

1 Scepticism

1.1 SOME DISTINCTIONS

Scepticism in its most interesting form always depends on an argument; the better the argument, the stronger the scepticism it generates. Since it depends on an argument, it must be able to be expressed as a conclusion. The sceptical conclusion is that knowledge is impossible. No one *does* know, because no one *can* know.

There is a contrast between the sceptic who offers an argument which has this as its conclusion and two other sorts of sceptic. The first is the person who responds to every assertion with the question 'How do you know that?' and then, whatever is offered, merely repeats the question until answers dry up. This repeated question is very successful in reducing others to furious but impotent silence, but there is very little to be learnt from it as it stands until we know what lies behind it. There are of course interesting possibilities here, for instance the following propositions:

- 1 No one knows that *p* unless he can say how he knows that *p*.
- 2 The attempt to answer the question 'How do you know that *p*?' by simply reasserting that *p* cannot be successful. It begs the question.

A sceptic who peddles the question without being willing to make an appeal to propositions such as these is not presenting an interesting philosophical position. Once the appeal is made, however, we are back with a scepticism that depends on an argument. It is worth pointing out in this connection that the two propositions above are dubious. The second, for instance, amounts to an assertion that one cannot answer the question 'How do you know that you are in pain?' by simply saying 'Because I am in pain'. Someone who gives this answer clearly takes it that in some cases it works, and we must not beg the question against him.

The second contrasting sceptic offers not so much an argument or a question but rather an attitude. This sceptic is a hard-nosed person who claims that most people allow themselves to be persuaded by what is really rather weak evidence, but that he needs more than that to convince *him*. This sceptic claims to have higher evidential standards than the rest of us; he stigmatizes others as gullible or as too easily persuaded. His position develops into scepticism proper (i.e. the view that knowledge is impossible rather than merely rarer than one thinks) when the standards are set so high that they cannot be fulfilled. But in order to become philosophically interesting (rather than simply an intellectual freak) he must do more than assert that higher standards of evidence are better. He must have some argument that the normal standards are inappropriate in some way. And that argument must be justified by appeal to our standards as well as his. There is then the danger of incoherence; is it consistent to provide an argument justified by normal standards of evidence, to the effect that those standards are inappropriate?

The conclusion is that with which we started; scepticism in its most interesting form always depends on an argument. In the next section we shall consider three sceptical arguments which are strong enough to be worth taking seriously. First we need to look at some distinctions between types of sceptical argument.

The first and least important distinction is between local and global sceptical arguments. Local scepticism maintains that, even if knowledge is possible elsewhere, it is for special reasons not available in this or that selected area. Favourite areas for local scepticism are ethics, religion and the future. We can know how things are in front of our noses, maybe, but it is not possible to know that an altruistic act is morally good, nor that God exists, nor that you will have eggs for breakfast tomorrow. Obviously local scepticism hopes to feed on special features of the areas it is concerned with. But my experience is that it is very hard to keep one's local scepticism local. A local ethical scepticism, for instance, tends very quickly to spread out and become a general scepticism about the unobserved or about the possibility of scientific knowledge. The problem is to find a convincing argument for local ethical scepticism which has no expansionist tendencies.

If local scepticism tends to collapse into global scepticism, this may be an advantage, since global sceptical arguments are generally more convincing and effective than their local counterparts. And the same is true of the second distinction I want to draw. Some sceptical arguments attack the notion of knowledge directly but leave

other related notions, crucially that of justified belief, untouched. Thus I might argue that to know you must be certain, but that one can never be really certain and hence one can never really know. Ignoring for the moment the force of the word 'really' in this argument, we can still feel confident that even if we give up talking of knowledge, granting that a necessary condition for knowing is unfulfilled, we can happily continue to talk about justified belief, distinguishing some beliefs as justified or as more justified than others and others as less justified or even completely unjustified. None of this talk of justification is threatened by the present sceptical argument. We may feel that the argument exposes idiosyncrasies in the concept of knowledge but that we can get by very well both for practical and philosophical purposes with the surviving notion of justified belief. A stronger form of sceptical argument would, however, threaten both notions at once and claim that any defect in the notion of knowledge is equally present in that of justified belief. Such stronger forms are available, as we shall see, and they are always more interesting than their weaker counterparts. The claim that none of our beliefs about the future are ever justified is more important and more interesting than the claim that although our belief that the sun will rise tomorrow is quite probably both true and justified, we cannot really be said to know that the sun will rise tomorrow.

The third distinction is even more important. We can distinguish sceptical arguments which, although they attempt to deprive us of knowledge (or even of justified belief) still allow that we *understand* the propositions whose truth we are no longer allowed to know, from those which claim that the reason why we don't know their truth is that we cannot understand them. An obvious instance would be the contrast between the suggestion that although we understand the proposition that God exists, we could have no evidence that it is true, and the suggestion that the proposition is incomprehensible to us, and hence *a fortiori* we can neither know it to be true nor be justified in believing it.

Of course the distinction drawn above only remains if we assume that it is possible for us to understand a proposition which we would or could never be justified in believing or could never come to know to be true. On a theory of understanding which linked what we can understand with what we could come to recognise as true, the distinction collapses and all the relevant sceptical arguments will be of the strongest type; that is, will claim that we do not even understand the propositions we claim to know.

It might seem that an argument to the effect that we don't understand the propositions we claim to know must be local rather than global. For a global argument would claim that we understand nothing, which is ridiculous first because we clearly do understand something and second because crucially we understand (and are expected to understand) the sceptical argument itself.

1.2 THREE SCEPTICAL ARGUMENTS

Brains in Vats

You do not know that you are not a brain, suspended in a vat full of liquid in a laboratory, and wired to a computer which is feeding you your current experiences under the control of some ingenious technician/scientist (benevolent or malevolent according to taste). For if you were such a brain, then, provided that the scientist is successful, nothing in your experience could possibly reveal that you were; for your experience is *ex hypothesi* identical with that of something which is not a brain in a vat. Since you have only your own experience to appeal to, and that experience is the same in either situation, nothing can reveal to you which situation is the actual one.

Is it possible, however, that though you do not know that you are not a brain in a vat you still know many other things, perhaps more important? Unfortunately if you do not know this there is not much else that you can know either, it seems. Suppose that you claim to know that you are sitting reading a book. You presumably also know that if you are sitting reading, you are not a brain in a vat. We can surely conclude that if you know that you are sitting reading, you know that you are not a brain in a vat, and hence (by simple *modus tollens*) that since you don't know that you are not a brain in a vat (agreed above) you don't know that you are sitting reading.

The principle on which this argument relies can be formalized as the principle of closure under known entailment:

$$PC^k: [Kap \ \& \ Ka(p \rightarrow q)] \rightarrow Kaq.$$

This principal asserts that if *a* knows that *p* and that *p* implies *q*, *a* also knows that *q*; we always know to be true any propositions we know to be the consequences of a proposition we know. (It is a standard practice to express principles of this sort in logical notation, but

for our purposes each logical formula will be followed by its English equivalent, as here. It is worth working through the logical notation, however, for ease of reference later. PC^k is a *closure principle* because it says that a move from something known to something known to be implied by it does not take us outside the closed area of knowledge.) The principle, then, given that *a* does not know that *q* ($\sim Kaq$), and that *a* does know that *p* implies *q* ($Ka(p \rightarrow q)$), allows us to infer that *a* does not know that *p* ($\sim Kap$). It seems therefore to show, more generally, that since you don't know that you are not a brain in a vat you cannot know any proposition *p* of which you know that if *p* were true, you would not be a brain in a vat. And there are similar slightly different arguments, for instance Descartes' version which takes *q* = you are dreaming and argues that since you don't know that you are not dreaming you don't know any proposition *p* of which you know that if *p* were true you would not be dreaming (see the first *Meditation* in Descartes, 1955).

How should we map this argument using the distinctions of 1.1? It is not entirely global; it admits that knowledge is possible, e.g. that you can know that $p \rightarrow q$, and indeed uses this fact as a lever. In fact its grip is restricted entirely to those propositions whose truth would mean that you were not a brain in a vat. It is however a strong argument in the sense that it aims to attack the notion of justified belief in just those areas where it attacks the notion of knowledge. This has not been shown yet, but it can be shown by running a complete analogue of the argument. All we need is to show that your belief that you are not a brain in a vat cannot be justified since nothing in your experience can count as evidence for that proposition, and then appeal to an analogue of PC^k :

$$PC^j: [JBap \ \& \ JBa(p \rightarrow q)] \rightarrow JBaq$$

which holds that if *a* is justified in believing that *p* and that *p* implies *q*, *a* is justified in believing that *q*. PC^j seems just as convincing as PC^k , if not more so.

However the argument does nothing to suggest that you do not understand the proposition that you are sitting reading a book. You are still allowed to understand it even if you can neither know it nor be justified in believing it. The argument only turns into an argument about understanding if we take a special view about understanding, as mentioned earlier.

The Argument from Error

You have sometimes made mistakes even in areas where you felt most confident; simple mistakes in arithmetic, for example. But nothing you can point to in your present situation tells you that this situation is not one in which you are mistaken. For all you can tell, it is relevantly similar to situations in which you have made mistakes. Since you clearly did not know then, how can you say that you know now? For all you can tell, the new situation is no better than the old.

This argument relies on an epistemological version of the principle of universalizability familiar in ethical theory (cf. Hare, 1963, pp. 7–16). A judgement that an action is morally good is universalizable in the sense that by making such a judgement one commits oneself to holding that any relevantly similar action is morally good. When a new and relevantly similar action occurs, one must either call it good or take back one's judgement that the first one was good. What makes a new action relevantly similar to an old one? An action is relevantly similar if it too has the properties which constituted one's reasons for the judgement in the first case. Being relevantly similar is not the same, then, as being completely indistinguishable. There is at least this restriction on the properties that count here, that they must be properties whose presence or absence can be registered by the person making the judgement. A difference between the two actions which that person is unable to recognize cannot justify a difference in judgement. The principle of universalizability tells us, then, that in the absence of an available difference we must make the same judgement again. There must be something we can pick out if a difference in judgement is to be justified.

There may perhaps be properties which are evidence-transcendent, by which we mean that it is always possible that they be absent even though we have the best possible evidence of their presence. Goodness is such a property, perhaps; and this is why the principle of universalizability has teeth in ethics. We cannot suppose that one action is good and another not good unless we can pick out a further relevant difference between them.

The argument from error exposes the consequences of this approach for epistemology. Suppose that I claimed yesterday to know that it would rain in the afternoon, on the normal grounds (weather forecast, gathering clouds, etc.), but that it turns out that I was wrong. At the time of my claim the fact that it was not going to rain was evidence-transcendent, as all claims about the future

must be. And this means that if on the same grounds I claim today to know that it will rain in the afternoon, I must continue to assert that I knew yesterday that it would rain that afternoon (in the teeth of the evidence). If, on the other hand, I abandon my claim to have known yesterday, I cannot make the claim to know today. For the only fact which would justify such a difference in claim is one which is not available to me; facts about the afternoon's weather are evidence-transcendent in the morning. Hence my acceptance that yesterday I did not know prevents me from claiming knowledge today.

We can run through the argument again from the point of view of an outside observer. There is the possibility that though *I myself* would not be justified in making different claims, another person might perhaps be justified in saying that yesterday I did not know, while today I do. This might be so after today's rain, for instance, when the facts about the rain have ceased to be evidence-transcendent to all. Although I couldn't tell the difference at the time, I was in fact wrong yesterday and right today, and this is sufficient to ground a difference in the outside observer's description of me as knowing today but not knowing yesterday. But it is easy to see what the sceptic will say at this point, quite apart from the implied oddity that there is a claim which I could not be justified in making but which another can be justified in making for me. For what is being suggested is that I knew today and not yesterday, despite the fact that there was no difference between the two days that I could tell at the time. But this shows that today, for all I knew, it was not going to rain, and how can we allow that I know that it will rain when for all I know it won't?

The conclusion seems to be that if I recognize that I have once wrongly claimed to know that *p*, then I cannot ever claim to know that *p* unless I can show a relevant difference between the two cases. And no one else can say of me that I know in one case and not in the other, because for all I know I am wrong both times.

So far, however, we have restricted the argument unnecessarily to cases where I have in fact made mistakes in the past. But we do not need to rely on actual mistakes in the past. For our purposes, possible mistakes will do just as well. This can be seen in the ethical case. An imaginary example can be so described that I am willing to say that the action it recounts is good. And that judgement of mine is as much universalizable, as binding on my future judgements about relevantly similar cases, as if the example had been real rather than merely imaginary. Similarly, an imaginary case in which I

would claim to know that p , but where p is false, will succeed in preventing me from claiming to know that p in a new case which is not relevantly (i.e., discernibly to me) different. So imaginary cases are as effective in the argument as actual ones.

It was of course a hallmark of our first sceptical argument that it started from an imaginary case in which you are a brain in a vat, which is not discernibly different from your present case. So the argument we have now reached seems to be a complex defence of the first part of that first sceptical argument. It seems to show that the imaginary case in which you are a brain in a vat being fed the experiences of reading a book is perfectly effective in showing that you do not know that you are reading a book. The difference between the two arguments seems to be in the route they pursue thereafter. The first uses PC^k to show that you don't know anything of which you know that if it were true you would not be a brain in a vat. The second argues more generally that since we have made mistakes, or would make them in imaginarily similar circumstances, we do not know now.

How strong is the scepticism which the argument from error would create if successful? If, as I shall argue in 4.2, there is no separable area in which we make no mistakes, then the argument from error will be global rather than local. But it is not immediately obvious how to write a similar argument against the notion of justified belief. We cannot argue straightforwardly that a false belief cannot be justified. So unless we can say, as we said above for knowledge, that you cannot claim your belief is justified unless you can tell the difference between cases where such beliefs are true and cases where they are false, it will be impossible to conclude that your true belief here is unjustified.

Perhaps, however, we can make this claim by appeal to the initial moves of the argument that you are not justified in believing that you are not a brain in a vat. There we claimed that if nothing in your experience could count as evidence that you were not a brain in a vat, your belief that you are not a brain in a vat cannot be called justified. The belief is unjustified because nothing that you can point to suggests that you are rather than are not a brain in a vat. And equally in our new case we can say that your belief is unjustified because nothing you can point to suggests that this is a case where your belief is true rather than one of the (admittedly rarer but still) indistinguishable (to you) cases where it is false. If this move is sound, our second sceptical argument attacks the notion of justified belief at least as much as the first argument does; in fact even more, because it is more global.

We must allow, too, that this argument as it stands leaves our understanding untouched. As before, unless we have a theory of understanding which links the possibility of understanding with the availability of justified belief, or of knowledge, our understanding survives the loss of justified belief. Such a theory might claim, for instance, that to understand a proposition is to be able to tell the difference between circumstances in which one would be justified in believing it and those in which one would not. There are such theories; more on them below.

The Justification of Arguments from Experience

Do we have any knowledge of events which we have not experienced or are not now experiencing? We normally suppose that our experience is a reliable guide to the nature of those parts of the world which we are not observing, and that in favourable cases it gives us knowledge. Thus I can know what is in the bottom drawer of my desk, or what I shall eat for breakfast tomorrow, by some form of inductive inference from what I have observed or am now observing. David Hume (1711–76), the Scottish historian and philosopher, raised in a special way the question of whether this is really so (Hume, 1955, ch. 4.2). He argued that I cannot know that my diary is in the (closed) bottom drawer of my desk unless I have reason to believe that my experience makes that proposition probable; we can suppose, perhaps, that my relevant experience is that I remember having put the diary there five minutes ago and that I do not remember having touched the drawer since, together with my general knowledge of the consistent behaviour of the experienced world. But I only have reason to believe that my experience makes that proposition probable if I have reason to believe quite generally that events which I have not observed are similar to events which I have observed. And Hume's point is that it is impossible to have any reason for that last belief. For that belief is not analytically or necessarily true; no contradiction is implied by supposing it false. And I cannot suppose that experience itself has given me reason to believe that the unobserved will resemble the observed, since the appeal to experience begs the question asked; it argues not *to* but *from* the crucial belief that our experience is a reliable guide, or that the unobserved will resemble the observed. Therefore I can have no reason to believe that my experience is a reliable guide, and hence have no reason for any belief about events beyond my experience and so cannot have knowledge of them.

It is worth stressing that Hume's argument does not attempt to

derive a sceptical conclusion from the fact that I might be wrong (as in a way the first argument does) nor from the fact that I have been wrong (as the second argument does). Instead, he maintains that our general belief that experience is a reliable guide cannot be justified, since all promising justifications assume what is at issue by supposing that experience can reveal that our experience is a reliable guide. There is an obvious weakness in this attempt to use an argument from experience to justify all arguments from experience.

The scepticism which Hume's argument creates is not global, since it concerns only our knowledge of the unobserved. The argument clearly attacks the notion of justified belief as well as that of knowledge in that restricted area, since it maintains that we have no reason in what is observed for any beliefs about the unobserved. It leaves the notion of understanding untouched; Hume seems to agree that we understand propositions about unobserved objects, although he does argue on independent grounds that they are mostly false.

1.3 A SHORT WAY WITH THE SCEPTIC

None of the three arguments mentioned above is of the strongest type, since none attacks our notion of understanding. There are local sceptical arguments of this strongest type, as we shall see in chapter 5 (our knowledge of other minds); and in our discussion of our knowledge of the past and of the future (chapters 10 and 11) we shall have to bear in mind arguments that it is impossible to conceive of an event as other than present, i.e. as being in the past or in the future. But we might suppose it impossible to provide an effective global argument of this type. It is not just that we clearly do understand *something*; rather we know in advance that it is only by understanding the sceptic's argument as we are clearly expected to, that we could be led to believe that we understand nothing. And even if we don't understand the argument, we surely understand the conclusion; and so the conclusion must be false.

This short way with the global argument can be copied with any global sceptical argument, whether it attacks knowledge alone or justified belief as well. Thus we might say that the sceptic implicitly claims to know his conclusion that knowledge is impossible, or that he claims that his premises justify his belief that justified belief is impossible. The former suggestion seems unconvincing, but the latter

is quite effective. What is the point of arguing that justified belief is impossible, for if you were right there could be no reasons for your conclusion?

These defences against the sceptic attempt to avoid detailed examination of the arguments put forward and focus instead merely on the conclusion. They do this in one of two ways. Either they dispute the sceptic's right to assert the conclusion, or to assert it as a conclusion; or they suggest directly that the conclusion cannot be true, and that hence they are excused from considering any suggested reason for believing it. Instances of the first sort have already been mentioned. An instance of the second sort could be derived from the claim that to understand this or any proposition is to know under which conditions it is true and under which it is false. If the sceptic's conclusion were true we could not have such knowledge; so if the conclusion were true we would not be able to understand it. It is impossible for us to understand the conclusion, then, without realizing that it is false. It is clear, I think, that this argument would, if successful, turn the argument against a global scepticism about understanding into an argument against a weaker but still global scepticism about knowledge. Since we do understand what the sceptic is saying, we must have the sort of knowledge required for that understanding.

In my opinion the sceptic should be entirely unworried by such arguments. He should insist that they provide no justification whatever for an unwillingness to consider his arguments seriously for what they are. To take the weakest case first; even if the sceptic were unwise enough to admit that any assertion involves a claim to knowledge and that he is asserting his conclusion that knowledge is impossible, he can still maintain his position. He takes himself to have true premises and a valid inference to a true conclusion; the premises might include the proposition that in the past he has made mistakes. He may admit (unwisely again) that in using that proposition as a premise he is implicitly claiming to know it to be true. But he insists that it follows from it and others of the same kind that knowledge is impossible. His argument then can be rewritten as follows: if I know this — and this is a central case of the sort of thing I know if I know anything — then I know that knowledge is impossible, and so if I know anything I know nothing. This argument can be seen in two ways. Either it is an instance of a proof by *reductio ad absurdum*, in which we assume something true in order to prove it false; or it is a way of exposing a paradox within the concept of knowledge, for the sceptic can surely insist

that if a central concept such as that of knowledge can be used to take us validly from true premises to a false or impossible conclusion, something is wrong with the concept; there is probably some internal tension which should be exposed rather than swept under the carpet. An answer which merely consists in pointing out the falsehood or impossibility of the conclusion must, therefore, be missing the point.

1.4 ANOTHER REPLY

One common reply to the first two sceptical arguments is to say that we have no reason to worry ourselves about them. Since it is admitted, or rather insisted, that as far as you are concerned there could be no difference between the hypothesis that you are currently sitting reading and the hypothesis that you are a brain in a vat being fed the experiences of one sitting reading, then it cannot matter to you which is really true and which is false. Nothing of any genuine interest or importance can depend on whether you are a brain in a vat or not.

This reply comes in a stronger and a weaker form, but in either form it has clearly got some point. It maintains that what the sceptic takes to be his strength is in fact his weakness. The sceptic insists that there is a difference between the two hypotheses, but that it is evidence-transcendent, i.e. that it is a difference which you cannot tell; and he concludes from this that you don't know which situation you are really in. The reply admits that the difference is evidence-transcendent, and uses that fact against the sceptic. But that fact can be used in two ways.

The weaker way is to say simply that although there is a radical and obvious difference between the two hypotheses, it is not one which could make any difference to you, and so that you can be exempted from paying any attention to it. This is rather like an attitude one might take to the philosophical discussions about whether we have free will or not. One might attempt to reject that entire discussion on the grounds that whether we have free will or not can make no difference to the way we live our lives. We act and will continue to act as if we have free will, whether our actions are determined or not. There seems to me to be something very wrong-headed about this move, both in the areas of free will and of scepticism. But we need not pause to examine its defects, since there is a stronger and more interesting move with which it may be confused.

The sceptic says that there is a difference between the two hypotheses but that you cannot tell it since it is evidence-transcendent. Our weaker move admitted this, but the stronger move denies it. The stronger position denies the existence of evidence-transcendent truth and evidence-transcendent differences, and so denies the sceptic the contrast he needs between his two hypotheses. If the difference suggested is one which *could* make no difference to us, then it is empty and does not exist.

The weaker move, then, says that there is a difference which does not matter. The stronger move says that there is no difference to matter. We could call the former a *realist* position; the realist believes that there are evidence-transcendent truths, truths whose obtaining lies beyond our powers of recognition. The stronger move could be called *anti-realist*; it denies the existence of evidence-transcendent truth and holds that differences which we are in principle incapable of recognizing do not exist.

Anti-realism of this sort does not arise gratuitously, nor is it intended initially as a method of countering scepticism. But its general thrust is clear. The realist and the sceptic think of the world as one on which we have at best a tentative grasp. There are many facts about the remote past and the remote future which we shall have no means ever of recognizing or verifying. And it is always possible that unknown to us the present world differs radically from the way it appears to us. The anti-realist does not believe in the existence of this further 'real' world which lies behind the world that we know and which may come apart from our world in ways which of course we could not recognize if they occurred. For him our world, the recognizable world, is the only world. So for the anti-realist the enterprise of epistemology is easier, since the objects of knowledge are brought closer to us; and there is no yawning gap between evidence and truth, since there can be no evidence-transcendent properties. For the question whether a property is present here can never be different from or lie beyond the question whether we have the best possible evidence that it is present.

Anti-realism is the theory of understanding which has been mentioned at various points in this chapter. (Its name and recent development are owed to Michael Dummett.) The anti-realist holds that our understanding of the sentences in our language must have been acquired in situations which we learnt to take as warranting the use of those sentences; situations in which those sentences are to count as true. It follows from this that if there is no such thing as justified belief, there is no such thing as understanding. For to

understand a sentence is to be able to pick out situations which justify us in believing that sentence to be true.

It might seem, then, that the anti-realist is in a peculiarly weak position. Every sceptical argument against the possibility of justified belief is an argument of the strongest form, and leaves us devoid even of the understanding we thought we had of our own language. But in fact the reverse is true. The sceptical arguments which would have this effect all require a move which the anti-realist would disallow, and hence never reach the annihilating conclusion. They all require us to make sense of the realist thought that it is always possible that, unknown to us, the world differs radically from the way it appears to us, and argue from this that we cannot know that the world really is the way it appears to us. But the anti-realist rejects this as impossible. For him, the sense of a sentence is determined by the sorts of situation we count in favour of the sentence being true, in such a way that the sentence with that sense (i.e. as we understand it) cannot be false if the sort of situation we count as making it true occurs. So anti-realism offers a perspective from which not only is there no possibility of a global scepticism about understanding, but also (and for the same reasons) there is no room for a global scepticism about justified belief either.

The trouble here is that the implausibility of the truth of scepticism is about as great as the implausibility of the truth of anti-realism. To see this, we need to see how much anti-realism requires us to give up and how strong our realist intuitions are.

One area which seems to demand a realist approach is that of other minds, which will be discussed in chapter 5. Our realist intuition here is that the sensations and thoughts of others, which do occur, are hidden from us. We observe their behaviour, of course, but the question whether they are actually experiencing sensations as we suppose is, for us, evidence-transcendent. It is always possible that despite all the behaviour there are actually no sensations going on there at all, or at least different ones from those we imagine. So there is a real question whether or not there are sensations which are not ours, but it is evidence-transcendent.

Another is that of the past. Whatever we may think of the future, we think of the past as having been in its time as determinate as the present now is. But we suppose that there are many propositions about the past whose truth is for us evidence-transcendent. Despite our lack of grip on these truths, we do take there to be a transcendent fact of the matter at stake, one that lies now beyond

all possibilities of being recognized by us. And this attitude about the past is a realist attitude.

These two areas will be considered in greater detail in later chapters. The point of our discussion of them so far is that the attempt to rebut scepticism by constructing an anti-realist alternative to the realism espoused by the sceptic is not going to be easy, even if it is possible. If there are some areas where anti-realism is comparatively easy to accept, well and good. But so long as there remain others where realism seems compelling, in those the sceptic's challenge bites. We may be unable to buy our way generally out of scepticism in the anti-realist market; the cost would be too high.

1.5 A BETTER RESPONSE

If no short way with the sceptic is possible, we have no alternative but to get involved directly with the arguments presented. Where might we look for help in a critical offensive? One hope might be that a satisfactory account of what knowledge is would have the effect of exposing errors in the sceptic's reasoning. And I shall be considering an account which claims to do this in chapter 3.

Meanwhile we should perhaps consider what state we would be in, were we to agree that the sceptic's argument is effective. It has sometimes been suggested that epistemology could survive the loss of the concept of knowledge, because all the important epistemological questions can equally profitably be rephrased using instead the concept of justified belief. Thus instead of asking whether we ever know what will happen in the future, we ask in which if any circumstances our beliefs about the future are justified. And the problem of other minds (ch. 5) can be presented as the problem of how, if at all, my observation of the behaviour of human bodies justifies my belief that those bodies are people; little extra is gained by asking whether and how I *know* that they are people.

The difficulty with this suggestion is that all the sceptical arguments presented, and indeed any interesting sceptical argument, seem to be directed as much against the notion of justified belief as against that of knowledge. And this means that the easy compromise position is unavailable. It really does seem to matter for epistemology that we find some reply to the sceptic.

FURTHER READING

Descartes' dreaming argument is in the first of his *Meditations*; use the Haldane and Ross edition (Descartes, 1955) or any reputable alternative.

Hume's questions about induction are raised in his *Inquiry Concerning Human Understanding*, section 4, part 2 (Hume, 1955; or any reputable alternative).

Stroud (1984, ch. 1) offers a very readable account of Cartesian (i.e., Descartes') scepticism.

The argument about brains in vats is given in Nozick (1981, pp. 167–71). His account will be examined further in ch. 3.

Stroud (1968) gives a helpful account of the anti-sceptical move considered in 1.3, commonly called transcendental. He links it to verificationism (the 'verification principle'), an earlier form of anti-realism.

Dummett, architect of anti-realism as successor to verificationism, offers the most introductory account I know in Dummett (1978, ch. 10) but if you are new to this area you would do better to wait until you have read ch. 9 of the present book before attempting it.

2

Knowledge

2.1 THE TRADITIONAL ACCOUNT

The standard account of knowledge, around which all recent work has been done, defines knowledge as justified true belief; it holds that *a* knows that *p* if and only if

- 1 *p*,
- 2 *a* believes that *p*,
- 3 *a*'s belief that *p* is justified.

Because there are three parts to this definition it is called the tripartite definition or the tripartite account; it defines propositional knowledge, knowledge that *p*; it does not define knowledge by acquaintance as in '*a* knows James' nor knowledge-how, e.g. knowledge how to ride a bicycle, unless these can be shown to reduce to knowledge-that.

The tripartite definition has obvious attractions. The first clause, that if *a* knows that *p* then *p* is true (which can be read as $Kap \rightarrow p$), is normally seen as stipulative. The second clause, that if *a* knows that *p* then *a* believes that *p* (we can read this as $Kap \rightarrow Bap$), seems minimal, and the third, that if *a* knows that *p* then his belief that *p* is justified ($Kap \rightarrow JBap$), is there in order to prevent any lucky guess from counting as knowledge if the guesser is sufficiently confident to believe his own guess. It is worth noticing, however, a consequence of this justification of clause 3; this is that a belief is not generally considered to be justified by the mere fact that it is true, for otherwise clause 3 would be unnecessary. If I decide on the toss of a coin which investment will provide the greatest yield, and fortunately turn out to be right, we suppose that my choice is vindicated by the outcome perhaps, but not justified by it; I had no real justification for making the choice I did. (Alternatively we

could distinguish between two forms of justification, justification before the event and justification after it, and run the tripartite definition in terms of the former; but then the question would be whether these really are two forms of the same thing.)

What are the problems for the tripartite definition? One might think that clause 2 is insufficient: to believe that p is not so strong as to be certain that p , and to know one must be certain, not just believe.

The best reason for wanting some account of certainty in our analysis of knowledge is that people are rightly hesitant to *claim* knowledge when they are less than certain. This hesitation seems to be due to something about what knowledge is, and there is no obvious way to explain it if knowledge is as the tripartite conception claims it to be. Thus although it is commonly suggested that the notion of certainty is relevant to the analysis of *claims* to knowledge, but not to the analysis of knowledge itself (e.g., in Woozley, 1953), this leaves us with no method of explaining why certainty should be required before one can claim knowledge when it is not required for knowledge itself, i.e., for the existence of what one is claiming.

Since we are going to discover other reasons for rejecting the tripartite definition, we have no reason to pursue this point here. The moral to be drawn is that if we are to give an account of knowledge which does not include a requirement of certainty, our account should make room for the notion of certainty somewhere; if it sees certainty as a requirement for a knowledge claim it needs to be able to explain in its own terms why that should be so.

But why should we be reluctant simply to change clause 2 to ' a is certain that p '? The answer is that we are prepared, in circumstances that are not particularly unusual, to allow that someone does in fact have knowledge when that person is so far from certain that he would not claim the knowledge himself. The classic example offered is that of the diffident schoolboy, who has learnt the dates of, say, the English kings the previous night but who is so alarmed by his hectoring schoolteacher that he becomes completely unsure that the answers that suggest themselves to him under questioning are in fact the right ones. Supposing, however, that those answers are correct, would we not allow that he knows them, even though he himself might not make that claim? And surely our reasons for allowing this are close to those suggested by the tripartite definition; he has the right answer, and not by luck.

There is a weakness in this appeal to the diffident schoolboy,

which concerns clause 2 again. To the extent that the schoolboy is less than certain of the answers that occur to him, can we allow that he still believes them? If we are not careful, use of this example to defuse the pretensions of a certainty condition will result in our losing the belief condition we were trying to defend.

2.2 GETTIER COUNTER-EXAMPLES

Henry is watching the television on a June afternoon. It is Wimbledon men's finals day, and the television shows McEnroe beating Connors; the score is two sets to none and match point to McEnroe in the third. McEnroe wins the point. Henry believes justifiably that

1 I have just seen McEnroe win this year's Wimbledon final.

and reasonably infers that

2 McEnroe is this year's Wimbledon champion.

Actually, however, the cameras at Wimbledon have ceased to function, and the television is showing a recording of last year's match. But while it does so McEnroe is in the process of repeating last year's slaughter. So Henry's belief 2 is true, and surely he is justified in believing 2. But we would hardly allow that Henry knows 2.

This sort of counter-example to the tripartite account of knowledge is known as a Gettier counter-example, after E. L. Gettier (1963). (I owe this particular example to Brian Garrett.) Gettier argued that they show the tripartite account to be insufficient; it is possible for someone not to know even when all the three clauses are satisfied.

Gettier here is not quarrelling with any of the three clauses. He allows that they are individually necessary, and argues only that they need supplementing.

It is worth formalising the situation, for reasons which will emerge later. Reading 1 as p and 2 as q , we have:

$\sim p, \text{Bap}, \text{JBap}, p \rightarrow q, \text{JBa}(p \rightarrow q), q, \text{Baq}, \text{JBaq}.$

So a Gettier counter-example is one in which a has a justified but

false belief by inference from which he justifiably believes something which happens to be true, and so arrives at a justified true belief which is not knowledge.

What response should be made to these infamous but slightly irritating counter-examples? There seem to be three possible routes:

- 1 find some means to show that the counter-examples do not work;
- 2 accept the counter-examples and search for a supplement to the tripartite analysis which excludes them;
- 3 accept the counter-examples and alter the tripartite analysis to suit rather than adding anything to it.

The remainder of this section is concerned with the first route.

On what principles of inference do these counter-examples rely? Gettier himself exposes two. For the examples to work, it must be possible for a false belief still to be justified; and a justified belief must justify any belief which it implies (or is justifiably believed to imply). This last is just the principle of closure PC^j mentioned above in the discussion of scepticism (1.2). So if we could show PC^j false this would have the double effect of undermining the Gettier counter-examples and (part at least of) the first sceptical argument. It might be possible, however, to construct new variants on the Gettier theme which do not rely on inference or on an inference of this sort, as we shall see in the next section, and if so no complaints about PC^j or other principles will be very effective.

One thing we cannot do is to reject Gettier counter-examples as contrived and artificial. They are perfectly effective in their own terms. But we might reasonably wonder what point there is in racking our brains to find an acceptable definition of 'a knows that p '. Is this more than a mere technical exercise? What, if anything, should disconcert us if we cannot come up with a trouble-free definition? Many of the innumerable papers written in response to Gettier give the impression that responding to Gettier is a kind of private philosophical game, which is of no interest except to the players. And hasn't Wittgenstein shown us anyway that a concept can be perfectly healthy without being definable, arguing that there need be no element common to all instances of a property (e.g. instances of knowledge) other than that they are instances (e.g. that they are knowledge)? (Cf. Wittgenstein, 1969b, pp. 17–18, and 1953, §§ 66–7.) So what on earth could depend on our success or failure to discover necessary and sufficient conditions for knowledge?

In many ways I sympathize with the general tenor of this

complaint, as may quickly become apparent. What sustains me in the search for a response to Gettier is the feeling that it may be possible to find an account of what knowledge is which will have a substantial effect on what we are to say about justification in later parts of this book. This could happen in either or both of two ways. We might find an account of what knowledge is which would suffice to undermine crucial sceptical moves, and hence confirm the possibility that some of our beliefs are justified; the account to which I give tentative support in chapter 3 has pretensions in this direction. Or we might hope to define justification in terms of knowledge. For instance, we might suppose that a belief is justified iff in certain circumstances (to be spelled out) it would be knowledge. (Jennifer Hornsby gave me this idea.) In the meantime we must consider some accounts of knowledge which seem less fruitful.

2.3 RESPONSES TO GETTIER

I rather obviously avoided, so far as I could, offering even the most tentative diagnosis of the defect in the tripartite analysis which Gettier exposed. This is because the different responses to Gettier all stem from different diagnoses of the way in which the tripartite analysis is lacking; once we know what is missing, it should be quite a simple matter to add it.

The Presence of Relevant Falsehood

The most obvious diagnosis is simply that the initial belief that p , from which the true justified belief that q is inferred, is false. So we might add to the tripartite analysis the fourth condition that nothing can be known which is inferred from a false belief, or from a group of beliefs of which one is false.

This simple suggestion has two defects. First, variants on the Gettier theme can be written in which, though there is falsehood, there is no inference. Suppose that I believe that there is a sheep in the next field because of what I see. I am not inferring from what I see that there is a sheep in the field; I take myself simply to see that there is one. The animal I see is a large furry dog, but my belief is not false, because there is a sheep there too, unknown to me, hidden by the hedge. Here we might admit that my belief is true and justified but refuse to grant that I know there to be a sheep in the field. (This example comes from Chisholm, 1977, p. 105.)

A reply might be that surely I am inferring that I see a sheep in

the field from my knowledge of my own present sensory states. This reply raises large issues; but chapter 5 contains a lengthy argument that if there is any non-inferential knowledge, some of it concerns things other than our sensory states — so why not sheep, for instance?

The second defect is that the suggestion is too strong and is likely to make it impossible for any of us to know anything at all. As we shall see, this is a danger with a number of responses to Gettier. In the present case, we all of us suffer from numerous false beliefs which have some role in our inferential processes, and so on this suggestion none of our present true justified beliefs would count as knowledge.

To eliminate these defects we must remove the reference to inference and tighten up the relation specified between the false beliefs and the true justified ones which are not to count as knowledge. Thus we could simply require an absence of relevant falsehood. This would get round the example of the sheep in the field because I presumably believe (falsely) that the animal I can see is a sheep even though this belief is not used in inference. But as a suggestion it seems rather to name the difficulty than to solve it: which false beliefs are to be counted as relevant?

An answer might be that a false belief that p is relevant in the required sense if, had the believer believed instead that $\sim p$, his belief that q would cease to have been justified. Not all false beliefs are relevant in this sense. Some will be so distant or insignificant that whether one believes them or their opposite would have no effect on what one believes here. For instance, among the beliefs in virtue of which I claim to know that Napoleon was a great soldier there may be one which is false, but which is so insignificant that my justification for believing that Napoleon was a great soldier would survive my changing my mind on that particular point. Such a false belief would not be relevant in our present sense.

But the new account faces difficulties, which can best be illustrated by an example. Suppose that I expect a colleague to give me a lift home this evening, but that her car has a flat battery; this won't stop us, however, because a friend's car is parked conveniently near with some jump leads which we can use to get her car to start. I now believe that she will give me a lift this evening, and this belief is justified. Do I know that she will give me a lift? The requirement that there be no relevant false beliefs suggests plausibly that whether I know depends on what other beliefs I have. But this suggestion raises difficulties. If, for instance, I merely believe

1 she will give me a lift home this evening,

I may be allowed to know this, but if I believe both 1 and

2 her car battery is not flat,

I may not, since I have a relevant false belief. But if, as well as believing 1 and 2, I happen also to believe

3 there is a friend's car conveniently near with jump leads,

then this apparently gratuitous belief makes it the case again that I know that I will be given a lift. For if I had believed the opposite of 2, my belief 1 would not be justified unless I also believe something like 3. It seems then that our present suggestion has the effect that whether I have knowledge will depend commonly on which other apparently gratuitous beliefs I may have. There is something unsatisfactory about this, and more work needs to be done to defend the account against complaints of this sort.

Defeasibility

A slightly different approach diagnoses the Gettier counter-examples as arising because there are some truths which would have destroyed the believer's justification had he believed them (cf. Lehrer and Paxson, 1969; Swain, 1974). Thus, for instance, suppose that Henry had believed that he was watching a recording of last year's final, as he was; in that case, his justification for his belief that p and thus by PC¹ for his belief that q would have been destroyed. The suggestion then is that a fourth clause be added requiring that there be no other truth such that Henry's believing it would have destroyed his justification for believing that q . This is the defeasibility suggestion; we require for knowledge that the justification be indefeasible, i.e. that the addition of further truths should not defeat it.

This will not imply that a false belief will never be justified, since the suggestion is that although some beliefs are defeasibly justified, we require indefeasible justification for knowledge. However, it is in danger of rendering the first condition for knowledge ($Kap \rightarrow p$) redundant. It looks as if a false belief could never be indefeasibly justified since there would always be some truth (even if only the negation of the false belief) whose addition would destroy the justification. But perhaps this is a strength in the theory rather than a weakness, since the new quadripartite analysis will have a

coherence that was previously lacking; it provides an explanation in the fourth clause of what was before included by mere stipulation, that knowledge requires truth.

The defeasibility suggestion could be said to provide an extension of the earlier requirement that there be no relevant falsehoods; we now look beyond those propositions actually believed by the believer to propositions which would have an effect if they were believed. But this extension is no real advantage. The sort of difficulty facing the notion of defeasibility can again best be illustrated by an example. Thus, perhaps, I believe that my children are at present playing in the garden at home, and I have good reasons for this belief. However, unknown to me, a neighbour rang up after I left home this morning to invite the children round for the morning. And if I had known this my justification for believing that they are playing at home would be defeated, because I also believe that they normally accept such invitations. However, my wife has become concerned about the health of one of them and refused the invitation. Do I know that my children are playing in the garden at home? If your intuition is that I do, you must reject the defeasibility criterion as formulated at present. If it is that I do not, on the grounds that had I heard about the invitation my justification would have been defeated, you have a duty to give some account of why the (unknown to me) truth that my wife has refused the invitation does not somehow redress the balance. Either way the defeasibility proposal needs to be altered.

The problem seems to lie, as it lay for the requirement that there be no relevant falsehood, in the way in which new true beliefs can be added piecemeal and overturn the existing justification, while there remain yet further truths waiting in the background to overturn the overturning. First we want to ask anyway whether there isn't likely always to be some truth which, if it alone were added and all others excluded, would defeat my justification. Even if this won't always happen, it will certainly happen often enough for the range of my knowledge to be considerably reduced, and this itself is some sort of an objection. Second, we need to find a way to counter the way in which the piecemeal addition of further truths seems to switch me into knowledge and out again.

We might achieve the second task by altering our account of defeasibility so that instead of talking about some one other truth (which caused the problem of piecemeal addition) we talk about all truths whatever. Thus we could require as our fourth condition that our justification would remain even when every truth is added to

our belief set, all at once. This new notion of defeasibility seems to allow (probably) that I now know that my children are playing in the garden, because the second added truth negates the defeating powers of the first. But there remain problems for this new notion of defeasibility. First, in talking of adding all truths at once we seem to have moved firmly into the realm of fiction. Indeed, do we have any suitable conception of 'all the truths'? Second, it seems on this criterion that we shall never have more than the slenderest of reasons to believe that we know something; for in believing this we are believing that when all the truths are in, our justification will remain, and it looks as if much more is required to support that belief than is required to support an ordinary claim to know.

Reliability

A different approach diverts our attention away from the relation between the proposition claimed as knowledge and other false beliefs which should have been true or other truths which should have been believed. It is sometimes suggested that a justified true belief can be knowledge when it derives from a reliable method (see Goldman, 1976; Armstrong, 1973, ch. 13; Swain, 1981). In the Gettier example, Henry does know that the Wimbledon final is being played that afternoon; this justified true belief derives from the reliable method of reading the newspapers, which are normally right about this sort of thing. However, his belief that *q* clearly derives from a method that is less than reliable. It would have led him badly astray here, had McEnroe suffered an unexpected lapse and succumbed to the efforts of Connors for once.

The reliability approach can be made more elaborate; in some ways it is closely related to the causal approach considered next, because we are clearly owed an account of what reliability is, and a causal answer is tempting (see, e.g., Goldman, 1979). However, we can already see difficulties for any variation on this approach. It is in danger either of making knowledge impossible or of walking straight into one of our sceptical arguments.

We may mean by 'reliable' that a suitable method, if properly followed, is perfectly reliable and never leads to a false belief. But, quite apart from the general difficulty of distinguishing between a defect in the method and a defect in the manner in which the method has been applied, it seems improbable that there are any perfectly reliable methods of acquiring beliefs. Man is fallible, and his fallibility is shown not just in the manner in which the methods are used but in the belief-gathering methods available to him. Hence

if knowledge requires an infallible or perfectly reliable method, it is impossible.

But if we retreat from the notion of perfect reliability and require only that the method be generally reliable, we invite sceptical arguments of our second type. How is it that a method which has failed elsewhere in relevantly similar circumstances suffices to yield knowledge this time? If we had any hope that our eventual account of knowledge would help us to reject the sceptical arguments, this particular account seems to make matters worse rather than better. Of course this won't show that the account is wrong. It may be that the correct account of knowledge does unfortunately give the sceptic the opening he is looking for. But we should not accept that this is how things are until we are convinced that there is no other account of knowledge which offers the sceptic less leverage. We can still hope for one which makes life harder rather than easier for him.

A final retreat would be to require only that the method be reliable *this time*. This has the effect of diverting our attention away from previous cases where the method has failed and hence of escaping the sceptical argument which takes its start from those cases. But we might reasonably doubt whether the requirement that the method be reliable this time amounts to any genuine addition to the tripartite account. If reliability is defined in terms of the production of truth, it adds nothing to the first condition once we restrict our attention to the particular case. If it is defined in terms of justification, it adds nothing to the third. And no other accounts seem very inviting. (It may be, however, that the causal theory amounts to a notion of justification in the particular case; see 2.4)

Conclusive Reasons

A different approach diagnoses Henry's failing in the Gettier case as due to the fact that his reasons were less than conclusive. If we require for knowledge that the justified true belief be based on conclusive reasons, all the Gettier cases, and indeed any case in which the believer is right by accident, fall to the ground.

All the work in this approach must go into a persuasive account of what it is for reasons to be conclusive. One suggestion would be that where beliefs A – M constitute conclusive reasons for belief N, A – M could not be true if N is false. This will exclude the counter-examples, but it will also make knowledge a rare phenomenon at best. Empirical knowledge, anyway, looks impossible now; in the empirical realm, our reasons are never conclusive in this sense.

A weaker account, owed to F. Dretske (1971), suggests that someone's reasons A – M for a belief N are conclusive iff A – M would not be true if N is false. This is weaker because to say that A – M *would* not be true if N is false is not to say that they *could* not be true if N is false, as the stronger account demands. It is so weak as not really to provide a genuine sense of 'conclusive', but this doesn't really matter. This weaker account seems to me promising in its general approach, and the theory I shall be supporting in the next chapter is distinctly similar. But it differs in not talking about reasons; and this is a virtue because it does seem possible that there should be justified belief without reasons. My belief that I am in pain may be justified, perhaps, but I can hardly be said to base it on reasons, conclusive or otherwise. I don't base it on reasons at all.

The Causal Theory

A. I. Goldman proposes a causal supplement to the tripartite definition (Goldman, 1967). An initial diagnosis of the Gettier counter-examples may be that it is just luck that Henry's justified belief is true. This diagnosis cannot itself provide a suitable answer. We cannot merely stipulate that there be no luck involved, because we all of us rely on luck to some extent. For instance, the fact that our reliable belief-gathering method provides here a true belief rather than a false one, as it sometimes does, will be just luck as far as we are concerned. And of course the fact that luck is always involved somewhere gives the sceptic a toehold too. But the diagnosis can suggest a better answer. Goldman's suggestion is that what made the belief true in the Gettier case is not what caused Henry to believe it. So he proposes, as a fourth condition for knowledge that *p*, that the fact that *p* should cause *a*'s belief that *p*. This excludes the Gettier cases because in them it is coincidental that the belief is true. We want a link between belief and truth to prevent this happening, and a causal link looks promising.

Attractive though this approach is, it faces difficulties. The first is that we may find it hard to suppose that facts can cause anything; surely they are too inert to affect the way the world goes, even where that world is the merely mental world of beliefs. What, after all, are facts? One's first idea is that facts are similar to, if not identical with, true propositions (which would explain why there are no false facts). But can true propositions cause anything? Surely facts (or true propositions) reflect the world rather than affect it. The prevalent analyses of causation seem justifiably only to allow events

and possibly agents as causes. Second, there is a problem about knowledge of the future; Goldman's suggestion seems to require that here either we have an instance of backward causation (the future causing the past) or that knowledge of the future is impossible since causes cannot succeed their effects. Third, there is the problem of universal knowledge, or more generally of knowledge by inference. My belief that all men are mortal is caused, but not by the fact that all men are mortal; if any facts cause it, they are the facts that this man, that man, etc., have died. And these men are not caused to die by the fact that all men die (which would restore the causal analysis, with an intermediary cause); rather, all men die because those men do (among others). How then can the causal analysis show that I know that all men die?

There are answers to some of these criticisms, of course. We are more used to talking about facts as causes than the first criticism allows. The fact that philosophers have not yet persuaded themselves that they understand the idea that facts can be causes should not cause us to rule out all appeal to fact-causation as philosophically unsound. (The preceding sentence is a case in point.) The second criticism, too, might be answered by complicating the theory by allowing facts to be known in cases where fact and belief are different effects of a common cause. The third criticism, however, seems more intractable. The admission that facts can be causes will not much improve our willingness to suppose that universal facts can cause universal beliefs.

There are promising aspects about the causal theory, and the theory which I shall support can in fact be seen as a generalization from it.

2.4 CONCLUDING REMARKS

The various proposals considered in the previous section were presented as if they were additions to the tripartite analysis, it being admitted that Gettier had shown that analysis to be insufficient. But we can find among them at least one which can be seen as a direct defence of the tripartite analysis. Any proposal which amounts to a new theory of justification may succeed in showing that in the Gettier cases the relevant true beliefs were not justified at all. And we could take the causal theory in this way. The causal theory could be telling us that a belief is only justified when caused (directly or indirectly) by the facts. It would then be adopting route 1, as

distinguished in 2.2. (Some versions of the reliability proposal could also be seen in this light.) Moving this way, then, we would be starting from a causal theory of justification; the causal theory of knowledge would simply be one of its consequences.

A possible way of arguing against a causal theory of justification would be to claim that we have no guarantee that there is only one way in which beliefs come to be justified, and in particular no real reason for supposing that any acceptable way must somehow be causal, so that all justified beliefs that p must be caused by relevant facts. Surely we don't want to rule out in advance the possibility that some moral beliefs, say, are justified, doing so just because we don't want to admit the existence of moral facts (if we don't). And we might still be doubtful of the existence of causally effective mathematical facts, without wishing to say that no mathematical beliefs can therefore be justified.

More importantly, however, the suggested causal account of justification is false because it denies the possibility that a false belief be justified. A false belief that p has no fact that p to cause it. This objection can only be evaded by finding a different account of the justification of false beliefs from that which is offered for true ones. But that cannot be right. Justification must be the same for true as for false beliefs, if only because we can ask and decide whether a belief is justified (e.g. a belief about the future) before we decide whether it is true or false.

This criticism leaves open the possibility of a different sort of causal theory, on the lines suggested at the end of 2.2. With a causal theory of knowledge and the thesis that a belief is justified iff if true it would be knowledge, we can give a causal account of justification which is not vulnerable to the existence of false justified beliefs.

FURTHER READING

Central papers in the area are Gettier (1963), Dretske (1971), Goldman (1967) and Swain (1974).

Perhaps the earliest discussion (and rejection) of the tripartite definition is in Plato's *Theaetetus* (Plato, 1973, 201c–210d).

The enormous industry recently generated by perceived defects in the tripartite account is painstakingly analysed in Shope (1983), with copious references. There are of course many approaches and variants on approaches to the Gettier problem which I have not discussed, including Shope's own.

Most of the papers referred to in the present chapter are collected in Pappas and Swain (1978), which also contains an analytical introduction to the area.

Prichard (1967) gives an interestingly different account of the relations between knowledge, belief, certainty and truth.

An important question which we have not discussed is whether knowledge implies belief. For this, cf. Ring (1977).

The papers by Gettier, Prichard and Woozley are collected in Phillips Griffiths (1967).

3

The Conditional Theory of Knowledge

3.1 THE THEORY

This theory, which we owe mainly to Robert Nozick, takes its start as others do: from the defects which Gettier exposed in the tripartite analysis. Nozick suggests that the reason why we take the justified true beliefs in those examples not to have been known is that *a* would have believed them even if they had been false. The reason why Henry's belief that McEnroe is this year's champion was too lucky or too luckily true to count as knowledge is that his route to this lucky truth was such that even if it had been false, he would still have ended up believing it. Nozick takes it therefore that for *a* to know that *p* we require that *a* would not have believed that *p* if *p* had been false.

This gives us, so far, the standard two conditions:

- 1 *p*
- 2 *a* believes that *p*,

with

- 3 if *p* were not true, *a* would not believe that *p*.

But Nozick argues that although this account may cope with the examples Gettier offers, there are other similar examples which would escape what we have so far. There are two ways in which it can be a coincidence that *a*'s belief is true, and both need to be ruled out. The first is that if it were false, *a* would still believe it; we have dealt with this already by the addition of clause 3. The second is that there may be slightly different circumstances in which it remains true, but *a* no longer believes it. Many examples are

Dummett (1978, ch. 21) is an attempt to make the best sense of the dispute between realist and anti-realist about the past.

Martin and Deutscher (1966) attempt a general account of memory, stressing the need for a causal connection between past perception and present 'representation'.

13 Induction

13.1 INDUCTION, PERCEPTION AND MEMORY

In preceding chapters we have considered our ability to gather knowledge about our surroundings, and our ability to retain or recover such knowledge later. In the present chapter we consider our ability (or lack of it) to move beyond that knowledge; to construct new knowledge on the basis of the old. We can extend our knowledge by reasoning; by seeing that things we already know provide reasons in favour of other beliefs. Where those reasons are strong enough, we can hope that in believing as they suggest, we have acquired new knowledge. A true belief, based on previous knowledge which provides sufficient inferential justification for that belief, will be knowledge, we may hope.

There are two styles of reasoning, deductive and inductive. Deductive reasoning occurs where we take our inferential justification to be conclusive, in the special sense that it is impossible, on pain of self-contradiction, for the beliefs which are our reasons to be true and the conclusion that we draw from them to be false. When we are right in taking this to be so, our deductive reasoning is valid; otherwise it is invalid.

Inductive reasoning occurs when we take our reasons to be sufficient to justify our conclusion, without being conclusive in the sense above, or when we think we have some but not yet sufficient reason for the conclusion, hoping perhaps that further reasons may yet be found so that the sum total of reasons will be sufficient. This can be most clearly expressed in terms of probability. A successful inductive argument is one which makes its conclusion probable, or more probable than any equally detailed alternative; the (relative) probability it gives its conclusion may not be sufficient yet to justify our believing it, because there may be stronger reasons on the other side, or because the degree of probability gained is not large enough

to justify more than an increased willingness to look further. But we suppose it possible that with further reasons we shall eventually be justified in accepting the conclusion.

Inductive reasoning is employed in almost every branch of human enquiry. We can ask whether what we know about the present provides inductive justification for beliefs about the past, as a detective might in a criminal investigation. Or we can ask whether our beliefs about the past justify certain beliefs about the future, as someone might who is considering an offer of marriage (though I do not recommend the approach in this instance). Induction, therefore, is not specially concerned with knowledge of the future. But we might be tempted by the reverse. Can knowledge of the future be gained other than inductively?

There are two ways in which we might hope to gain knowledge of future events other than by inductive reasoning. The first is by supposing that we sometimes know what will happen because we have a non-inferential knowledge of our own intentions. This raises complex issues which we shall do best to avoid for present purposes, merely noting that it is not obvious that this sort of knowledge of the future is entirely independent of inductive support. We may have non-inferential knowledge of our intentions, perhaps, but doesn't our resulting knowledge of the future depend upon a general knowledge that in certain areas our intentions are normally implemented?

The second way in which we might hope for non-inductive knowledge of the future is by supposing that in certain cases we can see what will happen. This suggestion may be made in extreme cases, as when we think of fortune-tellers observing future events in a crystal ball; or in more ordinary cases, as when we talk of seeing that the crash will happen or that the ball will go out. But there is a powerful argument that we can never observe the future in either sort of case.

There needs to be a strong argument, for some of the things that have been said in earlier chapters do seem to make room for the idea that, at least in a limited way, observation of the future is possible. Perceptual belief may be a matter of degree; some beliefs are more perceptual than others, some very perceptual, some less so, and perhaps none purely perceptual. This makes room for us to say, for example, that one can see that a climb is difficult or a cliff dangerous; we have not forced ourselves into a position in which we have to say that such things can only be known by inductive reasoning from what can more properly be said to be seen. Perhaps there is here sufficient latitude for us to avoid ruling out

talk of seeing what will happen, which is in fact quite common; for instance, talk of seeing that the ball will go out.

The crucial argument driving us to rule out perception of the future completely is an argument about the direction of causation. We suggested earlier that perceptual beliefs differ from others in that the facts which are their contents are able to be the main cause of the belief in favourable cases. But if there is to be perception of the future, the things that are the contents of the beliefs must lie in the future. But can something that is still entirely in the future be any part of the cause of a present event? Can the future cause the present? Even if we allow that a cause need not precede its effect but can be simultaneous with it, there are arguments that the present cannot be an effect of an event which has not yet happened.

The classic sort of example is that of someone who invariably wakes up five minutes before his alarm clock goes off, no matter what time he sets it for. What prevents us from saying that the cause of his waking is the fact that his alarm clock would go off five minutes later? To say this we have to say that the effect (the waking) can be over before the cause has even started. And this is supposed to be impossible. Suppose that one day he has set his alarm clock for eight o'clock and wakes at five to eight, but the alarm clock fails to go off at eight for some mechanical reason. On that occasion we have to look elsewhere for the cause of his waking, and we are likely to choose an earlier event such as his setting the alarm for eight. But if it is the earlier setting of the clock that caused the later waking in this case, mustn't we in equity say that this is the cause in every case? The ringing of the alarm which occurs (normally) five minutes after the man wakes can then be seen as a separate effect of a common cause, rather than as a later cause of an earlier effect.

This argument can be generalized. It is no reply to say that there might be someone who just didn't wake up when his alarm clock was defective and was going to fail to ring. This is because the structure of any example seems to require that the effect be over before the cause starts, or at least that there be a gap between the beginning of the effect and the beginning of the cause during which someone could intervene to prevent the occurrence of the cause. It is not necessary that anyone should ever intervene; it is enough that intervention be possible. For the possibility of intervention is incompatible with the thought that the waking is caused by the later ringing of the alarm. This is because of what it is for one event to be an effect of another. B is only an effect of A if the occurrence

of A is necessary for the occurrence of B. To see this, take a particular morning on which our man wakes up. In supposing this to be the effect of the impending ring, we suppose that if it were not going to ring, he would not have woken up. But the mere fact that we can intervene shows that this conditional is false. Because he has woken up and it might not ring, we know that he would have woken up whether it was going to ring or not. And to know this is to know that his waking is caused by something other than the impending ring.

By this 'bilking' argument, then, the causal order can never be the reverse of the temporal order (see Flew, 1954). And this shows, because of the way that perception requires a causal relation between belief and fact, that we cannot perceive the future. If we are impressed by the argument, then, we shall insist that one cannot really see that it will rain or that the ball will go out. Rather one sees the present weather conditions and trajectory, and reasons inductively to the future rain and the future path of the ball.

13.2 TWO CONCEPTIONS OF THE FUTURE

As before, realists and anti-realists offer competing accounts of the future. The anti-realist in this area takes it that our understanding of future-tense sentences is gained entirely in circumstances which we learn to take as evidence for the truth of those sentences, and that realist suggestions that we can subsequently come to an understanding of what it is for such sentences to be true in the absence of the sort of evidence concerned are fanciful. There are not, in our original understanding of such sentences, two distinct elements to be separated in this way. The sentence is assessed as true or false according to present and past evidence. If, then, there is no evidence either in favour of or against the sentence, it is neither true nor false. Most sentences about the future are in this category. Therefore most sentences about the future lack a truth value.

There is something very appealing about this position. The future has not happened yet, we feel, and hence there is nothing for future-tense sentences to describe and nothing about the future which can make such sentences now true or now false. The openness of the future is not an epistemological defect, there being an enormous number of facts about the future which we can have no hope of knowing (yet). It is a metaphysical necessity; and this distinguishes the future from the past. The past is complete, and its nature makes

past-tense statements determinately (though perhaps not determinably) true or false. We cannot say the same about the future.

If we find this sort of contrast between past and future compelling we are tempted by anti-realism about the future and by realism about the past. A more complete anti-realism would reject the contrast, and say much the same about our conception of the past as about that of the future. Leaving aside the argument about the past, we should consider one argument in favour of some form of anti-realism about the future. This is the claim that a realist about the future cannot allow that we have free will.

There are arguments against the possibility of free will that are independent of the present controversy. The causal determinist argues that every event and every action is caused by other events, and therefore that in no case is it true that we were able to act otherwise than we did. For what it is for one event to cause another is (or includes) that given the first event and the attendant circumstances, the second could not have been otherwise than it was. A free action is one which could in the circumstances have been otherwise than it was. So if all actions are caused, no action is free.

This determinist argument is disputed by the compatibilist, who holds that an action can be both caused and free. The argument we are concerned with, however, is that of the fatalist, not the determinist. The fatalist supposes somehow that it is already fixed what will happen tomorrow, but not for causal reasons. His reasons are those of the realist, who takes every future-tense sentence to have a truth value. A fatalist might reason as follows: either I shall die of cancer caused by smoking or I shall not. If I shall, then there is no point in giving up smoking. If I shall not, there is no point in giving up smoking. So there is no point in giving up smoking. This reasoning starts from the realist assumption that every statement about the future has a truth value; either true or false.

With less caricature, we can present the fatalist argument against free will as follows. It is already true either that I shall go to that restaurant this Friday or that I shall not. If it is already true that I shall go, how can I be said to have the option of not going, and vice versa? It looks as if the matter is already fixed, and nothing I could do could change it. Whatever the truth is is unavoidable, and the rest is impossible.

There is sufficient weight in this argument, despite its appearance of verbal trickery, to have persuaded Aristotle to abandon realism about the future (see Aristotle, 1962, ch. 9). Such a move reaps the usual advantages against the sceptic. For the anti-realist holds that

what it is for a statement about the future to be true is not to be separated from what it is for us to have the best possible evidence that it is true. Therefore there is no possibility that we have the relevant evidence but that the statement is false. The realist complaint that it is always possible, even given the evidence, that the statement be false is dismissed as inconceivable. So anti-realism about the future is very attractive. The difficulty is how to adopt it without adopting a general anti-realism. We have, after all, already argued against anti-realism in the theory of perception.

There is however a more severe sceptical argument in this area, which attacks in general the very notion of evidence that the anti-realist is relying on here. This is Hume's question about induction in general. In considering it we move away from exclusive consideration of knowledge of the future to a more general interest in our knowledge of the unobserved.

13.3 HUME AND HIS CRITICS

Hume's questions about inductive reasoning have already been raised in 1.2; they have lain dormant for a long time, and are now reawakened. We shall not repeat them, but begin immediately to consider possible answers to Hume.

Is the Circularity Vicious?

Hume complains that the only plausible attempt to justify our use of inductive inference involves a vicious circle, because it appeals to experience to justify appeals to experience. But some philosophers have been tempted to maintain that, though there is something like a circle in an inductive justification of induction, it is not vicious (e.g. Black, 1954, ch. 11, and 1958). The suggestion is that an argument such as:

Inductive reasoning has proved reliable in the past. Therefore inductive reasoning is (generally) reliable.

has as its conclusion a statement that the principle of inference which takes us from the premise to that conclusion is reliable. But this is no form of circularity; there is a crucial difference between principle of inference and statement. There would be a circularity if the argument really required as a premise the proposition which also stands as conclusion. But it doesn't require this. There are only

two reasons for thinking it does, both of which are mistaken. First, we might say that the argument is inconclusive otherwise. To make it conclusive, add the relevant statement to the premises; but then the argument is blatantly circular. The reply to this is that the argument is a perfectly sound instance of inductive reasoning as it stands. There is no need to add further premises in order to make it deductively valid, and hence conclusive. In general, it can be no complaint against inductive arguments that they are not deductive; the justification of induction is not the attempt to show all inductive arguments to be covertly deductive.

Second, we might think that generally a statement of the principle of inference on which an argument relies should, for full explicitness, be inserted as a premise. But this suggestion leads to infinite regress. For the resulting argument will depend upon some principle of inference, which will therefore need to be inserted; and the resulting argument will depend upon a further principle, etc., etc.

The argument is therefore held not to be covertly circular, because in no sense is the conclusion needed as one of the premises. An argument can therefore establish the reliability of its own principle of inference, when, as above, its conclusion asserts that reliability. We can justify induction inductively.

This is ingenious, but unsuccessful. To see why, we should try to look at it in Hume's way. Does the argument above give me reasons to accept its conclusion? I can only take it as doing so if I already accept, on independent grounds, the principle of inference on which it relies. Hence the argument could never give me a reason to accept its conclusion if I did not have sufficient reason to do so already.

Appeals to Analyticity

Hume's question is how we can have any reason for supposing that past and present observation provides us with evidence from which we can infer inductively. In the absence of such a reason, he insists, there can be no such thing as evidence, no such thing as having inductive reasons for belief.

A classic response is to hold that it is not possible to question whether past and present observations constitute evidence or provide reasons for further belief (see Edwards (1949) and Strawson (1952) ch. 9.2). The statement that observation does constitute evidence is true because of what we mean by 'evidence'. Someone who doubted whether the observed orbit of a planet were evidence about its future orbit would show by this that he didn't understand the

meaning of the word 'evidence'. There is no possibility that we might somehow be wrong about what 'evidence' means, and hence no possibility that observations should fail to provide evidence about the unobserved.

This is called the analytic justification of induction, because it amounts to holding that the statement 'the observed past is evidence for the future' is analytic; it is true solely in virtue of the meanings of the words in it. As such, of course, it is not directly available to those who, with Quine, reject the notion of analyticity and the traditional contrast between analytic and synthetic. But there is a further argument against the analytic justification, which can appeal to Quineans and to others alike; it is due to Urmson (1953).

This further argument starts from the remark that to call something evidence is in part to evaluate it, to see it as a reliable guide, which we are justified in following. Now terms that are used for evaluation in this way have a peculiar characteristic (though maybe this characteristic is present less noticeably in other cases), one which can be most easily seen in the case of the most general term of approval, 'good'. We learn the use of this term by appeal to such examples as our mentors (parents, scout-masters) offer us. But we are not thereby restricted to approving of only or all the things they approve of. However this may be, our understanding of the term 'good' can cast off entirely all reliance on the original examples. We can come to approve of a radically or even completely different set of objects or none at all, without thereby showing that we have forgotten what we were taught. And the same is true with the word 'evidence'. Our understanding of that word, gained no doubt in ordinary circumstances, is such that we can without abusing it come to take different sorts of things as evidence, or even to wonder whether anything is really evidence for anything else. And this is of course just what Hume did. It appears then, that the analytic justification does not succeed in ruling out Hume's question as incoherent.

It may be objected that we have missed the force of the analytic justification. It is possible that the analytic justification is intended as an anti-realist answer to a realist question; and that we have only rebutted it in realist terms, thus missing the point. Hume perhaps supposes that there are matters of fact about the unobserved past and future, about which our accumulated experience is normally taken to be inconclusive but relevant evidence. He then points out that experience can give us no reason for supposing that we can cross the gap between observed and unobserved. The anti-realist

reply is to insist that there is no gap between the facts of the matter and those propositions which we take to be relevant evidence. Our understanding of what it is for propositions about the unobserved to be true is tied indissolubly to the sorts of consideration which we take to be relevant evidence. The reason why it is an analytic truth that the observed past is evidence for the future, then, is the general claim that our concepts of truth and evidence go together. For a proposition about the future to be true is just for there to be evidence available that it is true; to understand what it is for a statement to be true is just to know what to count as evidence that it is true. Therefore we cannot suppose that we are wrong in taking observations to be relevant evidence; that they are relevant evidence is determined, not so much by the meaning of the word 'evidence', but by the meaning of the propositions for which they count as evidence.

This anti-realist version of the analytic justification of induction seems more formidable. To make a proper response to it we need to approach the matter from a different direction.

13.4 GOODMAN'S NEW RIDDLE OF INDUCTION

Hume posed his question in terms of a practice of inferring from observed regularities to the probable continuance of those regularities. He argued that this practice could not be justified, if justification requires reasons for thinking such inductive reasoning to be reliable. But he did not therefore conclude that inductive *practice* is irrational. He suggested that human nature is such that we acquire a habit of expectation after observing a sufficient regularity in nature. We have no reason to reason inductively, but we cannot help it. Given that we understand what it is to be rational not by reference to what we ought to do (and commonly fail to do) but by reference to what we do do, inductive inference is rational practice but not one grounded in reason.

Nelson Goodman argues that Hume's solution really raises even larger and harder questions of a similar type (Goodman, 1973). Hume's answer is that observed regular patterns create in us a habit of expectation. We so often observe objects falling when released that we naturally expect the next one to do so. And correct inductive inference is defined in terms of inference to events similar to those observed. The most reliable inference is supposed to be the one whose conclusion suggests that the world will go on in the way

most similar to its course up to now. But Goodman claims convincingly that this appeal to similarities conceals an assumption which it is hard to justify.

Suppose that up to now all observed emeralds have proved to be green; canons of inductive inference, as understood above, enjoin us to infer that the next emerald will be green and, with less probability, that all emeralds are green. But we can interrupt this cosy story by supposing there to be another predicate 'grue', with the following sense: an object is grue at a time t iff it is green and t is before 1 January AD 2000, or is blue and t is after 1 January AD 2000. Given this predicate, all our evidence that future emeralds will be green is equally evidence that they will be blue. For the emeralds we have observed are no more green than they are grue. It is as true that they are green as that they are grue, and so we are as justified in concluding that emeralds will be blue after t as that they will be green.

What this means is that correct inductive inference cannot be characterized in terms of inference to the continuation of previously observed similarities. We are inclined to suppose that one of the inferences about the emeralds is correct, and the other not (and of course there are infinitely many more such inferences, using predicates such as 'gred', 'grellow' and worse). But we have not yet provided any reason for preferring one to another. Nor have we any workable account of what a correct inductive inference is, by appeal to which we could hope to show that the use of such inferences is justified.

The natural response is to protest that there is something deeply suspicious about artificial predicates like 'grue'. But there is no agreement on what is wrong with them. (In a way the new riddle is simply the question what is wrong with them.) The mere fact that they are artificial is no complaint; it merely shows that we don't use them, not that we ought not to.

One common answer is to suppose that no 'sound' predicate could contain a reference to a particular point in time (or space). 'Grue' is defined in terms of greenness before t , and is therefore unsound. But there are two things wrong with this. First, it is equally true that 'green' is defined in terms of grueness before t ; an object is green at t iff it is grue at t and t is before 1 January AD 2000, or bleen and t is after 1 January AD 2000 ('bleen' is defined as grue is, except that 'blue' and 'green' are reversed). So each predicate appears unsound, on this criterion, from the point of view of someone using the other. Second, even if such predicates were to

be called 'unsound', we have not yet been told what it is about such predicates that makes inferences concerning them unreliable.

Instead of pursuing further attempts to solve Goodman's new riddle, we end by relating it to previous concerns. First, the anti-realist answer to Hume above offers nothing that will solve Goodman's riddle. The anti-realist appeals to the 'fact' that we do use past observations as evidence for future cases, or that our understanding of what it is for something unobserved to be true is linked to what we take as evidence that it is true. Our understanding of what it is for something to be green in the future is linked to what we take as evidence that it will be green; given the meaning of the relevant statement, then, there is no chance that we are wrong in what we take to be evidence for it. But these remarks do not discriminate between 'grue' and 'green'. The only hope that they would do so rests upon the thought that we have been using the predicate 'green' and not the predicate 'grue'. But this thought is of dubious usefulness, for two reasons. First, it is not clear what difference it makes which predicate we have used; we are not going to be able to show the other one to be somehow invalid on that account. Second, and more important, what about our previous practice makes it the case that we have been using 'green' rather than 'grue'?

It is true that we have been using the word 'green'. But what shows that we have not been thinking in terms of grueness all along? What shows that we won't suddenly begin to call blue objects 'green' on 1 January AD 2000? Until we can say which concept we have been using, we cannot hope to argue in favour of one on the grounds that it is the one we use.

This is important because an attractive reason for discounting concepts like grueness is that they cannot be acquired from examples (see Small (1961) for this suggestion). Greenness is a concept we can learn from examples, as we know because we have done it. But this argument collapses when we confront the possibility that we have been using the concept of grueness all along. If we have, we must have acquired it from our experience of grue objects.

The question we have now reached, and the attempt to answer it by appeal to examples, should remind us forcibly of Wittgenstein's rule-following considerations (5.5 – 7). Goodman's question is and should be similar to Wittgenstein's, for Goodman is asking what justifies one way of going on rather than another; this was exactly the question Wittgenstein was considering.

13.5 COHERENTISM AND INDUCTION

The position now is that we have two forms of inductive scepticism to deal with, Hume's and Goodman's. Anti-realism provided some sort of answer to Hume, as we saw, but completely failed to cope with Goodman's new riddle. We now consider the claims of coherentism.

Coherentists make large claims for the ability of coherentism to provide a perspective within which inductive scepticism collapses. Ewing held that "the coherence principle provides the only rational justification for induction" (Ewing, 1934, p. 247); Blanshard agrees (Blanshard, 1939, vol. 2, pp. 504–5). There are two constraints within which these claims should be assessed. The first is our present preference for internalism; externalist answers are insufficient. The second is that the answer should be as effective against Goodman as against Hume.

Suppose that we have a succinct statement of an inductive principle of inference, IPI. An externalist move would be to say that the adoption or use of IPI results in increased coherence in one's belief-set; this justifies our use of IPI. An internalist would add to this that we believe that the adoption or use of IPI results in increased coherence, and that so long as this belief is true we are justified in the use of IPI. Internalism does not have to show that this true belief is justified, as we agreed in considering the degrees of internalism in 9.3; to do so would lead to regress, and also here take us straight back into Hume's arms.

But still we have the question whether it is true that the adoption of IPI always results in increased coherence. It is no use trying to show this by appeal to the past; this would involve a well-worn circularity. The question whether the use of IPI leads to an increase in coherence cannot be entirely empirical, therefore. We need a reason in advance, a conceptual link between inductive inference and an increase in explanatory coherence.

Coherentists can provide such a reason. They do this by maintaining that the guiding principle of inductive inference just is "inference to the best of competing explanations" (Harman, 1970, p. 89). For instance, a detective inferring the guilt of one suspect from the evidence is reaching that hypothesis which provides the best explanation of all the evidence. So it is no accident that induction leads to an increase in explanatory coherence. And because both truth and justification are seen in terms of coherence, we can

say that the use of induction must take us nearer the truth. What better justification of induction could we hope for?

And there is more to it than this. We have not yet explained Ewing's claim that coherentism is the *only* successful position in this area. To see the basis for this claim, we need to retrace Hume's steps.

How can we know or be justified in any belief about what will happen next? We see the brick hurtling towards the window; our natural belief about the future is reached by inductive inference. What justifies this inference? There is no necessary connection linking the two events of this brick's flight and the destruction of the window. There is no contradiction involved in supposing that one happens and the other doesn't, which there would be if they were connected by logical necessity; and there is no other comprehensible notion of necessity than that of logical necessity. That being so, these two events are conjoined but not connected. Our inference from one to the other must therefore derive from experience of similar conjunctions in the past. And this sort of inference cannot be justified empirically without circularity (see 1.2).

Coherentists such as Ewing and Blanshard stop this train by denying a crucial step. They hold that there is another comprehensible notion of necessity, natural necessity, which links the individual events together; and maintain that anyone who denies this is doomed to inductive scepticism, so that their view offers the only hope.

The sort of natural necessity they are talking about is basic to the possibility of explanation; without it explanation is impossible. For to explain something is to see why, in the circumstances, it had to happen. Philosophers in Hume's tradition, for example Hempel (1942), take explanation to occur when we are satisfied that the relevant event is or was going to happen. But this is not enough. Unless we can see *why* it is going to happen, we have no explanation. And to see why, we need more than knowledge that in previous similar cases this sort of thing happened. That sort of knowledge may enable us to make bleak predictions, but does not help us to an explanation; for explanation we need understanding (cf. 11.1). Hume's argument is that since there is no such thing as a necessary relation between events, we are reduced to the sort of explanation he can offer. But the reply is that if explanation is possible at all, there must be necessary relations between events. In the simple example of the brick and the window, it isn't just that the window will break; it has got to, or is bound to. Our inductive knowledge

that it will break is knowledge by inference. The passage of the brick *entails* the breaking of the window; for, given the brick, the window must break. And we know this because it would be far harder to explain the window's not breaking than it would be to explain its breaking. (There are in fact links of mutual explanation.)

Ewing and Blanshard claim to escape Hume's argument by denying his atomism, the view that individual events are conjoined but never connected. It is the atomism that creates inductive scepticism rather than the nature of the case.

But this still leaves us with Goodman. Goodman points out that as well as our own inductive practice there are infinitely many others, each enshrined in its own language; and there seems to be nothing to choose between the languages. Each Goodmanian practice is equally inductive. So if one is justified, all are. Let us suppose for the moment that the previous argument about Hume leaves us with the conclusion that our inductive practice is justified. Why should the fact that other practices would also be justified disturb us? After all, the fact that many other belief-sets would be as justified as our own does not disturb us; justification can be shared. The plurality objection (8.2) only gets its bite when we consider claims to truth rather than justification. So isn't the sort of plurality Goodman points to equally acceptable?

Goodman needs to show that since other practices are possible, ours is not justified. If ours is justified, they are too. And the crucial point is that for each belief which we are justified in holding, there will be another practice, as well grounded as our own, which recommends the opposite belief. For instance, use of the predicate 'grue' leads us to conclude that emeralds are not green but blue after 1 January AD 2000. This means that the two practices are not just different but competing. They compete; nothing justifies the adoption of one rather than the other; so nothing justifies the belief that emeralds will stay green rather than the belief that they will not.

At this point we should turn for help to the analogy with Wittgenstein. What we have found is that there is no internal feature of our practice which can render it more justified than any competitor. And there is no external feature that will do this either; the only relevant feature was that our practice is the one we use, and the others are not. This doesn't help much. But we should not despair. The conclusion of Wittgenstein's thoughts about rule-following was that a practice does not need an independent external ground to justify it (5.6). The reason why we look for an external

ground is that we think that such questions are external questions; they ask about our practice from a point of view which purports to be outside it. But they are not. Just as questions about the justification of some part of our use of language, for example about our use of ethical terms, are questions asked within our linguistic practice; so questions asked about induction, which is an all-pervading form of life which seems almost to constitute rational behaviour for us, are still questions which we are asking and setting within our scheme of things. Admittedly they cannot be justified from outside that scheme. But this should not disconcert us. If we insist that questions about justification can only be asked from within our practice, we shall be less ashamed of the answers we are so tempted to give to them.

If this Wittgensteinian answer to Goodman is effective, we learn that Goodman's move beyond Hume was a move from an answerable question which the coherentist can answer to an unanswerable question which nobody can answer, in the sense in which it was intended, but which guarantees its own unanswerability and thus rules itself out of court. If we can put Goodman to one side, Ewing was right to claim that coherentism is the only position from which induction can be shown to be rational practice.

FURTHER READING

Hume (1955, ch. 4.2) and Goodman (1973, chs. 3–4) are two classic arguments.

Responses to Goodman include Barker and Achinstein (1960), Small (1961) and Quine (1969, ch. 5).

Swinburne (1974) is a good collection of responses to Hume, including those of Edwards, Russell and Black; the introduction is a helpful survey.

The 'bilking' argument against backward causation is in Flew (1954). Dummett (1978, ch. 18) shows greater inventiveness.

Cahn (1967) is a helpful discussion of Aristotelian and other attempts to come to terms with fatalism.

A good expression of determinism is Van Inwagen (1975). Hume's compatibilism is in Hume (1955, chs. 7–8).

There is an argument about whether Hume's question about induction should be seen as bleakly sceptical, rather than as an attempt to reduce an incorrect theory of rationality to absurdity. For this, see Stroud (1977, ch. 1).

Harré and Madden (1975, chs. 3–4) dispute Hume's conclusions. The coherentist response to Hume is given in Ewing (1934, ch. 4.3) and Blanshard (1939, vol. 2, pp. 504–11).