticism to refuse to set the problem up in the first place—little better than announcing that the mind embraces the relevant facts after all. So perhaps the best thing *is* to get away from the spatial image as radically as possible, and this includes avoiding the Protean notion of an informational state. Although I sympathize with this, the notion of a state does not have to be taken spatially, and there is no better general term to sum up the fact that at given times we are in positions (states) in which we form beliefs, and that the ways in which we do this, and their strengths and weaknesses merit investigation.

The lowest place, as it were, at which we could claim knowledge was where our state was what I called 'sufficiently' authoritative, meaning that it made a reliable source on the kind of matter in question, and where it was actually stable. This is a possible resting point: it depends, I think, on whether we read the condition as strong enough to mean that there is no chance, or virtually no chance, or only a chance that can be dismissed, of our being wrong. Read without that understanding, the condition that belief be true, authoritative and stable would be far too weak for

explicitly in McDowell's argument against his opponents: he believes that if the mind does not embrace past states of affairs, the sensations of others, and so on, then it must embrace only proxies of them 'interposing' between us and them (pp. 472–4), giving rise to insuperable problems of understanding and knowledge. In the theory of thought this is the analogy of the position that we either see physical objects 'directly', or we see proxies of them—sense data. Austin attacks this dichotomy at the beginning of *Sense and Sensibility*: 'In philosophy it is often a good policy, when one member of a putative pair falls under suspicion, to view the more innocuous seeming party suspiciously as well.' (p. 4). Dummett's 'challenge' to realists, to explain how things which are not 'manifest' can be understood, and which is met by McDowell by the strategy of making more and more of the world 'manifest', seems to me to be much better met by entirely refusing the terms of discussion.

knowledge. For it could coexist with a good chance of being wrong. But where we have a good chance of being wrong, we are in a kind of state which makes us unsafe sources of information. So we must read the condition so that it excludes any significant chance of being wrong. And this is the point which I want to focus upon in what follows.

If we said that knowledge can exist provided there is no significant chance, or real chance, of error then we can defend a title to knowledge in the face of an open, acknowledged, possibility that the world might not be as we have come to take it to be. The sceptic is apt to complain that when this is all we have, then for all we know, things are not as we take them to be. But this is wrong, for the whole issue is whether on the contrary we can know something through being in a state which is indicative, although not guaranteeing, provided the external condition is satisfied, that we are authoritative, and the state is stable. My suggestion is that the sceptic gets away with this bare citing of possibility, because of the normal implicature that we cite a possibility only if we also give it some chance of being realized. It is normally only to the point to cite possibilities which are 'relevant', and this is exactly what relevance is. So it can seem that mere possibility left open defeats knowledge, whereas in fact it may be that it doesn't, but that only real chance of error does. The externalist has it that we know because we are right, and because any improvement in our position would just confirm that we are, and because we exercise sufficient soundness to be a proper source of information on such a matter. Once this is so, the sceptical possibility can, as we naturally say, be ignored. It is the relation of this position to scepticism which I now wish to expose. For I hold that, although it may seem to cheapen knowledge, in fact it does considerable justice to sceptical doubt: it offers an explanation of the deep roots of those doubts, and it may enable us to place them even within the context of a general sympathy with 'anti-realism' or 'internalism'.

III

Reliabilists and justificationists think of themselves as forming two different camps. Now one element in the view I have been defending supports each of them. Reliabilists appear right, in so far as the soundness we require of informants need not imply any self-consciousness on their part. They could be like good instruments, and be deemed to know things just by being rightly tuned to the truth. But although an informant need not have views about his

own reliability, we need to do so. It is always a weakness to have no account of why an informant should be thought to be yielding the truth. It generates a bad *kind* of state to be in. And when our own title is in doubt, externalism does not help us unless we can properly see ourselves as reliable. To put the matter in terms of section *I*, when we are unable to see ourselves as forming belief reliably, but nevertheless form it all the same, we are doing a kind of thing which destroys our title as authorities. We cannot suppose that the mere fact of our being right removes this taint, any more than the man in the Royal Society case escapes the charge of irrationality or gains the title of knowledge, just because on that occasion his belief was true.

Now the power of scepticism is quite underrated, if it is seen as merely a forlorn attempt to shake confidence by invoking possibilities which can normally be ignored. Its real power comes with the absence of any sense of our own reliability. Crucially, we have a sense of there being a large number of possible worlds which appear as ours has done, but which contain scientific realities unlike ours, skew distributions of other minds, large elements of counter-inductive truth, and so on. We might try to say, blankly, that we know that these possibilities are not realized. But we have to be able to regard ourselves as reliable on just this *kind* of point. How can we? How could I have a better than chance propensity to tell when sceptical possibilities are realized? I can do nothing more than rehearse the very considerations governing belief; if they leave open a space of possibilities, then there is nothing more to say about which possibility is realized, and nothing more than chance to determine whether I am right. There is, for example, only one kind of world in which other minds distribute as I naturally take them to do, but there are lots where they distribute in other, partial, ways (no other minds, ones attaching only to . . . etc.). There is only one kind of world which is well-behaved with respect to my inductive regularities, but there are many which deviate in their different ways. If evidence leaves the possibility of such distributions, how can I be better than chance at telling when they are realized?

I think it is wrong, or at least misleading, to suggest, as Barry Stroud does, that scepticism here involves taking an 'external' view of our knowledge, as opposed to an 'internal' one in which such questions do not arise.[1] At least, this is wrong if it leaves open a ready way of suggesting that the external standpoint is optional, or even that it makes no sense. And the usual metaphor

[1] Stroud, op. cit., ch. 4.

of externality is dangerous in this respect. All that really happens is that normal ('internal') assessments of knowledge go on against a background of assumptions of general reliability—the generally truth-yielding nature of our procedures. But the *same* demand that there is no chance of error can be made of the procedures, even if it is only in philosophical moments that we think of raising it. Thus when Stroud diagnoses Moore as unable to hear the philosophical sceptic's question in the intended way, it is unnecessary, on my view, to suppose that there is a special, transcendental, inquiry or context which Moore cannot enter. Rather, there is a univocal query—about the chance of being wrong—which is normally answered against a background of common-sense theory (and is so answered by Moore), but which can equally be raised about the procedures used in creating and sustaining that theory, or the principles upon which it seems to depend (induction, trust in the senses, etc.). I suggest that this better explains Moore's peculiarity, which is that he seems blind to the point at which any grasp of our own reliability fails us. It also gives us reason to be cautious with the metaphor of externality, for it is not as if the philosophical undertaking demands quite different tools or perspectives from the everyday assessments of chance: it just has a different topic. Similarly, our everyday financial standing may be settled by considering the credit we have at the bank; this does not rule out a sensible query about the financial standing of the bank itself.

To show the query to be improper, in the philosophical case, we would need to show that the relevant notion of 'chance' is inapplicable, when we consider the chance of the sceptical possibility being realized. Unless this is done, an airy assertion that there is no chance of things being like that will sound quite unsupportable, and the sceptic wins. There is only one way that I can see of respecting the possibilities, but avoiding scepticism, and that is to improve the theory of truth, for modal assertions and for assertions of chance. We have to say that although there is a real space of possibilities, as the sceptic maintains, there is also no chance that any of them are realized: it is known that we are not unreliable. Are there doctrines in the theory of truth which enable us to say this?

IV

The sovereign proposal is to think of truth as some kind of construct out of our conception of the virtues of methods of inquiry, and the consequences to which they lead. 'Realism', in

at least one good use of the term in this connexion, thinks that we can explain the virtues of method by certifying that they are midwives to truth; 'anti-realism' sees truth as that which ought to be established, or would be established, by the best use of the best methods. The one philosophy sees the virtues of right reasoning as a precipitate from an antecedent notion of truth, and the other reverses the priority.

It is often suggested that the anti-realist direction makes for an easy dismissal of scepticism.[1] For instance, it is supposed that on the philosophy of the later Wittgenstein, our procedures and practices of ascribing pain to other people, together perhaps with the consequences we attach to such a description, determine what we mean by it. This leaves it open, it is supposed, that such ascriptions are defeasible, so that any finite evidence for the ascription can lead us to be wrong. But it is then supposed that the priority of assertibility conditions forbids us from making sense of the sceptical possibility that the world contains no consciousnesses but mine (or those of some favoured sub-group including me): stoicism, pretence, and so forth can only exist against a background of general correctness, and this correctness is supposedly guaranteed by the criterial, practice-governed conception of meaning. I find this obscure. The practice of attributing mental states to others leaves open the possibility of error in the face of finite behavioural evidence. If it leaves this possibility in each case, then even if it does not follow that it may do so in every case, still we must ask why it does *not* do so in the conjunction of individual cases—that is, as regards the world in general. The concept of virtue attached to such ascriptions may leave us quite unable to reject this bare possibility. Rejecting possibilities of error may be no part of the practice, and not entailed by the virtues or ways of reasoning which are integral to the practice.

There is another way of raising this problem. Once more, suppose we sympathize with the anti-realist priorities. Then I might be confident that the best possible system of belief about other people, formed by the most virtuous dispositions, should contain the belief that others see colours as I do. But it might *also* contain the proposition that it is possible, in spite of any of the evidence I have or could have, that they do not. It would contain this proposition if the idealized increases in information or virtue do not rule it out, or in other words, if even a supremely virtuous

[1] For a typical assessment see Colin McGinn, 'An a priori Argument for Realism', *Journal of Philosophy*, 1979. But see K. Winkler, 'Scepticism and Anti-Realism', *Mind*, 1985.

cognitive agent, using information as admirably as possible, should still allow or respect that possibility. And perhaps he should, for perhaps none of the increases in virtue which we can imagine to ourselves would ever lead to a proper refusal to countenance it. And in that case, it will be correct to allow it—on the anti-realist's own construction, it exists. In the same way it would be correct to allow Horatio's possibility, that there may be truths which would lie for ever outside our comprehension. It would be the part of virtue to admit as much. Even God would be right to doubt the guaranteed nature of what is, nevertheless, his own knowledge; it would be part of the 'final best science' to not regard itself as the final best science (this is a truth which the final best science could not acknowledge, disproving that definition of truth, or alternatively disproving the existence of both a final best science, and truth).

It will be evident then that I differ from both Putnam and Dummett, in holding that the question of priority does not coincide with the issue of whether we can understand 'verification transcendent' truth-conditions, or in other words, allow sceptical possibilities. To that both my sides reply that we do so in so far as our practices contain as a legitimate element an enterprise of wondering whether, in the largest respects, the world is as we take it to be. This result accords with what I call 'quasi-realism', for it is another respect in which someone who approves of the anti-realist instinct over the priority of truth or virtue, yet ends up with the very thoughts that the realist took for his own. In turn the suggestion casts doubt on whether we really have an issue between global anti-realism, and global realism, for if each side ought to end up following the same practice, there may be nothing to dispute over. But I would urge that sometimes, in local areas, we can make sense of the divergence of priorities, and even award the victory to one side. For instance, I believe that in the theory of morality, or modality, or chance, there is an advantage to the side which starts off by regarding method as fundamental. Moral and modal propositions, and most notably for present purposes, those about chance gain their identity, and the identity of the concept of truth to associate with them, from their place in a two-sided practice—that of coming to them, on the one hand, and of using them to guide the conduct of life and thought on the other. It is therefore particularly attractive not to try to explain their role by postulating an antecedent notion of truth to which they answer—a layer of facts about distributions of possible worlds, or of chances over them—but rather to explain what is to count as truth in their

case by thinking of what it is for them to perform their role successfully. One might try to say that this is always the case, so that if these propositions are better understood this way, any proposition would be. But this does not follow. These propositions may find their place in steering us around the facts, as it were, rather than in describing new layers of fact, but it would not be possible for all propositions to be like that. And these propositions (and those of mathematics) share peculiarities that contrast them with others from the outset, and which make the notion of truth so problematic in their case. They have no recognized epistemology, and their truth is not the starting point of any serious explanatory theory of our experience. So they do not serve as an attractive model for a general debate.

So far, things look even better for scepticism (again, remember that this is in spite of the relatively weak account of knowledge I am offering). The sceptic is not silenced by the highly abstract changeover from, say, 'metaphysical' to 'internal' realism (anyone who has ever taken the problem of induction seriously may wonder why he should have been thought to be put out by such a change). But we are owed an account of the relevant assessment of chances, and here there may be more scope for a response. Suppose we put some fledgling, anti-realist, thoughts about the truth about chance alongside the position in which we now find scepticism. Hume denied, rightly, that any durable good can come of extreme, or Cartesian, scepticism. But that is not at all the same as denying that durable good comes from disallowing the sceptical possibility. We have already urged that virtue may involve respecting general possibilities of error. But we have not yet seen that the virtuous method of forming belief about our reliability should leave us any sense of a real chance that those possibilities are realized. Can we hold the line against scepticism at just this point?

Scepticism invites us to 'stand back' and think of a logical space of possibilities, many of which accord with our evidence, but only one of which accords with the way we take actuality to be. When we do this we are apt to think that there is a real probability measure, meaning that some such possibilities have a better chance of realization than others. On an anti-realist line about chances, matters are the other way round: proper confidence itself determines what we are to say about any such measure. The ordinary considerations in favour of induction, other minds, and so on, have to give us a title to say that there is simply no chance (or, dismissably small chance) of things not being as we take them to be.

Saying that there is no such chance will sound like saying that we know that a particular ticket will not win a lottery—something which is usually false, since we can have no authority on whether such an event will happen. But it is not like that. For there we have a kind of thing that does happen (individual tickets win) and our reliability over whether it is going to happen in a particular case cannot be better than chance. Here we know the chances—they supervene upon natural facts in our world. But nobody has any right to say that massive undetectable facts according with sceptical suggestions are the kinds of thing that happen: we are not flying in the face of an actual empirical kind with a given frequency of realization when we deny them any chance at all.

There is now an opening for the sceptic to ascend a level. What I am doing, he will say, is denying that there is a real 'trans-world' probability metric, giving his possibilities a real chance of being actual. I am saying that chances are properly to be evaluated in the actual world, in which, I suppose, things like massive undetectable failures in the mentality of others, or failures of induction or memory, do not happen. But, he will complain, this is not sufficient. Suppose that the idea of trans-world chances is incoherent. Perhaps if we knew enough about the world, we could also say that there was no chance of bizarre possibilities being realized: *contingently* the chance of his possibilities is actually zero. But his claim is precisely that we do not know this much about the world: for all we know about contingent reality, the chances of (say) there being no other minds, or of the world conforming to Goodmanesque bent predictions, is quite high. In other words, denying a trans-world metric on chance is not enough, for it still leaves us ignorant of the distribution of chances at our actual world. And, the sceptic continues, the chances may be pretty unfavourable.

I think this admits of no refutation, for it depends entirely upon who has the onus of proof. If the sceptic's task were to prove that there is a chance of massive falsity in common-sense beliefs, then he fails. We can maintain, in one breath, both our normal beliefs and the corresponding title to knowledge, since there is no chance of their being false. The sceptic cannot dislodge us, since he cannot prove the existence of the disturbing, knowledge-defeating chances. On the other hand, we cannot prove against him that the relevant chances are zero—not without helping ourselves to the very contingent knowledge that he wishes to deny us.

Perhaps there is more to be said here along these lines. I have been emphasizing the continuity between discussion of scepticism

and everyday discussion of chance, albeit that the latter is heavily constrained, empirically, in a way that the former is not. Now someone might urge that this continuity is spurious: it hides the crucial difference that in the ordinary case we have *procedures* for assessing chances, but all of these, as we have seen, can be disallowed by the sceptic. Can this be used to urge that the position, and the debate, is indeed taking words out of the contexts in which alone they make sense, so restoring suspicion of the whole issue? I do not think so. For example, wrestling inductive procedures into an anti-sceptical shape is a perfectly recognizable activity, bound by the same rules of chance, the same kinds of argument, that hold sway elsewhere. The trouble is, that it tends not to issue in anything that impresses the sceptic.

So all we are left with to urge against him is a refusal to accept the onus. Allowing chances, like allowing possibilities, is something that we do. Sometimes, when the actual world affords stable frequencies, doing it well is heavily constrained by natural facts. But when we think of the chance of the world being as sceptical thoughts suggest it might be, we are not so constrained. Then, we are only aiming to express the proper, best way in which confidence is to follow on from the use of the ordinary evidence in favour of common-sense beliefs. So we can properly allow the bare sceptical possibility, and properly disallow any chance whatsoever that it is realized. Durable good comes from both policies. Neither flies in the face of a real, trans-world distribution of possibilities or chances, for there is no such thing.

V

This may not be a particularly glorious victory over the sceptic. I do not mind this—indeed, like Stroud I would mind more if the victory had been gained by the kind of dismissal which refuses to acknowledge the deep legitimacy of sceptical worries. The deep roots of scepticism lie in the need to see ourselves as reliable over as many matters as possible. It is not the use of inappropriate standards, nor shifting to a different and doubtful external point of view, nor yet accepting an unbearably strong cut-off point for knowledge, which leads us to focus upon our most general methods and ask for their reliability. It is, as we might put it, not trying to hover above our boat with a new and unfamiliar set of a priori instruments of inspection; it is just using the same instruments on the same boat, but on a little-visited and basic part of its structure. Unless this is seen scepticism will not have had an

adequate answer. For that very reason the problem of knowledge, as we have inherited it, or the very subject of epistemology, should not be seen as the parochial, historical outcome of mistaken conceptions of mind or experience. It may produce stunted side-shoots because of such mistakes, but when they are lopped off the problem of knowledge remains. For, given problems such as those of observation and induction, we have no stable way of imagining a body of knowledge, which protects all the exposed surfaces where the title to reliability is vulnerable, and needs questioning, protecting, reconceiving.

Wittgenstein imagined that the philosopher was like a therapist whose task was to put problems finally to rest, and to cure us of being bewitched by them. So we are told to stop, to shut off lines of inquiry, not to find things puzzling nor to seek explanations. This is intellectual suicide. If the philosopher is indeed like a therapist, then his task is to insist upon constant exercise: the inspection of the bindings, the exposed surfaces, the possibilities and chances, the dangerous places where a sense of our own reliability takes no place in the rest of our scheme of things.

Note. I owe thanks to E. J. Craig, and Lindsay Judson, for conversation and comments on these themes.