

Chapter 5 Reliabilism, Induction, and Scepticism

5.1 Introduction

At the end of chapter 3 I pointed out that the teleological theory of representation has radically anti-verificationist consequences. The contents of belief are fixed by their consequences for action, not by the circumstances that lead believers to adopt them. So it is perfectly possible that a judgement should have a given truth condition, and yet human beings be systematically prone to form this belief when it is false.

Such a realist¹ account of representation might be thought to open the door to scepticism: if truth-conditions transcend evidence, then what assurance do we have that our beliefs are free of error? In this chapter I want to show how this sceptical threat to knowledge can be met from within a realist perspective. Accordingly, in what follows I shall take the implications of chapter 3 as read, and assume without further argument that judgements about the natural world answer to non-verificationist truth conditions. My focus here will rather be on the notion of knowledge, and on how a proper understanding of this notion enables us to give an adequate response to scepticism, even within the framework of a realist theory of representation. At the beginning of the next chapter, however, I shall return to the general debate between realism and anti-realism, and compare my overall realist attitude to representation and knowledge with the anti-realist alternative.

In more detail, the plan of this chapter is as follows. Sections 5.2-8 will defend a reliabilist account of knowledge. Such reliabilist theories are nowadays fairly widely accepted; but a distinctive feature of my defence will be its appeal to the point of the concept of knowledge, rather than to intuitions about test cases. In sections 5.9-13 I shall then show how this reliabilist account of knowledge provides an answer to the traditional sceptical problem of induction. Sections 5.14-17 will then generalize this answer and address some other arguments for scepticism.

5.2 Knowledge and the Project of Enquiry

Let me start with a question raised in Chapter 2 of Bernard Williams' *Descartes* (1978). Williams asks: why do human beings want knowledge? He takes it as given, as I shall, that humans want true beliefs. But, as we all know, a belief can be true and yet not be knowledge, as when it is a mere hunch or some other lucky chance. So the point of the question is: why do we want our beliefs to be known, in addition to being true?

Williams' answer goes as follows. Human beings are prone to false beliefs. So, if our desire for true beliefs is not to be idle, we will need to exercise ourselves to achieve it. It is no good, however, to start checking through all your beliefs with the intention of discarding the false ones. To have a belief is to take that belief to be true. So once you have formed your beliefs, internal inspection will not serve to distinguish the true from the false ones. Insofar as you are prone to error, the damage will already have been done.

The only effective way for us to ensure that our beliefs are true is to block the error at source, by bringing it about that the processes by which we acquire beliefs in the first place are ones that generally yield true beliefs. So Williams argues that the desire for true beliefs itself generates the desire that our beliefs should issue from processes that generally produce truths. And then, finding it independently plausible that beliefs

produced by such processes should count as knowledge, Williams has an answer to his original question as to why we should want knowledge: our desire for knowledge derives from our desire to avoid error, in that attaining knowledge is the only effective means by which humans can avoid error.

I want to draw something more ambitious from this analysis. I think that, in addition to explaining why we should want knowledge, Williams' story also shows us what knowledge is. Williams takes it as given, from outside his analysis, that beliefs generated by truth-producing processes will count as knowledge. But I think that his story also explains why we have this concept of knowledge, why we pick out beliefs generated by a truth-producing process as knowledge, as an especially good kind of belief. My idea here is that our concern to avoid error makes us especially interested in the state we need to get into as a means to avoiding error, and that this is why we call that state "knowledge" - the state, to repeat, of having acquired a true belief from a process which generally produces true beliefs.²

5.3 Certainty and Reliability

The above remarks prompt an immediate question: how truth-productive does a belief-forming process need to be in order to be an effective means of avoiding error, and therefore to qualify as a source of knowledge? In particular, is it enough that it merely be reliable, in the sense that it generally delivers true beliefs as a matter of contingent fact in this world? Or does it need in addition to yield certainty, in the sense that it should be impossible for a belief issuing from that process to be false?

Much traditional philosophical thinking assumes that knowledge requires certainty. But from the point of view of my remarks in the last section it is not clear why certainty should be necessary. Knowledge, I have suggested, is the state that we need to get into if we are to succeed in avoiding error. But we will succeed in this aim as long as we have belief-forming processes which are reliable in this world. That such processes would lead us astray if things were different does not mean that they will lead us astray, as things are. This line of thought suggests that the traditional demand for certainty may be a mistake, perhaps fostered by an over-optimistic view of what human thought can achieve, but inessential to knowledge itself.

I shall return to the idea that the demand for certainty may be a mistake in section 5.5 below. But first, in the rest of this section and the next, let me say a bit more about the contrast I have drawn between reliability and certainty. Note that I have defined certainty objectively, rather than psychologically: the issue is whether it is in fact impossible for a given belief-forming process to produce a false belief, not whether the subject is aware of this, nor whether it yields some feeling of absolute security. It is this objective notion that matters to the arguments of this chapter. However, there are obvious links between it and subjective requirements on knowledge. For, as Descartes so forcefully argued, the only plausible source of certainty in the objective sense derives from various operations of the conscious mind -- in particular, from introspection and intuition. And so, if we can achieve knowledge with objective certainty, then we will also, as it happens, have "subjective warrants" available, in that we will always be able to tell introspectively that our knowledge has come from these putatively infallible conscious sources.

ÊÊ From the reliabilist point of view, by contrast (henceforth I will use "reliabilism" to mean the view that only reliability is required for knowledge), any subjective

requirements on knowledge are gratuitous. For, in order for a belief-forming process to be reliable, there is no need for its reliability, or even its existence, to be available to consciousness. According to reliabilism, we will know, say, that there is a table in front of us, just in case the unconscious visual processes that give rise to such perceptual beliefs generally deliver true beliefs, whether or not we are aware of this. There is therefore no pressure, given reliabilism, to reconstruct such perceptual knowledge as first involving some infallible introspection of some sensory idea, and then some intuitively compelling inference from this idea to the presence of a table. If the demand for certainty in knowledge is unmotivated, then so too is this reconstruction of perceptual knowledge as involving infallible inferences from infallible introspections.

Ê
Ê

5.4 Knowledge and Normativity

It is sometimes felt that reliabilist epistemology changes the subject. As I have just pointed out, reliabilism implies that whether or not we know will often hinge on matters, such as the reliability of some visual process, which lie quite outside our consciousness. But this seems to imply that we are at the mercy of nature, that we cannot do anything to affect whether or not we know. And this then makes reliabilist epistemology seem a quite different subject from the traditional version (henceforth I shall use "traditional" to refer to views according to which knowledge requires certainty or subjective warrants³). For surely a central concern of traditional epistemology was the normative question of what we should do in order to ensure that our beliefs are knowledge.

Ê However, this reaction to reliabilism involves a fallacy. It is true that traditional conceptions of knowledge offer advice on how to achieve knowledge: roughly, you should consciously monitor your thought processes, and avoid any which are not necessarily infallible. And it is true that reliabilism does not concur in this advice. But this is not because reliabilism has stopped offering advice on how to know, but simply because reliabilism offers different advice.

Ê Where traditionalists advise aspirant knowers to monitor what goes on in their conscious minds, reliabilists will simply advise them to take whatever steps are needed to bring it about that their beliefs come from reliable processes. Such steps may well call for us to influence processes which lie outside consciousness, but that is no reason to conclude we cannot succeed. After all, most of the things we influence lie outside consciousness, like our environments, our physical health, and so on. Similarly, there is no reason why we cannot influence non-conscious aspects of our belief-forming processes, by such means as rote learning, adjusting the working of instruments we rely on, and so on⁴. (Reliabilists will allow that conscious monitoring is one way to improve the reliability of our belief-forming processes. But it is not the only way.)

Ê To guard against a possible misunderstanding, let me emphasise that I am not suggesting that it is a requirement on knowledge that knowers must take active steps to bring it about that their beliefs are knowledge. I am aware that I began this chapter by identifying the concept of knowledge as the state someone concerned to avoid error (a "concerned enquirer" henceforth) wants to get into as a means to avoiding error. But it does not follow from this that the only way to be a knower is

to take active steps to get into that state. Either you may already be in the requisite state, not because you did anything to make your belief-forming processes reliable, but simply because they were reliable to start with. Such passive knowers will already be in the state concerned enquirers aim to get into, even though they do not themselves share the concern to get into that state.⁵

Ê

Ê

5.5 Rationales versus Intuitions

I have just argued that reliabilism does not abandon the traditional normative issue of how best to acquire knowledge. In this section I want to consider a rather different argument for thinking that reliabilism changes the subject. This appeals, not to considerations of normativity, but directly to intuitions about knowledge. Many philosophers take it to be intuitively obvious that subjective warrants are part of the concept of knowledge. And so they conclude that reliabilism, which dispenses with such requirements, must be wrong.

Ê My response to this is that intuitions are not the only way to evaluate a theory of knowledge. There is a vast contemporary literature which aims to decide between reliabilism and traditional theories solely by appeal to intuitions about ingenious test cases.⁶ Unfortunately, however, these intuitions weigh on both sides, and the literature based on them is notoriously indecisive. By contrast, my approach in this paper has not appealed to intuitions, but has tried to identify an underlying concept of knowledge, by locating the role it plays in our thinking, by trying to understand why knowledge is such a matter of concern to human beings.

Ê My suggestion has been that knowledge is tied up with our desire to avoid false beliefs: it is the state a concerned enquirer needs to get into as a means to achieving this desire. Given this identification of the concept of knowledge, so to speak, we can then investigate more detailed conceptions, or theories, let us call them, of what that state is, more detailed theories of exactly what state an active enquirer needs to get into as means to avoiding error. The theory I am defending is that the requisite state is acquiring a belief from a reliable processes. The theory that certainty and hence subjective warrants are required is a different theory, but still a theory, in the terminology I am using, of the same concept of knowledge.

Ê This is why I think that any intuitions which may favour the traditional theory over the reliabilist alternative are indecisive. If I have shown that the reliabilist theory is the right theory, in that acquiring a belief from a reliable process is indeed what is needed as the means to avoid error, then I have therewith shown that the traditional theory and the intuitions that support it are mistaken.

Ê Perhaps there remains a gap here. Suppose it is granted that I have identified the underlying concept of knowledge successfully, and have shown that reliabilism is the theory of knowledge that best fits it. An obvious question which then arises is why anybody should have had contrary ideas about certainty and subjective warrants in the first place. For, as I have told the story, reliabilism, as opposed to the traditional theory, follows pretty quickly from the concept of knowledge, thus making it mysterious why anybody should ever have thought anything more was needed. We cannot simply rest with the suggestion, offered briefly in passing earlier, that traditional ideas about knowledge may have been fostered by excessive optimism

about what can be achieved. For the possibility of achieving something does not explain why we should want it, if it is not already desirable.

I shall return to this issue in section 5.8. But first let me comment briefly on the similarities between the strategy outlined in this section and that defended by Edward Craig in "The Practical Explication of Knowledge" (1986). At a detailed level, Craig's views differs from mine: he offers a third-person account of the concept of knowledge, by contrast with my first-person account, arguing that knowledge is the state our informants need to be in, for us to avoid error, not the state that we ourselves need to be in. But at the level of general strategy, Craig and I are in accord, in that he too seeks to offer an account of the point of the concept of knowledge, and to use this account to explain the nature of the concept, rather than simply trying to identify the concept from intuitions alone.

As to our differences, my objection to Craig's line would be that he is in effect focusing on the special case in which we succeed in avoiding error by acquiring beliefs from informants who succeed in avoiding error. I accept that this special case may well have been of primary significance in the historical development of the everyday concept of knowledge, in that worrying about your informants' reliability calls for rather less sophistication than worrying about your own reliability (cf. Craig, *op cit*, p 215). But, even so, Craig's third-person focus seems to me to have the disadvantage of cutting the link with the traditional normative issue of what we should do to avoid error. That is, even if it is unfaithful to the history of the concept of knowledge to view the desire for good informants as a special case of a general desire to have good belief-forming processes, I would argue that the more general perspective I have adopted nevertheless has the advantage of showing how the concept of knowledge relates to familiar philosophical worries about knowledge.

Ê
Ê

5.6 Knowledge and Percentages

How much reliability should a reliabilist require for knowledge? I shall consider two dimensions to this question. First, I shall consider whether we need 100% reliability, or whether some lesser percentage, such as 95%, say, is enough. Second, I shall ask over what range of possibilities the relevant percentage is to be assessed.

On the percentage question, it would be a mistake to think that the rejection of certainty has already decided this question in favour of something less than 100%. For "certainty", as I have been using it, implies that a belief-forming process cannot go wrong, will deliver 100% true beliefs in all possible worlds. This is a much stronger requirement than 100% reliability in this world. So we can reject certainty and still uphold a requirement of 100% contingent reliability. On the other hand, even the latter seems a fairly strong requirement. So perhaps we should consider arguments in favour of some lesser percentage.

On the second question, about the range of possibilities, the argument so far has shown it would be a mistake to require this degree of reliability to hold up across all possible worlds. But, as we shall see, there may remain reasons for wanting it to hold up across at least some counterfactual situations.

Let me deal with the percentage question first. I shall return to the question about the range of possibilities in the next section. Up to a point, it is possible to by-pass

the percentage question. Suppose a given belief-forming process delivers beliefs which are true 95% of the time. Then the appropriate output from that process would not be a full belief in the first place, but a 0.95 degree of belief. After all, if you believe it is going to rain tomorrow on the basis of a 95% reliable method of forecasting, you would be ill-advised to bet a million pounds to a penny, or indeed to stake anything more than nineteen to win one, on its raining tomorrow. So if knowledge unqualifiedly requires belief, as I have implicitly been assuming throughout, then this in itself seems to call for belief-forming processes which deliver truths with 100% reliability.

Still, perhaps it is a bit quick to assume that knowledge requires strictly full belief. After all, in everyday discourse we certainly refer to beliefs of high, but less than strictly full, degree as "beliefs" simpliciter, and to that extent we should expect the notion of knowledge also to encompass sufficiently well-founded beliefs of high, but not full, degree. But, having said this, there is then an obvious answer to the question of how well-founded such a belief of high, but not full, degree needs to be to qualify as knowledge. For, under the present suggestion, everyday discourse has certain standards, perhaps varying from context to context, of how firmly a belief has to be held to qualify as a belief simpliciter. So why not simply incorporate those standards into our analysis of knowledge, and say that for a belief to be knowledge it should come from a process whose reliability is at least sufficient to warrant the degree of confidence required for the belief to qualify as a belief simpliciter in the first place, and not just as what even everyday discourse would consider as a partial belief?

It will be helpful for what follows to observe that, while it is certainly true that we often allow beliefs of less than strictly full degree to qualify as knowledge, there is also a practical sense in which it is always better to get a belief from 100% reliable processes. To put it simply, the reason is that such beliefs will then be true, and so decisions informed by them will succeed with probability one, whereas, if those beliefs came from less than 100% reliable processes, then the actions they informed would be less likely to succeed.

Actually, this puts the point rather too simply, in that even if your beliefs are proof against error, they may still not be informative enough to tell you how to achieve some result; and, even if your beliefs are informative enough, you may fail to draw the inference correctly. But these two caveats would apply equally even if the same set of beliefs came from less than 100% reliable processes, and so do not affect the point that it is always preferable, from the point of view of achieving your desires, to get a full belief from a 100% process, rather than a less than strictly full belief from a less reliable process.

It should also be admitted that in many cases the extra costs of getting 100% reliability will not be worth the extra probability of success, in which case we will do better to settle for a partial belief. This is no doubt why everyday discourse does not make strictly full belief a precondition of knowledge in general. But this merely calls for a yet further qualification, and still does not affect the underlying point that, when costs are equal, full belief from 100% reliable processes is always better.⁷

Ê
Ê

5.7 Nearby Possible Worlds

Let me now turn to the second question raised at the beginning of this section: what range of possibilities is relevant to the reliability of belief-forming processes? At first sight it might seem to follow from my overall argument that reliability in the actual world is all that matters. After all, as I observed earlier, reliability in this world is all that we need in order to avoid error. However, there are good reasons why knowledge calls for more than merely this-worldly reliability.

Let us return to the idea that knowledge is the state concerned enquirers need to get into in order to avoid error. It is true that concerned enquirers have no interest in reliability in non-actual worlds as such. Nevertheless, in acquiring processes which are reliable in this world, concerned enquirers will inevitably acquire processes which are reliable in a range of non-actual situations as well.

The reason is that, if you are an concerned enquirer, you will not be able to anticipate the future in enough detail to be able to tell exactly when you are going to use any given belief-forming process, and so will not know exactly which truths that process needs to deliver in order to be reliable in the actual world. Instead, you will inevitably have only limited knowledge about the general nature of the world and your particular situation in it⁸, information which will narrow down the range of circumstances you may in future find yourself in, but it certainly will not tell you exactly what they will be. So, in aiming for reliable belief-forming processes, you will inevitably be constrained to aim for belief-forming processes which will reliably deliver true beliefs across the entire range of possible circumstances that your current information leaves it open you may end up in. Since not all the possibilities in that range will become actual, you will inevitably be aiming to get into a state which would deliver true beliefs in various nearby possible worlds, as well as in the actual one.

So my overall approach to knowledge accommodates the requirement that knowledge should have a certain degree of counterfactual reliability. Note, however, that this is still a long way short of requiring reliability in all possible worlds, or even reliability in all causally possible worlds. For the information already possessed by concerned enquirers will still in general be enough to rule out the possibility of their being in most possible worlds (such as the world where you are manipulated by Descartes' evil demon) or even in most causally possible worlds (such as the world where you are a brain in a vat).⁹

Ê
Ê

5.8 The Attractions of Certainty

I promised to return to the question of why certainty should be intuitively plausible as a requirement for knowledge. A serious answer to this question would include an historical dimension, examining the development of Western epistemological notions, with particular reference to the mediaeval distinction between demonstration and opinion, to the seventeenth and eighteenth century struggles to find a place for the newly emerging scientific knowledge within this distinction, and, perhaps most important of all, to the religious dimensions which so animated the participants in these debates. However, any such historical investigation is beyond the scope of this book. Instead let me offer a possible philosophical explanation of the pull of certainty, not as a competitor to an historical account, but as a possible complement.

I have just argued that knowledge requires not only reliability in the actual world, but reliability across all worlds which are possible, relative to the information open to concerned enquirers. However, in discussing this issue of counterfactual reliability, I have so far implicitly been taking it for granted that our notional concerned enquirers are aiming for full beliefs from 100% reliable processes. But, as we saw earlier, in many practical contexts it will often be more efficient to settle for less than full beliefs, delivered by belief-forming processes of appropriately high, but less than perfect, reliability.

Now, an enquirer who was concerned to acquire such a less than full belief would be entitled to ignore, when assessing the reliability of the relevant belief-forming process, not only all worlds which are impossible relative to his or her current information, but also any worlds which fall below an appropriate threshold of probability relative to that current information. For clearly the fact that the process would be unreliable in such unlikely circumstances does not give a concerned believer sufficient reason to withhold a high degree of belief from its deliverances.

Add to this last point the consideration, elaborated in section 5.6, that, although it is often perfectly sensible to settle merely for a high degree of belief, it is always better, especially where it is important that your actions will not fail, to get full beliefs from 100% reliable processes. Putting these two points together, it follows that your knowledge will get better the more possible circumstances with any non-zero probability your belief-forming processes are reliable across.

This does not of course amount to a good argument for thinking that the best thing would be to acquire beliefs from processes that are reliable across all possible worlds. Even if it is a good thing to be reliable across all worlds with non-zero probability relative to current knowledge, this is a long way short of reliability across all possible worlds. For, as I have said, most possible circumstances will be downright impossible relative to the information available to any concerned enquirer. Nevertheless, one can see how it would be easy to slide, from the thought that you need ideally to guard against any possibilities that your information leaves with any non-zero probability, to the thought that you need ideally to guard against any possibilities whatsoever.

Ê
Ê

5.9 The Problem of Induction? What Problem?

I turn now to the problem of induction. Let us suppose, for the sake of the argument, that the general form of induction is simple enumeration. (I do not really think this is a good model for inductive inferences. But it will help the exposition to assume so for a while.) So, for example, from the premises, that N ravens have been black so far, we conclude that all ravens are black. Schematically,

Ê

(1) Ê Fa1 & Ga1

Ê .

Ê .

Ê FaN & GaN

Ê _____

Ê
ÊÊ All Fs are Gs

Ê
ÊÊ The traditional complaint about this form of inference is that it is logically invalid.Ê The conclusion does not follow logically from the premises.Ê It is logically possible that the premises be true but the conclusion be false.Ê For reliabilists, however, this complaint has no force.Ê Since a form of inference can well be contingently reliable without being logically guaranteed, reliabilists can simply respond to the traditional complaint by arguing that the illogicality of inductive inferences is no reason to deny that such inferences yield knowledge.

ÊÊ Perhaps it is worth pausing briefly to explain how the notion of reliability applies to inferences.Ê Though I have not treated this explicitly so far, the appropriate notion is obviously conditional reliability:Ê the conclusion should always be true in the actual world, if the premises are.Ê (This will then ensure, for a reliabilist, that reliable inferences will transmit knowledge, that they will yield known conclusions when applied to known premises.Ê For if the premises are known, in the sense that they are true and reliably arrived at, then any conclusion derived from a conditionally reliable inference will also be true and reliably arrived at, and so known.)

ÊÊ It is tempting to leave the problem of induction here, with the observation that the logical invalidity of induction does not mean its conclusions are not knowledge.Ê However, I suspect that most readers will be unpersuaded by this quick way with inductive scepticism, even if they are persuaded by the general arguments for reliabilism.Ê So in the next X sections I shall consider whether there are any further reasons why a reliabilist should worry about induction.Ê Accordingly, I shall now take it as given that reliabilism is the right account of knowledge in general;Ê the issue to be considered is whether any sceptical doubts about induction still arise within this assumption.

Ê
Ê

5.10Ê Is Induction Reliable?

One possible worry about the simple reliabilist response to the problem of induction sketched in the last section is that it seems little different from the "analytic justification of induction" proposed by Edwards (1949) and Strawson (1952, Ch. 9).Ê Yet it is now widely agreed that inductive inferences cannot be shown to be legitimate simply by observing, as the "analytic justification" does, that most people would characterize induction as a central case of "rational" thinking.Ê For such facts of common usage leave it open that there may be underlying requirements for a form of reasoning to be rational, which are not in fact satisfied by induction, are that most people may therefore be in error in holding induction to be rational.

ÊÊ However, the reliabilist response to induction is quite distinct from the analytic justification.Ê Reliabilists do not accept a form of reasoning as rational just because it is widely regarded as "rational", but only insofar as it satisfies the underlying requirement of reliably delivering truths.¹⁰ In particular, reliabilists will deem induction to be rational, and its conclusions therefore knowledge, not because it is called "rational", but because they believe that it is in fact a reliable method of getting new truths out of old ones.

However, this now points to an obvious problem. That induction reliably generates truths is itself a substantial contingent claim. Yet no support has so far been offered for this. We reliabilist friends of induction seem simply to be taking it for granted that induction is a reliable method of inference, and then concluding, in virtue of our general reliabilism about knowledge, that induction yields knowledge. But what basis do we have for the initial assumption that induction is reliable?

Some reliabilists are inclined to respond, at this stage of the proceedings, that we do not need to know that we know in order to know.¹¹ I think this is the wrong move. It is perfectly true, of course, that ordinary non-philosophical knowers do not need to know that they know. But the present demand for a defence of the claim that induction yields knowledge is not being made of ordinary knowers who are using inductions, but rather of us philosophers who are talking about inductions, and in particular about the question of whether inductions yield knowledge. We reliabilist friends of induction are explicitly claiming that inductive inferences yield knowledge because they reliably yield truths. Given this, it is perfectly reasonable for someone to challenge us to provide support for this claim.

Of course, if we fail to meet this challenge, this will not necessarily show that induction does not yield knowledge. To lack any grounds for accepting the reliability of induction is not yet to have grounds for denying it. But such a stand-off would be a failure for us friends of induction, and a success for our sceptical challengers. The point at issue is whether induction yields knowledge, that is, given reliabilism, whether induction reliably generates truths. We friends of induction say yes, our sceptical challengers ask for support for this claim. If we can't answer them, then they will have succeeded in showing we aren't entitled to our stance.

Ê

So to uphold induction as a source of knowledge we need to show that inductive inferences are reliable. However, now we are clear about this need, I do not think it is hard to satisfy. The obvious way to find out whether induction is reliable is to examine such evidence as bears on the matter. When people make inductions, do their conclusions turn out to be true? There are plenty of past examples of people making inductions. And when they have made inductions, their conclusions have indeed turned out true. So we have every reason to hold that in general inductive inferences yield truths. That is:

(2) When person1 induced, from N observations of A going with B, that All As are Bs, this conclusion1 was true

When person2 induced, from N observations of C going with D, that All Cs are Ds, this conclusion2 was true

Ê.

Ê

When personN induced, from N observations of L going with M, that All Ls are Ms, this conclusionN was true

Ê_____

Whenever someone induces, their conclusion is true.

Ê

Ê

5.11 The Legitimacy of Normal Methods

Let me first put to one side two obvious worries about the premises of this argument. First, aren't there plenty of past examples of unsuccessful inductions with false conclusions, as well as successful ones with true conclusions? Second, how can we know that even the successful inductions are successful, given that observation of the past will only show, for example, that As have been Bs so far, not that all As are Bs? I shall deal with both these points in due course (in section 5.17 and footnote 14 respectively). But for the moment it will be helpful to ignore them, and attend instead to the move from the premises to the conclusion of (2). For it will not have escaped the notice of most readers that this is itself an inductive inference, of just the kind whose reliability we are presently concerned to investigate.

However, is there anything wrong with this? It is not as if the discussion so far has identified some flaw in induction, of a kind which would imply that it ought to be eliminated from the battery of procedures by which we normally arrive at our beliefs. In particular, we have agreed that the logical invalidity of inductive inferences in itself casts no discredit on induction. Given this, when a certain question of fact is raised -- namely, are inductive inferences always reliable? -- what is more natural than to try to resolve this question by means of our normal procedures of investigation, which include, as it happens, our inductive procedures?

I know that to some philosophical sensibilities this will seem unduly complacent: surely we aren't entitled to any methods of investigation, until we have demonstrated their worth. But where are we supposed to start? We certainly need to begin with some methods of thought, lest we lapse into philosophical catatonia. Many philosophers, I realize, will want to follow Descartes, and restrict our initial methods to introspection and intuition. But Descartes' rationale for this restriction was that it promised certainty, and we have already agreed that this is an unnecessary desideratum on our methods of thought. It would seem equally sensible to continue with our normal methods of thought, at least until we uncover some reason to distrust them. And these methods will include induction, since, to repeat, we have not as yet been given any reason to distrust induction.

5.12 Varieties of Circularity

Still, even if nothing has as yet been shown wrong with induction in general, it may well be felt that there is something wrong with the inductive argument (2) in particular. For isn't (2) a circular argument, and therefore illegitimate? This objection needs to be treated carefully. It is true, as we shall shortly see, that circularity of a certain sort is present in (2). However, provided we keep firmly in mind the specific argumentative task to which (2) is directed, we shall also see that this circularity is not damning.

As a first step, we need to distinguish between "premise-circularity" and "rule-circularity".¹² An argument is premise-circular if its conclusion is contained among its premises. An argument is rule-circular if it reaches the conclusion that a certain rule of inference is reliable by using that self-same rule of inference. Clearly premise-circularity is a vice in an argument. The point of an argument is to take us from old beliefs, which we already accept as premises, to some new belief as a conclusion. But if the conclusion is already contained in the premises, then the argument will fail

in this primary task. However, argument (2) is clearly not premise-circular. It is a genuinely expansive argument, whose conclusion, that all inductions yield true conclusions, manifestly outruns its premises, that N inductions so far have done so.¹³

On the other hand, argument (2) is rule-circular. Even if the claim that induction is reliable does not appear among its premises, it does use an inductive inference to reach its conclusion that induction is reliable. I have a number of comments to make about the rule-circularity of (2). But first let me make a wider comment, not about argument (2) in particular, but about rule-circularity as such: namely, that it can scarcely be a general requirement, on all legitimate forms of inference, that it be possible to show that they are reliable in some non-rule-circular way. For this would disqualify even deduction as a legitimate form of inference. (While it is possible to demonstrate that deductive inferences are reliable -- indeed necessarily reliable -- by means of the standard semantic soundness proofs, these demonstrations themselves unquestionably employ deduction.) So the fact that induction can only be shown reliable in a rule-circular way, as in (2), certainly does not in itself yield any immediate reason to distrust induction.

But this is merely to repeat the point that we have as yet been given no good argument for distrusting induction. Our current concern, however, is whether (2) takes us beyond this, and gives us a positive basis for trusting induction, despite the fact that it is admittedly rule-circular.¹⁴ Let us recall the context of argument in which (2) was put forward. We agreed, on general reliabilist grounds, that induction does not need to be logically valid to yield knowledge, but will yield knowledge just in case it is reliable. However, the sceptic then pointed out, we cannot just take the belief that induction is reliable for granted. To which we responded that we are not taking this for granted, but have a good argument, based on empirical evidence, for the conclusion that induction is reliable, namely, argument (2). It seems to me that, in this specific context, the context of showing a sceptic who accepts reliabilism that we are not just helping ourselves to the belief that induction is reliable, (2) does just the job it is required to do.

Ê

5.13 Who Needs Persuading?

Perhaps the best way of showing this is by detailing some of the tasks argument (2) is not intended to fulfil. For a start, we should recognize that argument (2) would be no good for persuading people who do not make inductions to start making them. While the conclusion of (2), that inductions are reliable, would certainly be a good reason for such people to start inducing, if they accepted it, they clearly will not be persuaded to do so by (2), for the route from (2)'s premises to (2)'s conclusion requires just the kind of inductive inference that they eschew. In particular, then, (2) will be no good for persuading people who have already reflected on the reliability of induction, and have been persuaded, for whatever reasons, that they ought to stop performing inductions, that they ought to start again.

However, in the present context of argument, this is no demerit in (2). Argument (2) is not addressed to people who avoid inductions. We may yet discover good reasons for avoiding inductions, and indeed in the next few sections I shall examine some possible such reasons, but right now we are assuming that nothing has yet been shown wrong with induction, and are considering whether, given this, argument (2)

can show us whether induction is reliable. So (2) should be thought of as addressed to people who have not yet been given any reason to distrust induction. And (2) ought surely to persuade such people at least of its conclusion.

I realise that many readers will feel that, if (2) is a good argument in defence of induction, then it ought to be capable of persuading any intelligent being, with whatever epistemological habits. But this is an extremely strong demand, and it is not at all clear why we should accept it. The only plausible rationale, once more, seems to stem from the assumption that knowledge requires certainty, together with the assumption that the only kinds of belief-forming processes which can plausibly deliver certainty are conscious operations whose logical infallibility is introspectively available. Together these assumptions imply that any source of knowledge ought in principle to be recognizable as such by any conscious beings, in virtue of their introspective abilities; and hence these assumptions imply that a good argument for the legitimacy of some source of knowledge ought to persuade all people whatsoever, however wrong-headed their starting position. However, once we reject the assumption that knowledge requires certainty, then this whole line of reasoning falls away, and the strong demand that a good defence of induction ought to persuade any conscious being is left without any obvious means of support.

These last remarks bear on the question of "counter-inductive" arguments for "counter-induction". It is often observed that inductive arguments for induction, like (2), have counter-inductive mirror images. Counter-inductivists, when they observe that a number of As are all Bs, conclude that the next A will not be a B. When it is pointed out that this is illogical, they can respond, "So what? Illogical it may be, but this doesn't show that it is not in fact a good way of reliably reaching true conclusions". And when we say, "All right. But what basis do you have for supposing that counter-induction does in fact deliver true conclusions?", they reply, "Ah, so that's what you're worried about. Let us then look at the evidence that bears on the question. On a large number of occasions in the past people have counter-induced, and have been led to false conclusions. So we conclude -- counter-inductively -- that the next time we counter-induce we will get a true conclusion".

There is room to dispute whether this is in fact a perfect mirror image of (2) (Cf Van Cleve, *op cit*, footnote 16.) But let that pass. The more important point is that, even if counter-inductivists can mirror (2), this does nothing to discredit (2) itself. I have already conceded that (2) is not going to persuade people who do not make inductions to start making them. Counter-inductivism now simply gives us a further example of people who have abnormal inferential dispositions, and who will not therefore be persuaded by (2). Except that the parable of the counter-inductivists adds an extra twist, namely, that counter-inductivists will be persuaded, by their mirror of (2), to the conclusion that their abnormal counter-inductive dispositions are reliable. But all this leaves (2) untouched. We should not expect it to perform the impossible task of knocking imaginary non-inductivists out of their non-inductivism -- its task is only to allow normal people, like ourselves, to resolve the issue of whether induction is reliable.

By this stage, some readers may be feeling that argument (2) does not do very much. Indeed, if it only works for people who already make inductions, is it really doing anything at all? My answer is that it is not supposed to do very much. Nearly all the serious work was finished before (2) came on the scene. Most importantly,

the general arguments for reliabilism have already shown that the logical invalidity of induction is not a problem. Argument (2) is just supposed to show that, given that there is nothing problematic about induction, then there is no barrier to our concluding that it is reliable, and hence that it yields knowledge.

Trained philosophers naturally expect a "justification of induction" to do something to rehabilitate induction, in response to an argument that there is something wrong with it. But (2) is not meant as a "justification" in this sense. So we should not condemn it for its failure to be one.

5.14 The Strategy Generalized

As a first step towards generalizing the anti-sceptical strategy outlined in the last few sections, let us be more realistic about induction. I have already noted one way in which the above discussion of induction has involved an idealization, namely, in respect of the assumption that in our experience all past enumerative inductions have been successful. This assumption is of course manifestly false. There are plenty of good examples of enumerative inductions leading to false conclusions, from Russell's chicken who expected to be fed every day, to the Newtonian physicists who expected acceleration always to be inversely proportional to rest mass.

In any case, apart from such direct evidence, there is also a principled argument to show that simple enumerative induction cannot possibly be a reliable method of inference. I refer to Goodman's "new problem of induction". Goodman (1954) shows that there are far too many ways of classifying events, far too many As and Bs, for every instance of schema (1) to yield a true conclusion. Indeed, Goodman shows how to construct, for every instance of (1) that might yield a true conclusion, an infinity of other instances which will then yield false conclusions.

These are good arguments against enumerative induction. It is important to recognize, however, that they are quite independent of the traditional objection to induction. They do not just make the point that enumerative induction is logically invalid. On the contrary, they show that enumerative induction is not just invalid, but downright unreliable.

The moral, for us reliabilists who want to resist scepticism about induction, is that we had better not take our stand on simple enumerative induction as schematized in (1). Rather, we need somehow to show that our actual inductive practice has a more sophisticated structure, perhaps involving restrictions on the As and Bs which are candidates for projection, and perhaps limited in the degrees of belief which we extend to its conclusions.

I shall say a bit more about such an alternative model of inductive inference shortly. But first let me observe that such a model will open the way to the anti-sceptical strategy outlined above once more. Imagine that we can show that our actual inductive practice is more sophisticated than simple enumerative induction, and that it therefore cannot be discredited as unreliable by either Goodman's new problem or past performance. And imagine, furthermore, that when we investigate the reliability of our inductive methods, using existing methods of investigation, including those inductive methods themselves, we find ourselves able to conclude that

it is reliable. Then this defeats scepticism about our inductive practice. As before, neither the fact that this practice may be invalid, nor the fact that its reliability might only be discoverable in a rule-circular way, will be a barrier to our concluding that it yields knowledge.

In this chapter I have been concentrating on induction. But the anti-sceptical strategy I have used can be generalized to apply to our belief-forming methods in general, including such non-inferential methods as perception and memory. If the only objection to them is that they are not certain, in the sense that it is possible that they should yield false beliefs, then this is no reason to believe that they are not reliable. And if, moreover, investigation shows that those methods are reliable, then, any rule-circularity notwithstanding, we will be in a position to conclude that they yield knowledge.¹⁵

5.15 Non-Enumerative Induction

I now want to outline a more realistic model of induction, with the intention of showing that induction is not in fact discredited as unreliable by either Goodman's argument or by past performance.

To start with Goodman, note that any solution to Goodman's problem is likely to lead us to view induction in terms of elimination rather than enumeration. For any solution will involve some limitation on the As which are possible candidates for being associated with any given B. But if we have such a limited range of possible As, then it should be possible for us to find experimental data which will identify the actual antecedent of B by eliminating the other candidates, rather than by providing repeated instances of the relevant generalization.

Within the context of deterministic assumptions, J.S. Mill's methods of induction show what kinds of data are required to reach such eliminative conclusions, given various kinds of assumptions about limited ranges of possible antecedents. Much contemporary science, it is true, does not assume determinism. But there are probabilistic analogues of Mill's methods, which use techniques of analysis of variance and multiple regression to discriminate, among the factors which might in principle be correlated with some effect B, those which are genuinely rather than spuriously correlated with it.¹⁶

Mill's and related methods are nowadays little discussed by philosophers. This is unfortunate, for Mill's methods are clearly far more in accord with actual scientific practice than the standard philosophical model of induction by enumeration of instances. Science does not need a large number of repeated observations to establish that copper melts at 1084°C, or that chickenpox is caused by a herpes virus, or that water is H₂O. Rather, since there are only so many possible melting points, or infectious agents, or combinations of elements, a few relatively simple observations will suffice in each such case to discriminate the actual law from the initially possible alternatives.

No doubt part of the reason philosophers have been uninterested in Mill's methods is that Mill himself does not offer any satisfactory response to sceptical doubts about induction. But the argument of this chapter shows how we can defend Mill's

methods against such sceptical doubts: namely, by showing that those methods are a reliable source of true beliefs, and so of knowledge. Moreover, the argument of this chapter shows that there is no reason why such a demonstration should not be rule-circular: what we want is some route to the conclusion that those methods are reliable, but not necessarily a route that avoids those methods themselves.

In the rest of this section I will sketch out one possible route to the conclusion that Millian methods of eliminative induction are reliable.¹⁷ But before I do so, let me observe that a demonstration of the reliability of Mill's methods will also constitute an implicit answer to Goodman's new problem of induction. For, in order for Mill's methods to be reliable, only a certain limited range of As can possibly be associated with any given B. So any demonstration of the reliability of Mill's methods will need to show that these As are indeed the only candidates for projecting along with B. That is, if Mill's methods work, then there must be general reasons why only some sorts of generalizations -- like "all emeralds are green" -- are on the cards, and that others -- "all emeralds are grue" -- are not. A demonstration of the reliability of Mill's methods will thus show why green is projectible with respect to emeralds and grue is not.

Let us look a bit more closely at Mill's methods. It is important, in thinking about the reliability of these methods, not to view the kind of background assumption which tells us, say, that B has a deterministic antecedent, and that A_1, \dots, A_n are the only possibilities, as a premise to a Millian inductive inference, to which we conjoin the further, observational premises that A_1, \dots, A_{n-1} have been found without B, to conclude that A_n is the actual antecedent.¹⁸ For this would just make Millian induction a special case of deduction, and moreover it would leave the scientists who engaged in Millian inferences with an undischarged premise, namely, the premise that one of the A_1, \dots, A_n is the deterministic antecedent of B.

Rather, we should think of the observation of A_1, \dots, A_{n-1} without B as the sole premise to a Millian inductive inference. We don't need to suppose that the scientists themselves know why this inference works (though they probably will); all that matters, in the first instance, is that they are disposed, on observing A_1, \dots, A_{n-1} without B, to conclude that A_n is the actual antecedent.

The further thesis that B has a deterministic antecedent, and that A_1, \dots, A_n are the only possibilities, need only come in at the philosophical meta-level, when we address the question of the reliability of the Millian inductive inference. It is we philosophers, who want to ascertain whether the Millian inference is reliable, who need to know that one of the A_1, \dots, A_n must be the deterministic antecedent of B, not the scientists who actually make this inference.

So the scientists who make Millian inferences are not necessarily guilty of helping themselves to undischarged premises. The only premise they need is that A_1, \dots, A_{n-1} have been found without B, and that they can get from observation. On the other hand, the complaint of undischarged premises can reasonably be levelled at a philosopher, like myself, who explicitly argues that such Millian inductive inferences are reliable, on the grounds that one of the A_1, \dots, A_n must be the deterministic antecedent of B. For I at least then owe some account of my basis for this latter claim.

When we were still thinking of induction as simple enumeration, this was the point where we turned induction on itself, and used an inductive inference to arrive at the conclusion that induction is reliable. If we could do this in the present context, then we would once more have an answer to sceptical questions about induction. However, it is not so obvious that we can make the same move within the context of an eliminative approach to induction.

Consider the inference by which medical scientists establish that chickenpox is due to a herpes virus. We can construe the scientists as inferring this from a set of Mill-style observations about the presence and absence of various viruses in people with and without chickenpox. And we can explain the reliability of this inference on the grounds that viral agents are the only possible causes of infectious diseases that do not respond to antibiotics.

But now what about this latter claim, that viruses are the only possible causes of antibiotic-resistant infectious diseases? We need this claim to explain the reliability of the medical scientist's Millian inference. But where does it come from? The trouble is that we can't get it from the Millian inference in question itself, since this form of inference can only tell us which viruses are the causes of which antibiotic-resistant infections, not which agents are responsible for antibiotic-resistant infections in general. That is, the Millian inference in question isn't self-supporting in the way enumerative induction is.

Perhaps we can still establish the general claim that viruses always cause antibiotic-resistant infectious diseases by another type of eliminative induction, different from the first. What we would need would be a type which uses suitable observations to eliminate all the other possible agents apart from viruses as the cause of such diseases, not one which eliminates all other viruses apart from the herpes virus as the cause of chickenpox.

Now, it is plausible enough that there is such a reliable mode of eliminative induction. But then the same problem arises again. For the reliability of this new mode of eliminative induction will now rest on some further fact, such as that only a certain range of invasive physical agents that disrupt the biochemistry of the sufferers are possible candidate causes of diseases. And so the philosopher who wants to assert the reliability of this new form of eliminative induction has a new undischarged premise to cope with.

A kind of regress threatens. Eliminative inductions seem to fall naturally into a hierarchy, with the reliability of each being explicable only with the help of assumptions which derive from a form of inference higher in the hierarchy. We can explain the reliability of the procedures which show that a herpes virus causes chickenpox by invoking the assumption that antibiotic-resistant infectious diseases are always due to viruses. But this assumption in turn needs to be established by a different procedure, whose reliability depends on some such assumption as that all diseases are due to invasive physical agents that disrupt the biochemistry of the sufferers. And perhaps this assumption too can in turn be established by an eliminative induction, which uses relevant observations to discriminate between this assumption and other physically possible models of disease. But this then leaves us philosophers with the task of explaining how we came by the assumption that the physical possibilities are the only possibilities.

ÊÊ We can expect to find this kind of hierarchy repeated in different areas of science.Ê Perhaps the fact that such sequences will characteristically go from subject-specific assumptions to general assumptions about physical possibility offers a way out of our problem.Ê If the regress is to stop, at some point we will need to find a form of eliminative induction which can establish its own reliability, as enumerative induction did earlier.Ê I would like tentatively to suggest that such a form of eliminative induction might be found at the level of basic physical science.Ê At this level the task is to discover the limits of physical possibility itself.Ê Basic physics aims to decide between different theories of force, matter and spacetime, between different theories of what is physically possible.Ê Now, it is plausible that at this level the inductive strategy used by physicists is to ignore any such theories which lack a certain kind of physical simplicity.Ê If this is right, then this inductive strategy, when applied to the question of the general constitution of the universe, will inevitably lead to the conclusion that the universe is composed of constituents which display the relevant kind of physical simplicity.Ê And then, once we have reached this conclusion, we can use it to explain why this inductive strategy is reliable.Ê For if the constituents of the world are indeed characterized by the relevant kind of physical simplicity, then a methodology which uses observations to decide between alternatives with this kind of simplicity will for that reason be a reliable route to the truth.

ÊÊ It should be emphasized that this story does not depend on any a priori notion of simplicity.Ê To this extent the term "simplicity" is perhaps a misnomer.Ê The account simply depends on the existence of certain general features which characterize the true answers to questions of fundamental physical theory.Ê Far from being knowable a priori, these features may well be counterintuitive to the scientifically untrained.Ê Thus circular motion is not especially "simple", in the relevant sense, compared to the kinds of motion that results from inverse square force laws.Ê Discontinuity is not "simple", notwithstanding the fact that everyday experience shows us sharp boundaries between physical objects and their surroundings.Ê Explanations in terms of observable causes is not "simple", compared with explanation by microscopic hidden mechanisms.

Ê

5.16Ê The Historical Contingency of Knowledge

It follows from this last observation that the methods of inference which enable us to find out about the physical world are not native to human thinking.Ê At some point in human history people acquired the ability, which they did not have before, to focus specifically on certain kinds of explanations of physical phenomena, and to ignore others.

ÊÊ This might seem to raise a problem for my overall response to scepticism.Ê My suggestion has been that we should check our existing methods of thought to confirm that they are reliable sources of true belief and hence of knowledge.Ê And, in answer to the objection that these checks themselves employ those same existing methods of thought, I have observed that we have as yet been given no reason to distrust those methods.

ÊÊ However, some readers may feel that this strategy is markedly less plausible when applied to methods of thought which are historically happenstantial, rather than to

those which are an innate part of the human cognitive endowment. Demonstrations of reliability that depend on such innate methods of thought at least have the virtue of persuading all actual human beings; it is only purely though-experimental individuals, like the counter-inductivists, who will be unpersuaded by such demonstrations of the existence of human knowledge. But demonstrations of reliability by appeal to historically contingent methods of thought like the preference for "physically simple" hypotheses will fail to persuade many actual human beings as well -- namely, all those individuals in the historical, geographical, or cultural distance who have not adopted the contingent methods of thought in question.

Still, does this matter? That other people foreswear some method of thought is not in itself a conclusive reason for distrusting that method of thought, nor, therefore, for distrusting vindications of human knowledge that depend on it. Maybe our reasons for thinking that we have knowledge, and that the sceptic is therefore wrong, will fail to persuade various actual human beings, as well as the imaginary counter-inductivist. But this scarcely shows they are not good reasons. After all, there is plenty of evidence to show that many methods of thought which are native to human thinking are untrustworthy, and that we therefore do better to replace these native habits by historically contingent alternatives.

Perhaps a kind of generalization from variability lies behind the widespread feeling that historically contingent methods of thought are epistemologically suspect. I have in mind the following line of argument: different people in different times and places have adopted many different procedures for deciding questions about, say, the causes of observable phenomena; since these different procedures standardly deliver mutually inconsistent answers, only one of which can be true, most of these procedures must be unreliable; so it is highly probable that any given such procedure must be unreliable; in particular, therefore, it is highly probable that our currently preferred procedure is unreliable.

However, once it is spelt out, the weakness in this line of argument is apparent. For it presupposes that such a set of historically varying procedures forms a homogeneous category, whose overall inadequacies therefore detract equally from all members of the category. However, there may well be relevant differences between members of the category, differences which block the inference from "most of the procedures in the group are unreliable" to "this particular procedure is probably unreliable". For instance, it may be that some currently preferred procedure has been adopted as a result of critical reflection, rather than mere deference to tradition; perhaps it makes use of controlled experiment, rather than mere hearsay; maybe it involves mathematical precision, rather than mere guesswork. These features do not of course provide a conclusive demonstration that the procedure in question is a reliable source of truth. But they are surely sufficient to invalidate the argument that, since various procedures which lack these features are unreliable, our preferred procedure must be unreliable too.¹⁹

Ê

5.17 The Pessimistic Meta-Induction

The argument from relativism considered in the last section needs to be distinguished from the well-known "pessimistic meta-induction from past falsity". The "pessimistic meta-induction" also calls in question the reliability of the inductive

strategies of modern science. But instead of maintaining that there is nothing to choose between these strategies and the incompatible alternatives which have been practised by historically and culturally distant humans, the pessimistic meta-induction focuses directly on the output of modern scientific method, and argues straight off that this scientific method must be unreliable, since it characteristically issues in false beliefs.

After all, the pessimists can point out, we now take it that Newtonian physics, the phlogiston theory of combustion, the theory that atoms are indivisible, and so on, are all false. So doesn't it immediately follow that method by which these theories were reached cannot be a reliable route to the truth?²⁰

This pessimistic meta-induction is unquestionably an important argument, which indicates that caution is necessary in scientific theorizing. However, it raises many questions of detail, and it would extend our discussion unduly to treat it fully here. I shall restrict myself to a few brief comments, which I hope will suffice to show that, even if the pessimistic meta-induction advises caution, it does not mean that we should withhold belief entirely from all scientific claims.

The basic flaw in the pessimistic meta-induction mirrors that in the argument from variability discussed in the last section: it lumps into one homogenous category items that deserve separate treatment. The argument from variability lumps together all human thought processes. Similarly the pessimistic meta-induction lumps together all scientific theories.

In so doing, the pessimistic meta-induction ignores important differences between scientific theories, differences which matter to the question of whether the historical record casts doubt on their truth. For a start, the tendency to theoretical falsity is much more common in some areas of science than others. Thus it is relatively normal for theories to be overturned in cosmology, say, or fundamental particle physics, or the study of primate evolution. By contrast, theories of the molecular composition of different chemical compounds (such as that water is made of hydrogen and oxygen), or the causes of infectious diseases (chickenpox is due to a herpes virus), or the nature of everyday physical phenomena (heat is molecular motion), are characteristically retained once they are accepted. So the testimony of past form counts against some kind of theories more than others. Past scientific failures indicate caution about the thesis that quarks and leptons are the ultimate building blocks of matter. But they give us no reason to doubt that water is made of hydrogen and oxygen.

Nor need we regard this differential success-rate of different kinds of theories as an inexplicable historical datum. It may simply be a result of the necessary evidence being more easily available in some areas than others. Paleoanthropologists want to know how many hominid species were present on earth three million years ago. But their evidence consists of a few pieces of teeth and bone. So it is scarcely surprising that discoveries of new fossil sites will often lead them to change their views. The same point applies on a larger scale in cosmology and particle physics. Scientists in these areas want to answer very general questions about the very small and the very distant. But their evidence derives from the limited range of technological instruments they have devised to probe these realms. So, once more, it is scarcely surprising that their theories should remain at the level of tentative hypotheses. By contrast, in those areas where adequate evidence is available, such as chemistry and

medicine, there is no corresponding barrier to science moving beyond tentative hypotheses to firm conclusions.

Ê I do not necessarily want to suggest that there is a level of evidence which will ensure sure-fire inductive inferences, inferences that are 100% reliable in the actual world.Ê Maybe even the best achievable evidence will on occasion lead scientists astray.Ê However, by distinguishing between well supported and badly supported theories, we can at least avoid tarring all scientific conclusions with the failures of poorly supported speculations.

Ê If even the best evidence is less than sure-fire, then we ought never to accord strictly full belief to scientific conclusions.Ê Rather, we ought to tailor our degree of belief to the reliability of similarly evidenced conclusions, in the way indicated in section 5.6 above.Ê But as long as the success-rate of well evidenced inferences is high, this degree of belief can still be close to one.Ê Speculations based on meagre evidence may often turn out to be false.Ê But this is no reason to think all inductive conclusions will suffer the same fate.

1. This terminology derives from Dummett, who introduced the term "anti-realist" for theories of meaning according to which truth is not evidence-transcendent.Ê (Cf Dummett, 1978, p 146 and passim.)

2. Gettier cases show that we need an extra requirement.Ê Not only must you acquire a true belief from a process that generally produces true beliefs, but the truth of your belief must not be an accident relative to its coming from that process.Ê This extra requirement is a natural upshot of the hypothesis that knowledge is a state which is a means of avoiding error;Ê for if the truth of your belief is a result of a Gettier-style accident, then your avoidance of error won't be a result of your embodying a truth-conducive process, in the sense that it won't be because the process generally produces truths that you have avoided error.Ê Rather, it will be because of some lucky fluke.Ê Cf. Williams, op cit, pp 43-4.Ê Further analysis is needed, of course, to make the relevant notion of accident precise.

3. What of the many philosophical views, especially contemporary views, which favour subjective warrants, but not certainty?Ê I have two excuses for running the two requirements together.Ê First, I can think of no good rationale, as opposed to intuitions, in favour of the demand for subjective warrants, except as a corollary of the desire for certainty.Ê Second, most of my arguments will apply as much to former demand taken separately as to the latter.

4. For more on the extent to which we can alter our non-conscious belief-forming processes, see my *Reality and Representation* (1987, ch 7.4, 7.6).

5. Perhaps passive knowers should have at least this much in common with concerned enquirers:Ê their belief-forming processes should not just happen to be reliable, but should be present because they are reliable.Ê This suggestion will rule out purely fortuitous reliability, but will allow in reliability due to evolution, learning, and education, alongside the case where concerned enquirers adopt processes because they consciously recognize the reliability of those processes.Ê This suggestion also indicates a possible answer to the delicate question of how to individuate belief-forming processes for purposes of assessing their reliability:Ê namely, as a first

approximation, we should individuate them by the same characteristics as are needed to explain their adoption. Cf Papineau (1987, pp 136-8).

6. See Shope *The Analysis of Knowing: A Decade of Research* (1983).

7. I am here assuming beliefs about non-chance matters. Where genuine non-unitary chances are involved, the best degree of belief about any outcome will, of course, be different from one. But even here it will always be better to get true beliefs of full degree about chances, belief-forming costs apart, for such beliefs will then ensure that in general your decisions maximize your objective chances of success.

8. Some readers may be unhappy with these assumptions about the knowledge available to concerned enquirers. There are two possible worries here, one about a possible circularity on my part, the other about a possible circularity on the part of concerned enquirers. If you are worried about a circularity on my part, let me observe that my present concern is not to define knowledge, so much as to identify the role the concept plays in our thinking: I think it is helpful in this task to consider the predicament of a concerned enquirer who already has some knowledge; it would be a further task, which I shall not attempt here, and which would indeed preclude mention of the knowledge of concerned enquirers, to specify necessary and sufficient conditions, in non-epistemological terms, for someone to be a knower. If, on the other hand, you are worried about some kind of circularity on the concerned enquirer's part, on the grounds, perhaps, that any seriously concerned enquirer ought to start by assuming nothing, then your worry should be assuaged by the points made in sections 5.11-13 below.

9. Craig, *op cit*, pp. 218-21, argues similarly from the limited informational situation of an enquirer to a counterfactual requirement on knowledge. But he is concerned with the kind of counterfactual reliability we want of our informants, given our limited information about the particular situation at hand, whereas I am interested in the counterfactual reliability we want of ourselves, given our limited general information about the situations we are going to be in. The precise degree of counterfactual reliability required for knowledge is a complex issue, which I shall not pursue any further here, except to observe that this is an area where one good theoretical rationale seems to me to be worth a thousand delicate intuitions.

10. This claim perhaps deserves further discussion. Some reliabilists would hold that, while reliability suffices for knowledge, some kind of extra subjective warrant is needed for rationality. In my view, however, the arguments about knowledge go through for rationality. For a defence of the analogous point about "justification", see Goldman (1979).

11. Cf Van Cleve (1984) p. 559, 562. Much of my following defence of induction is influenced by Van Cleve's important article. However, in the passages referred to, Van Cleve seems to deny that reliabilists need to defend the reliability of induction, on the grounds that reliability is an "external" requirement. This seems to me an unfortunate slip, given that evaluations of reliabilist defences of induction are highly sensitive to prior judgements of exactly what the reliabilist needs to do.

12. Cf Braithwaite (1953, pp 276-7), Van Cleve, *op cit*, p 558.

13. There is a problem about the notion of premise circularity: if "contained among the premises" just means logical implication, then all deductive arguments will be premise circular. Some philosophers, most notably Descartes and Mill, take this to show that deduction is uninformative. The majority prefer to understand "containment" more strictly. This debate is irrelevant to our current concerns, however, since everybody will agree that the non-deductive argument (2) is not premise circular.

14. If it does, then this will answer the second of the questions raised about the premises of (2) at the beginning of this section, namely, the question as to how we know that the conclusions ("All As are Bs") of past inductions are true. The answer is, by induction. For, if the distinction between rule- and premise-circularity legitimates the inductive move from (2)'s premises to (2)'s conclusion, then it will also serve to legitimate antecedent inductive moves from instances of As being Bs to "All as are Bs". Cf Van Cleve, *op cit*, pp 560-1.

15. One of the aims of my *Reality and Representation* (1987) was to show how we might investigate all of our standard methods of belief-formation, as we might investigate any other natural phenomena, and discover that those methods are by and large reliable sources of truth.

16. These probabilistic analogues of Mill's methods involve two inferential steps: first from sample data to objective correlations, then from these correlations to causes. While this second step can be 100% reliable, the first cannot avoid the uncertainties of statistical inference. (Cf Papineau, 1993.) I have hopes for a reliabilist-style account of statistical inference, but there is no question of going into this here.

17. This will expand the discussion in Papineau (1987, pp 196-8).

18. This simple method of difference is just one of Mill's methods. For a full catalogue see the Appendix to Mackie (1974).

19. In *Reality and Representation* my response to the argument from variability was that there is no variability among self-certifying batteries of belief-forming procedures, that is, batteries of procedures which generate conclusions about their own reliability (1987, ch 10). I still think this argument can be made to work, but it now seems rather a big hammer to crack a small nut. From my present perspective, self-certification is merely one example of the kind of relevant difference which can block the generalization from variability.

20. Doesn't this argument undermine itself, by casting doubt on the present-day scientific theories whose truth it assumes when judging past theories to be false? But the argument can be cast as a *reductio*: if current theories are true, then past theories are false; so, by the pessimistic meta-induction, current theories are false; so current theories are false. (Cf Jardine, 198x, p x.)