

OUP Copyright

CHAPTER 16

Knowledge, research, and evidence-based medicine

Chapter contents

- Session 1 **Evidence-based medicine, Hume, and the problem of induction** 435
- Session 2 **Philosophy of science responses to the problem of induction** 440
- Session 3 **Epistemological responses to the problem of induction** 449
- Session 4 **Evidence-based medicine and clinical trials** 457
- Reading guide** 462
- References** 463

OUP Copyright

Introduction

Like other areas of medicine, treatment and management within psychiatry has increasingly been influenced by the growth of evidence-based medicine or evidence-based practice. A recent review article in the *British Journal of Psychiatry* starts with this observation.

Clinical effectiveness, evidence-based medicine (EBM) and related terms were the politically correct medical slogans of the 1990s. For many they are 'buzz-words' conveying a modern progressive approach and in some circles it is unwise to express scepticism. Evidence-based medicine is being embraced by all specialities and there has been a strong signal that psychiatry is joining the movement by the introduction in 1998 of a psychiatric journal dedicated to evidence-based practice.

Williams and Garner (2002, p. 8)

The authors go on to express some guarded scepticism about the scope of the application of EBM to psychiatric care. They argue that 'too great an emphasis on evidence-based medicine oversimplifies the complex and interpersonal nature of clinical care' (ibid., p. 8). What is striking is the degree of caution expressed in making this modest criticism. Why is EBM so influential and what limitations might it possess?

At heart, EBM concerns how best to learn from experience. So this chapter will start by looking back to a very influential if abstract way of raising that general question: the eighteenth century Scottish philosopher David Hume's notorious 'problem of induction'. It will then examine two different kinds of response to the challenge raised there. This will help shed light on the nature of knowledge itself and the social character both of the production and transmission of knowledge. In turn this will suggest that while evidence plays a bottom-up role in supporting scientific theories, there is also a 'top-down' influence on evidence by theory. This interplay of factors calls for judgement by scientific practitioners.

One response to the problem of induction is contained within the philosophy of science. It involves methodological routes to scientific conclusions. (One cannot say routes to scientific *knowledge*, however, because Popper's influential falsificationist account denies that positive reasons for belief can be given in the face of Hume's problem.) Examining some of the conflicting responses made within the philosophy of science will highlight the role of the broader scientific context to the evidence that can be offered for scientific medical claims.

The second kind of response is contained within a different branch of philosophy: epistemology. This branch concerns discussion of the nature (logos) of knowledge (episteme). Responses within epistemology to the problem of induction have tended to argue against the problem. That is, they have attempted to show that there is no real problem once the nature of knowledge is properly understood. By examining some responses made within this branch of philosophy we aim to show how knowledge can be a product of social processes and can be socially transmitted.

The connection between practical debates surrounding EBM and the more abstract problem of induction is that EBM concerns the issue of how best to learn from past experience and to apply past findings to future practice. The strategies generally favoured by EBM are based on the use of evidence derived from randomized clinical trials (RCTs). Now it may seem obvious that randomized trials comprise the best method of arriving at clinical findings, and indeed there are good arguments for this method. However, as the chapter will argue, RCTs should not be thought of as ways of simply harvesting available data in accordance with a priori reasoning about proper scientific method. Instead they have their place within a broader context of scientific theory that influences both the methods used for gathering data and the interpretation of that data. Indeed, as Chapter 12 argued, it is misleading to talk of data in this way as though there were a way of gathering facts *before* interpreting them. A proper understanding of the role of evidence in scientific psychiatry suggests the need to balance both bottom-up and top-down interdependence of evidence and theory.

Hume's problem of induction

The first session will explore how Hume's sceptical argument appears to undermine knowledge based on induction. Inductive reasoning—reasoning from particular facts to generalizations—plays a *prima facie* important role in science (although there are philosophers, as we will see, who deny this). Deductive reasoning—reasoning from general principles to other general principles or to particular facts—is important as well (as in the Hypothetico-Deductive method deployed in diagnosis), but it is induction that appears to be closely tied to the very idea of an *empirical* science. Even on the assumption that induction plays a role in empirical science, it cannot be the whole story, because scientific reasoning also involves the postulation of new and sometimes unobservable entities. But it does play a role in grounding or warranting (as it is sometimes called in philosophy) scientific laws.

Hume's problem of induction is a direct sceptical attack on the justification of this sort of reasoning. He argues in effect that induction can only be justified by induction itself and that that is question-begging.

Induction and the philosophy of science

Session two will then examine one strand of thinking in the philosophy of science, particularly associated with the late Sir Karl Popper, which starts with an assumption that inductive inferences cannot support knowledge claims and thus that science should be construed as a systematic structure of conjecture and *refutation*. The positive idea of evidence supporting theories is replaced by an emphasis on the use of evidence to reject false theories. This turns on the simple idea that while no finite amount of confirming instances can entail the truth of a universal generalization, a single counter-instance can entail that it is false. In fact, however, this simple thought cannot yield the precise methodology

that it might, at first, suggest. For reasons that build on the discussion of the theory dependence of observation of Chapter 12, even refutation is more difficult than it might seem.

Responding to these difficulties, another famous figure in modern philosophy of science, Imre Lakatos, developed falsificationism into a sophisticated model of a historically extended process of developing theories within competing research programmes. But as Lakatos admits, this does not yield a *prescriptive* methodology.

Partly as a response, historians and sociologists of science, including most famously Thomas Kuhn argued in the 1970s that there were no applicable trans-historical patterns of scientific rationality and that understanding of scientific practice should instead concern itself with the local perspective on what was deemed rational by the practitioners of science themselves. Most interestingly, perhaps, they have the courage of their conviction and also say that this is the perspective which should be applied to the sociology of science.

But, whatever their efficacy in resolving the problem of induction, what these discussions do help to show is that there is more to theory testing than simply gathering data and forming inductions on the basis of it. Evidence is interpreted in the light of background theories that are supported by, but also support, evidential claims.

Induction and epistemology

While, within the philosophy of science, the response to the problem of induction has focused on methodological strategies, philosophical epistemology has concentrated instead on rethinking the assumptions about knowledge that make induction appear problematic. One such assumption is that knowledge is self-intimating: that when one knows something, one knows that one knows it. Another is that knowledge requires that the knowing subject can supply a justification for their knowledge claims. Session 3 will examine the success that more recent work in epistemology has had in resolving the problem of induction.

Although the focus will be responses to the problem of induction Session 3 will highlight the dependence of individual inquirers or scientists on other people in acquiring knowledge. Knowledge, in the phrase of the philosopher the late Gareth Evans rubs off on other people 'like an infectious disease'.

In fact, what should become clearer during the course of this chapter is that there is something to be said for both approaches: epistemological approaches show how Hume's scepticism can be defused by adopting a more realistic picture of knowledge while the philosophy and sociology of science show the kind of practical arguments that count in actually selecting theories.

Evidence-based medicine

Session 4 will examine the consequences of these discussions for the role of evidence in medicine. Starting with an examination of some of the assumptions underlying clinical trials, the final session will look at recent emphasis on EBM and the nature of

the evidence in question and the strengths and weaknesses of this approach in light of advances in our understanding of the nature of science provoked by Hume's original problem. It will also summarize how issues raised throughout this part have an impact on a better understanding of EBM. Underlying the simple idea that medical practice should be based on the best available evidence is the real complexity, and the element of uncodifiable scientific judgement, in the generation of such evidence.

Plan of the chapter

- ◆ *Session 1* puts discussion of EBM into the context of the philosophical origins of the problem of induction in Hume's work.
- ◆ *Session 2* charts responses to Hume's problem drawn from the philosophy and sociology of science. Such responses suggest a number of rival models of scientific research and rationality, which emphasize the top-down influence of theory on evidence as well as more expected bottom-up dependence of theory on evidence.
- ◆ *Session 3* sets out some responses to the problem of induction drawn from philosophical epistemology. These help suggest a connection between a more realistic contemporary understanding of the nature of knowledge and its social transmission.
- ◆ *Session 4* looks to Mill's methods as a way of examining the origins of research evidence in controlled trials and thus the ineliminable role of good judgement in assessing evidence.

But before starting that detailed analysis we will turn to a positive expression of the virtues of EBM for psychiatry.

Session 1 Evidence-based medicine, Hume, and the problem of induction

EXERCISE 1

(30 minutes)

Read the whole of the short paper

Geddes, J.R. and Harrison, P.J. (1997). Closing the gap between research and practice. *British Journal of Psychiatry*, 171: 220–225

Link with Reading 16.1

- ◆ What is the purpose and rationale for EBM in mental health care?

The aim of evidence-based medicine

Geddes and Harrison (1997) present a very clear account of the motivation for adopting EBM. There is, they argue, a 'knowledge gap' between the accurate information generally employed by clinicians and the decisions that they make partly on the basis of it. In other words, clinical decisions are underdetermined by the readily available (pre-EBM) evidence. This is not to say that the

decisions are therefore made arbitrarily. The gap is instead filled by other factors: 'the conceptual aetiological school to which we subscribe' and 'the combination of experience and habits which we accumulate' (p. 220). However, these factors vary from clinician to clinician, which thus undermines the inter-rater reliability of diagnosis and treatment. Nor are they empirically tested guides. Hence, instead, the need for EBM to guide diagnosis and treatment.

Evidence-based medicine and randomized clinical trials

'EBM is the "conscientious, explicit and judicious use of the current best evidence in making decisions about the care of individual patients"' (p. 220) The kind of evidence in question is that provided through clinical trials. But there are issues about how to assess or rank conflicting evidence. The paper asserts that RCTs, or better still systematic reviews of RCTs, are the most reliable study design for the evaluation of treatments' (p. 221). But because such trials are not always available, Geddes and Harrison (1997), following widely accepted principles of EBM (see, e.g. Sackett *et al.*, 2000), suggest that there is a hierarchy of kinds of evidence that can also be pressed into service. It is as follows:

- 1a. Evidence from a meta-analysis of RCTs
- 1b. Evidence from at least one RCT
- 2a. Evidence from at least one controlled study without randomization
- 2b. Evidence from at least one other quasi-experimental study
3. Evidence from non-experimental descriptive studies, such as comparative studies, correlation studies and case-control studies
4. Evidence from expert committee reports, or opinions and/or clinical experience of respected authorities.

Complexities in the use of evidence-based medicine

It is worth noting that despite the presentation of this hierarchy of forms of evidence, the whole tenor of the paper is that EBM is not a mechanical procedure that removes the need for intellectual effort on the part of the clinician. To the contrary, he or she must be able to search out sources of evidence, assess their quality or reliability, assess their relevance to the clinical case at hand, and interpret what conclusions should be drawn from the evidence. But despite this emphasis, there still remains a danger that the concentration within EBM on RCTs disguises some of the complexities of scientific method discussed so far. An awareness of these helps to immunize against the assumption that things are as clear-cut as less nuanced accounts of EBM may sometimes seem to suggest.

The origins of the problem of induction: David Hume

We will start with Hume's (1975) sceptical attack on the foundations of induction, although Hume himself does not use that

term. It is contained in a section that precedes his discussion of causation, which was the subject of Chapter 15, but it is related. As we will see, Hume suggests that knowledge can be divided into either demonstrative relations of ideas or matters of fact. The latter are established on the basis of causal relations that are in turn founded on experience (by contrast with deductive demonstration). He is then concerned to establish their credentials for establishing matters of fact (causal relations) through experience.

EXERCISE 2

(15 minutes)

Read the extract from

Hume, D. ([1748]1975). *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. Oxford: Oxford University Press, section iv (extract pp. 37–38)

Link with reading 16.2

- ◆ What is the kernel of Hume's sceptical argument?

Hume's fork

Hume starts the section from which this extract is taken with the bold claim that all the objects of human reason can be divided between relations of ideas and matters of fact. This distinction, generally known as 'Hume's Fork' sets up the contrast that will be important between the status of knowledge claims that can be arrived at through deductive demonstration and those for which a merely inductive warrant can be provided.

It is worth noting briefly two features that characterize the distinction.

1. Truths that comprise relations of ideas do not depend on, or presuppose, any existence claims. Thus the claim that the square of the hypotenuse is equal to the sum of the squares of the other two sides is a truth that is independent of whether there are any right-angled triangles in the universe.
2. The negation of truths that comprise relations of ideas produces claims that could not have been true and cannot be 'distinctly conceived by the mind'. By contrast, the negation of matters of fact could have been true and thus can be so conceived. This amounts to saying that relations of ideas express necessary rather than just contingent truths (see also Chapter 5).

Three kinds of truth

A distinction of this sort has been influential throughout empiricist philosophy. Consider the three binary distinctions:

1. epistemological: a priori versus a posteriori
2. metaphysical: necessary versus contingent
3. semantic: analytic versus synthetic.

One appealing assumption has been that these different ways of sorting truths all sort them into the same sets. Thus all truths that

can be known a priori (that is, without experience) are necessarily true and their truth is fixed by the concepts used to frame them (they are analytically true in virtue of their meaning). Equally, all truths that require experience to be known, and are thus a posteriori, are also contingent and synthetic (i.e. their truth requires both a contribution from their meaning and from the world).

There are also *prima facie* plausible arguments for the alignment of these distinctions. For example, if a truth is a priori then one does not need to know which possible world one inhabits in order to know its truth. (Experience teaches us which of the many possible worlds we actually live in.) This suggests in turn that it must be a necessary truth (one that holds in *all* possible worlds). Furthermore it seems plausible that it must be analytic, because its truth clearly does not require a worldly contribution and thus must be fixed entirely by its meaning.

Despite these arguments, the neat alignment of truths has also come under attack. The Prussian philosopher Immanuel Kant (1724–1804) argued that there were synthetic truths which could be known a priori including mathematics. In this century, the American philosopher W.V.O. Quine (1908–2000) attacked the assumption that the distinctions are well founded clear-cut distinctions of kind and the logician Saul Kripke (1940–) argued that some a posteriori truths are, nevertheless, necessary (such as that water is H₂O). We will return to the consequences of Kripke's arguments in Part V. For now all that matters is noting that drawing the distinction allows Hume to focus the issue of the foundations of matters of fact. (If the distinctions do not align or are not even firm this will not solve the general problem Hume raises for justifying empirical knowledge. It merely changes the form it takes because there ceases to be a clear contrast between fallible a posteriori reasoning and the supposedly certain a priori and deductive reasoning.)

Knowledge by induction versus knowledge through the testimony of the senses

Having drawn the distinction between relations of ideas and matters of fact, Hume (1975) then focuses his attention on the status of knowledge of matters of fact. In fact, the focus is narrower still: on knowledge of what lies 'beyond the present testimony of our senses, or the record of our memory' (p. 26). Thus he does not here investigate the status of direct observational knowledge of particular matters of fact—knowledge made available by my opening our eyes to the world—or even our memory of such particular facts. On the one hand, this sets up an implicit contrast between observation and induction. The latter appears less reliable than the former. And indeed, as we will see later, induction in general will never have a stronger justification than observation in general because inductive inferences take observations as their premisses. On further reflection, however, the concepts used to frame observation reports are typically laden with theory (see Chapter 12). They can in individual cases thus be overturned on the basis of enough inductive counter-evidence.

Hume (1975) suggests that such reasoning is founded on the relation of cause and effect. It is this relation that underpins reasoning beyond our direct observations and binds unobserved facts to observed facts. 'Were there nothing to bind them together, the inference would be entirely precarious.' (p. 27). However, Hume goes on to question what grounds our knowledge of cause–effect relations. He argues that this cannot be a piece of a priori reasoning and is instead based on our experience. Hume argues both that cause–effect relations concern separate events and thus no amount of inspection of the cause–event can yield knowledge of what effect it will lead to and also that the negation of cause–effect relations does not produce any logical contradiction or a state of affairs that cannot be distinctly conceived. Having cleared the ground in this way, Hume goes on to discuss how experience rather than a priori demonstration can ground our knowledge of cause–effect relations.

Hume's problem of induction

Section iv, part II contains the sceptical discussion of induction. Hume begins by asking, on the assumption (for which he has just argued) that the foundation of our knowledge of matters of fact (aside from the case of direct perception) is knowledge of cause–effect relations, what underpins that relation? His answer is experience. This seems like a good answer and in everyday contexts would mark the end point of inquiry. But Hume, like a good philosopher, then asks a further question: 'what is the foundation of all conclusions from experience?' (p. 32). He suggests that he will argue that our knowledge of matters of fact is not founded on *reasoning* from past experience.

Taking as his example the connection between the sensible or observational properties of bread and its 'secret power' to provide nourishment, Hume argues that experience can play only a direct role in establishing that there has been such a connection in particular cases in the past, but questions how experience can underpin the extension of a more general connection 'to future times, and to other objects'.

These two propositions are far from being the same, I have found that such an object has always been attended with such an effect, and I foresee, that other objects, which are, in appearance similar, will be attended with similar effects. I shall allow, if you please, that the one proposition may justly be inferred from the other: I know, in fact, that it always is inferred. But if you insist that the inference is made by a chain of reasoning, I desire you to produce that reasoning. (p. 34)

Hume goes on to argue that inductive inferences cannot be demonstrative because their negations make good sense. The course of nature could (logically possibly) change. Thus a deductive defence of induction appears to be unconvincing because too strong. But on the other hand inductive defences of induction also appear hopeless because they are circular. Hume suggests that empirical reasoning:

proceed[s] upon the supposition that the future will be conformable to the past. To endeavour, therefore, the proof of this

last supposition by probable arguments, or arguments regarding existence [ie inductive arguments about matters of fact], must be evidently going in a circle, and taking that for granted, which is the very point in question. (pp. 35–36).

The problem in a nutshell

So stepping back from the details of Hume's argument we can set out the problem as follows. Suppose that the premiss is that all bread previously tested has been nourishing and the conclusion is that all future bread will be nourishing. Hume's challenge is to explain what form of inference justifies the conclusion.

The natural suggestion is that experience grounds the *rule* of inference as well as the *premiss* (in this case that all bread previously tested has been nourishing). It does this because of the more general piece of direct experiential knowledge that correlations between (in this case) sensible qualities and secret powers have held over time. This general experiential finding is then used to ground the inference from past to future in the specific bread case. But as Hume points out: using experience to ground the rule itself presupposes that very rule as an inference. Why should the fact that such correlations have held in the past support the claim that they will hold in the future unless an inductive inference is justified here as well?

Hume (1975) says:

When a man says, I have found, in all past instances, such sensible qualities conjoined with such secret powers: And when he says, Similar sensible qualities will always be conjoined with similar secret powers, he is not guilty of a tautology, nor are these propositions in any respect the same. You say that the one proposition is an inference from the other. But you must confess that the inference is not intuitive; neither is it demonstrative: Of what nature is it then? To say it is experimental, is begging the question. For all inferences from experience suppose, as their foundation, that the future will resemble the past... It is impossible, therefore, that any arguments from experience can prove this resemblance of the past to the future; since all these arguments are founded on the supposition of that resemblance. (pp. 36–37)

Inductive justification of induction

This passage suggests the following alternatives. For experience to yield a principle that will underpin reasoning from observed to unobserved cases we need an argument of the following sort.

- ◆ Premiss 1: Correlations have held in the past (observed cases)
- ◆ Conclusion: Correlations will hold in the future (unobserved cases)

Now by itself this argument is neither demonstrative nor deductive (it is not a valid argument). If experience of the stability of past correlations is to warrant the claim that they will continue to be stable, it does this in virtue of an *inductive* inference. But that was the very rule that this argument was supposed to justify rather than presuppose.

Deductive justification of induction

On the other hand, Hume suggests that such arguments presuppose that the future will resemble the past. Now putting aside the details of how it resembles it, if this were true, inductive arguments would be successful. But how is this fact supposed to help justify those arguments? The obvious answer is that if this is introduced as a second premiss, the argument becomes a valid demonstrative argument.

- ◆ Premiss 1: Correlations have held in the past (observed cases)
- ◆ Premiss 2: The future resembles the past (unobserved resemble observed cases)
- ◆ Conclusion: Correlations will hold in the future (unobserved cases)

But if it is to underwrite a true conclusion, this valid argument requires the truth of the second premiss, which is, as Hume points out, not a necessary truth or a truth that can be arrived at demonstratively. It requires a further inductive argument to justify it. So in either case, there appears to be no non-circular argument for inductive reasoning from experience.

EXERCISE 3

(15 minutes)

Take some time to think through the issues raised by Hume's argument. What are we to make of broadly Humean sceptical arguments directed at inductive reasoning? Hume appears to show that a substantial fraction of our knowledge claims about empirical matters are unjustified. But does this matter? Is it merely a game, a piece of word play?

Knowledge and justification

Hume's (1975) argument suggests the following worry. If we make a knowledge claim about unobserved events on the basis of observed events we open ourselves up to the challenge of providing grounds. Now there are first and second moves available to us. We can justify our claim about the *unobserved* by citing our experience of the *observed*. Hume can then challenge us to say how exactly experience of what has been *observed* has a bearing on what has so far *not been observed*. We can then say (as our second justificatory move) that we have more general experience (i.e. we have *observed* the following) that *observed* matters have always (in the past) been a good guide to initially *unobserved* but subsequently observed matters.

Prior to reading Hume this looks to discharge the burden of providing grounds for our original claim. However, Hume shows that this second justificatory claim only works if we can take it for granted that past experience is a good guide to future experience, or that what is observed is a good guide to what is unobserved, and that is just what he is challenging us to justify. At this point we have apparently no further arguments to make and the chain of justifications comes unstitched backwards from here. If we cannot justify induction in general, we cannot justify it in the particular case at hand. Thus we cannot say what the relevance of

our past experience is to our claim about the unobserved and thus we cannot justify that supposed knowledge claim. Thus we must admit that we do not know it despite our pre-philosophical inclinations.

The 'justified true belief' analysis of knowledge

This way of setting out the virulence of Humean scepticism reveals a hidden assumption that will be questioned later in the chapter: in order to know something it appears that we must either *know* that we know it or at least be *able* to know that we know it. Why is this at all plausible as an assumption? The answer is that at first sight it seems to be implicit in our concept of knowledge.

We distinguish between knowledge and mere true belief. If someone knows something it is not enough that they have true beliefs, something else needs to be added. (Note that truth and belief both appear to be *necessary* features of knowledge. One cannot know something that one does not believe. And one cannot know something falsely, although one can *think* one knows it: a case of merely *believing* falsely.)

Suppose, to use a well known example, I claim to know that the name of the President of the United States starts with a 'C' and that this is in fact true (as it was when the President was called 'Clinton'). Now suppose my reason for believing this is that I think his name is 'Churchill'. Although I have a true belief (the President's name does, at the time considered, start with a 'C'), it is not *knowledge* because it is held for the wrong reason. In this case the supporting belief is false and that is what undermines the first knowledge claim.

But knowledge claims can also be undone by faulty reasons even when the reasons involve true beliefs. Suppose I correctly believe that the President's name is Clinton but I believe this because, on the night of his election, a superstitious supporter of an opposing candidate who did not want to prejudice the final result declared to me that Clinton had already won. As a result, I believe that the name of the President of the United States starts with a 'C' and this is in fact true. And I believe it because I believe that his name is 'Clinton'. That is the justification, and it is a true belief. But it does not justify the first belief because I have *arrived at it* for the wrong reason: I would have believed it even if Clinton had not eventually won the election. So it is not itself justified.

Knowing that one knows

This sort of reasoning suggests that if knowledge is a belief that is true, and in addition justified, then that justification must be a belief (which is had or is available to be had by the subject of the first belief), which must itself both be true and be justified. This leads to the initially attractive assumption that if knowledge is justified true belief—a view that dates back to Plato—then the justification must also be known (it must be a justified true belief). In that case, however, its justification will also have to be known. This leads to a regress. And it is this regress that feeds Humean scepticism because of his argument that one cannot

provide an independent justification for induction. If one cannot do that—at the metalevel—then one cannot justify the claim that experience justifies ground level claims about future sunrises and the nourishing value of bread because knowledge at that level requires in addition that one knows that one knows. (The importance of this assumption for Hume will become clearer in the third session where it will be questioned.)

The point of philosophical scepticism

So the problem can be put like this. Some fairly low level and intuitive reasoning about knowledge leads us to believe that it requires justification and that justifications themselves have to pass muster. But an equally simple argument from Hume seems to show that inductive knowledge cannot pass this test. Thus it should not be counted as knowledge. Because the arguments concerned are so simple minded but concern something as central to us as the concept of knowledge, it is not convincing simply to dismiss this as a mere word game. If the arguments are wrong, it should be possible to show where they go wrong. If not, however, the only other conclusion is that paradoxically our reasoning itself leads to an unpalatable result. (Hume's own response here was precisely to play down the role of reason in general: to conclude that sceptical reasoning showed the limitations of reasoning.)

Now one response to any case of philosophical scepticism is to refuse to take its conclusion seriously in the sense of taking it to be true. Typically, a key feature of philosophical scepticism is that such conclusions are so radical that they cannot be accepted in practice in daily life—we could not put one step in front of the other if the past were not, or could not be assumed to be, for practical purposes, a guide to the future. But the point of philosophical scepticism is rather to force us to think more critically about such everyday assumptions of daily life. Philosophical scepticism is in this respect a kind of conceptual probe or microscope that helps us to open up and gain a more critical understanding of these everyday assumptions.

In Part I, we saw how the sceptical attacks of Szasz and others on the concept of mental illness led to a deeper understanding of that concept with important consequences for research and practice in mental health. Hume's 'problem of induction', similarly, has led to a deeper understanding of what is involved in relying on knowledge derived from experience (as in EBM). It is precisely, therefore, *because* the consequences of Hume's sceptical attack on induction are so radical that we should take his arguments seriously.

One way of taking Hume's arguments seriously is to question whether, if his conclusion is so radical, there must be something wrong with the premisses that lead to it. That is the approach that will be taken in the third session on epistemological approaches. But the next session will look instead at approaches to taking Hume's arguments seriously in the philosophy of science, approaches which, as we noted in the introduction to this chapter, are broadly methodological rather than epistemological in character.

Reflection on the session and self-test questions

Write down your own reflections on the materials in this session drawing out any points that are particularly significant for you. Then write brief notes about the following:

1. What, in general terms, is the aim of EBM? How does it relate to the philosophical discussion of induction?
2. What is 'Hume's fork' and how does it relate to what can be known? What types of empirical knowledge are there?
3. What is the source of the problem of induction.

Session 2 Philosophy of science responses to the problem of induction

Overview of the session

This session will discuss four responses from the philosophy of science that can be seen, in part, as responses to Hume's problem of induction. The first two are related. Both Sir Karl Popper (1902–94) and Imre Lakatos (1922–74) support forms of falsificationism. This is a position in the philosophy of science developed by Popper and then redeveloped (not entirely to the satisfaction of Popper) by Lakatos. It concentrates *not* on the positive use of evidence to support (by induction) generalizations and theories but rather on the refutation or falsification of theories by negative or disconfirming evidence. The underlying idea here is not so much to solve the problem of induction (despite what Popper says in the first case) but to bypass it. But it also aims to preserve a *rational* methodology for empirical science even without induction. The third and fourth responses differ in two respects. First, they argue that no such abstract rational methodology can be imposed by philosophical inquiry and instead argue that scientific method has to be investigated by local historical or sociological investigation. Secondly, however, they suggest that there is something much closer to an inductive practice albeit one that has to be investigated in a much more piecemeal empirical way.

What each of these four responses have in common is that they concern the attempt to characterize the general features of scientific method; they seek to chart the general relation between successive theories and evidence, to show how decisions are made as to which is the best or the better theory without concentrating on an unproblematic connection between evidence and a single theory.

As we will see in Session 3, this contrasts with responses that have been made to the problem of induction in the branch of philosophy called 'epistemology'. There, rather than attempting to characterize the relation between successive theories, philosophers have concentrated instead on the very idea or concept of knowledge and then applied their conclusions to the problem of whether inductive support can exist between evidence and theory.

In fact, as we anticipated in the introduction to this chapter, there is something to be said for both approaches. The readings in this session highlight the sort of structure presupposed by theory testing and thus undermine a naively inductivist picture of mere data gathering. To this extent, therefore, they deepen our understanding of how science actually works (we return to the implications of this for EBM at the end of this chapter). However, the arguments about scientific method in this session do nothing to disarm scepticism about induction as such. That is the goal of Session 3, which will in turn shed light on the nature of knowledge itself.

Falsificationism

The first response is by Sir Karl Popper, the philosopher of science (associated with although not a member of the Vienna Circle) who first developed falsificationism. We have already introduced Popper's work in Chapter 11. Perhaps the clearest statement of how falsificationism impacts on induction is given in: 'Conjectural knowledge: my solution to the problem of induction' in his *Objective Knowledge* (1972, chapter 1, pp. 1–31).

EXERCISE 4

(15 minutes)

Read the extract from:

Popper, K. (1972). Conjectural knowledge: my solution to the problem of induction. In *Objective Knowledge*. Oxford: Oxford University Press, pp. 1–31 (extract pp. 7–9)

Link with Reading 16.3

- ◆ Does Popper really solve the problem of induction?
- ◆ What is the cost of his 'solution'?

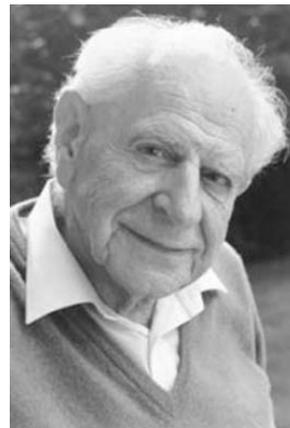


Fig. 16.1 Karl Popper

Induction and scientific rationality

Popper begins 'Conjectural knowledge' by qualifying how the problem of induction should be construed. He rejects both its key presuppositions: that the future resembles the past (unless this is construed so flexibly as to be vacuous), and that there are inductive inferences to be justified. He suggests that Hume responds to two different problems with two different and clashing answers. One concerns a logical problem of whether we are justified in reasoning from observed to unobserved instances. Hume responds that we are not. The other is the psychological problem of why we nevertheless have confidence in this form of reasoning. Hume's response to this is to invoke custom or habit. This response, Popper suggests, makes Hume an *irrationalist*. Popper flags early in the paper that his response, unlike Hume's, will not lead to a lack of rationality in scientific method.

Falsification not confirmation

The route to Popper's resolution of the problem of induction is to make a subtle change in the question that he thinks needs a positive answer. Instead of asking whether empirical reasons can be given to show that a universal theory (i.e. something that goes beyond observed facts) is true, one should ask whether such reasons can be given to show whether a theory is true or false. Theories can sometimes be shown to be false, he argues, and this is important because in general the question is asked within the context of science when there is more than one competing theory in play. He goes on to emphasize that this implies that 'we must regard all laws and theories as hypothetical or conjectural; that is, as guesses' (p. 9) and, later in the chapter, to suggest that his response remains within the domain of deduction rather than induction (p. 12).

This last point turns on a key distinction for falsification (more so on the more simple-minded versions; less so for more sophisticated versions as we will see). '[F]rom the point of view of deductive logic there is an asymmetry between verification and falsification by experience' (p. 12). The asymmetry is in the fact that, in principle, even a single observation can refute a universal claim—its falsity can be *deduced*—but no finite amount of observations can logically confirm such a claim. For this reason Popper stresses the role of severe tests of theories, or crucial experiments designed to disprove them. Nevertheless, as he points out later, this process of elimination of false theories by falsification is not a method of establishing truth because 'the number of *possibly* true theories remains infinite' (p. 15).

Given this focus on the negative side of theory testing, Popper emphasizes the intellectual virtue of theories that are, perhaps surprisingly, most easily falsified. Such theories, the boldest and most specific conjectures, are also the most informative and least vacuous. The reason for this is that they are true in the smallest number of possible worlds and thus if true, would more precisely specify the nature of the actual world. More probable theories are less informative because they rule out fewer alternative ways the world might be.

Practical theory choice

These logical conclusions leave open, however, the practical consequences for the guidance of action through beliefs arrived at by induction. Popper distinguishes between two such pragmatic issues:

- ♦ Pr1 Upon which theory should we *rely* for practical action, from a rational point of view?
- ♦ Pr2 Which theory should we *prefer* for practical action, from a rational point of view? (p. 21, emphases added)

He suggests as an answer to the first, that we should not rely on any theory 'for no theory has been shown to be true, or can be shown to be true' (p. 21). This reflects the fact that Popper does not so much solve the problem of induction—thus justifying inductive inferences—as side-step it. But this leaves no positive relation of support between evidence and theory and thus nothing to underpin reliance on theories.

Nevertheless, Popper argues that, while it is not rational actually to *rely* on any theory, it is entirely rational to *prefer* the best tested theories: theories that have so far survived rigorous attempts to falsify them. This last point is sensible advice from our normal induction-trusting perspective. It is consistent with the place of rigorously conducted research at the top of the evidence hierarchy in EBM (p. 436 above, this chapter). A study in which the researchers are 'blind' as to test and control results is more likely to detect a false theory (e.g. that a new medication is more effective than an established treatment) than a study in which the researchers know which is which and hence may (consciously or unconsciously) bias the results in favour of the theory they believe to be true. In this respect, then, Popper's observations on scientific methodology, are surely right. Science does proceed more by falsification than by confirmation. This is one powerful way in which we may genuinely learn from experience rather than being misled by our expectations.

It is worth asking, however, why, from the point of view of a purely falsificationist perspective, is it better to prefer a theory that has been severely tested and survived than any other theory that would have passed those tests? There will, after all, be an infinite number of these. Some can be generated by simply conjoining, with a previous theory, some additional claims about the future. Furthermore, this preference for specific theories can never amount to a justified reliance on them given the anti-inductivist stance of falsificationism. But how plausible an account of either science or our daily life is that? We do not merely *guess* that the roof will not collapse in the next few minutes. Much of clinical practice, for example, rests on lengthy trials where success in replication of results is as important as the elimination of false hypotheses. A Popperian can only explain the negative results of trials—showing that particular treatments do not work and can thus be rejected—but cannot offer an explanation of why repeated successful results should be taken as positive support for efficacy. On a strict falsificationist view of the matter, that is irrational.

Conjecture and refutation

As we discussed in Chapter 11, at the heart of falsificationism is the idea of conjecture and refutation. There are never *good reasons* to adopt or hold a theory. Reason works only negatively: pruning away at the bold conjectures that scientists advance. Once falsified by disconfirming evidence, theories should be rejected.

But as Popper himself realized things are more complicated than that in actual scientific practice. The observations that are to serve as tests of theories are themselves laden with theory. Successful theories often face disconfirming anomalies that take time to explain away. Theories are articulated within competing broader explanatory stances or programmes.

Consider that second point: theories are often born refuted but are not rejected for that reason. This is an important point in medical research. It is often very difficult to establish connections such as that between smoking and cancer. From a naïve falsificationist perspective a negative study falsifies or refutes a putative connection. However, that is not how good medical science progresses. Such results might themselves be scientific mistakes in hindsight. To be a realistic model of scientific practice, the falsificationist model of conjecture and refutation has to be augmented.

Lakatosian falsificationism

Such augmentation was provided in the work of the Hungarian philosopher of science, Imre Lakatos. Lakatos' model of scientific research programmes has been widely influential both in its own right and as an exemplar of the power of studying the history and philosophy of science as a combined discipline. The philosophical model was taken to provide the structure for a rational reconstruction of the history of science, demonstrating the rational methodology of science at play.

The key statement of Lakatos' position is given in his paper 'Falsificationism and the methodology of scientific research programmes' (Lakatos, 1970).

Three theories that presuppose hard data

Lakatos begins his discussion with three ways of understanding the relation of theory and evidence: justificationism, probabilism, and dogmatic falsificationism. These three are similar in their commitment to a class of hard data that can be used in one way or another as the basis for theories. In the first and second, observations are used either to confirm the theory (in the first) or at least to make the theory more probable (in the second). Dogmatic falsificationism is based on the logical point that Popper raised: that a no finite number of individual facts can logically or demonstratively confirm a universal claim (such as a statement of theory), but, in principle, a single fact can refute a universal claim. Thus falsification proposes that the purpose of evidence in theory testing is negative: the falsification by refutation of theories so as to eliminate false theories.



Fig. 16.2 Imre Lakatos

But are there any hard data?

But as Lakatos goes on to argue, this picture is too simplistic to characterize real science. The main problem is that rather than comprising an infallible class of foundational claims, observational claims are themselves fallible and theoretically charged, as we saw in Chapter 12. Lakatos argues for this by denying two claims that he suggests are implicit in dogmatic falsificationism:

1. there is a natural and psychological distinction between theoretical claims and observational claims, and
2. observational claims must be true.

His argument against (1) recapitulates some of the arguments from Chapter 12: observational claims presuppose the truth of theories about observational aids or psychology.

Can experience justify an observation statement?

The argument to deny the second claim is less convincing. The conclusion is that no factual proposition can ever be *proved* from an experiment. Lakatos's argument for it runs as follows: 'propositions can only be derived from other propositions, they cannot be derived from facts: one cannot prove statements from experiences—"no more than by thumping the table" (p. 99).

The underlying idea here is that if one wishes to prove a proposition, only other propositions have the right logical properties to stand in justificatory (or for that matter contradictory) relations to it. Events, such as a table being thumped, do not and neither, Lakatos suggests, do experiences. From this Lakatos concludes

that factual propositions cannot be proved from an experiment (and neither could they be disproved).

The problem with such a line of reasoning is that it makes the contribution of experience to the justification of our beliefs completely mysterious. What is the point of our opening our eyes in the first place? The problem of explaining how experience can have a *rational* impact on our beliefs is the key issue in McDowell's *Mind and World* (1994), discussed in Chapter 12. McDowell, you will recall, argues that our conceptual abilities are always passively drawn into play in even the most basic cases of experience. Thus experiences can have the right logical 'shape' to confirm or undermine beliefs because they are always already conceptualized. But lacking any such account of how experience can rationally constrain beliefs, Lakatos can say the right thing about observation claims—that they are fallible—but for the wrong reason: namely that he can say nothing sensible about the *groundings* of observation claims in experience.

One can, however, deploy a different argument to Lakatos's conclusion based on his first point. If observational claims presuppose the truth of theories, and if the only evidence for theories are observational claims, then there can be no definitive test either of theory or of observational statements. The support for both will be provisional.

Lakatos goes on to argue that even if these objections to a distinction between theory and observation were not true, dogmatic falsificationism would still fail to take account of a different connection, which makes falsification much less clear-cut than simple examples (like those provided by Popper) suggest.

EXERCISE 5

(20 minutes)

Look at the extract below from

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In *Criticism and the Growth of Knowledge* (ed. I. Lakatos and A. Musgrave). Cambridge: Cambridge University Press, pp. 91–195 (extract pp. 100–101)

Link with Reading 16.4

Here Lakatos tells a story about an imaginary 'misbehaving' planet and how scientists respond to it.

- ◆ What general conclusions can be drawn from the story?

A radical conclusion for refutation

Lakatos argues that the observed planetary 'misbehaviour' might be taken as a refutation of the existing Newtonian celestial mechanics but on the other hand it might be taken as evidence of the perturbing effect of another planet. Such a hypothesis suggests possible crucial observations of the postulated planet concerned. But if no such planet is observed, this still need not refute the theory that there is such a planet (itself deployed to defend existing celestial mechanics), which can in turn be resisted by

deploying the further hypothesis that there is opaque intervening material. And so on ad infinitum. In other words, *no apparently falsifying evidence is ever decisive*. Lakatos concludes that: 'Scientific theories are not only equally unprovable, and equally improbable, but they are also equally undisprovable.' (p. 103).

Rejecting dogmatic falsificationism, Lakatos goes on to consider two further revisions of falsificationism each of which provides insights into the nature of scientific methodology, though neither of which, he argues, sufficiently describes it. The two further versions are:

- ◆ Naive methodological falsificationism: this introduces the idea that observations and theory have to be considered holistically in science, i.e. as working together in groups rather than as isolated observation theory pairs.
- ◆ Sophisticated methodological falsificationism: the key idea here is of what Lakatos calls a 'research programme', i.e. groups of observations and theories operating over an extended period of time.

We will look briefly at each of these. Interestingly, Lakatos attributes all three versions of falsificationism (dogmatic, naive methodological, and sophisticated methodological) to Popper. His conclusion is that sophisticated methodological falsificationism gives the 'best fit' to what goes on in science. This conclusion, however, as we will see has been challenged by both historians and sociologists of science.

Naive methodological falsificationism ...

Rejecting dogmatic falsificationism, Lakatos introduces the second of three pictures of science he discusses: naive methodological falsificationism. Lakatos attributes this position to some of the writings of Popper, although he later goes on to argue that Popper subsequently developed a yet more sophisticated methodological falsificationism (the third of Lakatos' pictures of science). Surprisingly (or perhaps not given the antecedent arguments), Lakatos reports that this is a form of *conventionalism*.

... is a form of conventionalism

This is a surprising concession because it appears to give up a key feature of empiricism: that scientific theories answer to the facts. Conventionalism in the philosophy of science is (in general) the label for a view of scientific theories in which they are taken to be adopted for their usefulness and convenience without any regard for their truth. Think of the convention that everyone should drive on the same side of the road. This is useful but in no sense answers to a Platonic 'fact' of motoring and is itself neither true nor false. Conventionalism is thus antirealist about scientific methodology. It is also (though again only in general) antirealist in the metaphysical sense: that is to say, conventionalism denies that it even makes sense to speak of scientific theories being true of the world.

Lakatos's subscription to conventionalism is not so radical, however. Conventionalism enters Lakatos's account of methodological falsificationism through two *decisions* that correspond to the two assumptions of dogmatic falsificationism discussed above. Some statements are deemed observational and thus themselves immune from falsification *by decision* (p. 106). This is a matter of using successful theories as extensions of our senses by allowing them to go unchallenged when they are used to challenge other theories. Thus, rather than relying on a decisive foundation (whether used to confirm or refute theories), methodological falsificationism is based instead on 'piles driven into a swamp' (p. 108). The foundation of observation is only relatively stable rather than completely grounded. Thus even if 'falsified' by statements deemed observational, a theory may still be true. Lakatos mentions as an example of historical rashness the fact that Galileo and his disciples accepted 'Copernican heliocentric mechanics in spite of the abundance evidence against the rotation of the earth' (p. 115). In fact it turned out that that 'evidence' was misleading and Galileo was right to adopt the heliocentric view. Thus the rational advice—that falsified theories should be rejected—is fallible advice. It may result in the rejection of true theories.

Without a firm empirical basis of incontestable observation statements, the original falsificationist demarcation of science from non-science is undermined. However, Lakatos suggests that a conventionalist reworking of it is still acceptable. Theories will, however, only be 'falsifiable' in the light of methodological prescriptions about putting them to the test (cf. p. 111). A theory is scientific if it is 'falsifiable', but this now requires a kind of intellectual honesty from scientists to treat it as falsifiable and to test it accordingly as in principle the world can never decisively show that it is false.

Sophisticated methodological falsificationism

Despite the fact that naive methodological falsificationism is an improvement over dogmatic falsification in its attempt to come to terms with the holism implicit in theory testing, Lakatos suggests that it is still an implausible account of actual scientific practice. This leads him to suggest his third falsificationist model of scientific rationality, sophisticated methodological falsificationism.

The key development in moving from naive to sophisticated methodological falsificationism is to take the unit of assessment in science not as an individual theory standing in relation to evidence, but as a series of theories grouped within what Lakatos calls a research programme. A research programme (i.e. a series of theories) stands in relation both to the evidence and also to competing research programmes. Assessment is thus explicitly temporally extended.

Within a research programme, one theory can replace another if two conditions are satisfied:

1. the new theory predicts facts disallowed by the former theory while explaining the successes of the former theory, and

2. some of its 'excess content' (i.e. the additional facts predicted by the new theory) has passed empirical tests.

Within a research programme, therefore, theories are falsified *by other theories* in the light of evidence, rather than, as in Popper's original and simple formulation, by the evidence directly. A research programme, i.e. a series of theories, is said by Lakatos to be *progressive* if both the above conditions are met and to be *degenerating* if not. Crucial experiments are thus not just any experiments that produce results that are inconsistent with a research programme (Lakatos calls such results 'anomalies'), but only those that distinguish between competing theories in a series.

Research programmes: hard cores and protective belts, positive and negative heuristics

A research programme, as envisaged by Lakatos, comprises a hard core of central assumptions that are shared between different theories in a series; and also a 'protective belt' of further assumptions. The latter are those further theories and hypotheses that are needed to relate the hard core to observations.

For example, current research on changes in neurotransmitters in the brain in conditions such as depression is a research programme, in Lakatos' use of the term. There are many and competing particular theories of how this or that change in this or that neurotransmitter might be related to depression. But all these theories share, as one hard core of assumption, the belief that neurotransmitters and mood are related. This assumption is surrounded by a penumbra of other beliefs about, for example, neuroimaging techniques, particular brain chemistry and so on. This means that if, for example, a hypothesis about the connection between a particular mood and neurotransmitter were to fail to gain experimental support it could be explained as simply the wrong choice of mechanism, the failure of the imaging technique or whatever without being taken to cast doubt on the general idea that neurotransmitters and mood are related.

Lakatos then specifies two methodological rules: the positive and negative heuristic. The negative heuristic is the prescription that the hard core should be preserved even in the face of observational anomalies. The positive heuristic determines the kind of changes that should be made when there are anomalies. Thus the negative heuristic determines the continuity between different theories in a research programme in that it ensures constancy of the hard-core assumptions. This is deemed irrefutable by empirical evidence: a more or less conventionalist decision. The positive heuristic—which remains sketchier in Lakatos's account—concerns how potential anomalies are to be dealt with: what sort of auxiliary assumptions are built into the refutable 'protective belt'.

We can see both heuristics at work in our example (above) of the current research programme on neurotransmitters and depression. Thus, the negative heuristic protects the programme from evidence (such as that produced by Brown and Harris, see end of Chapter 15) that social factors have a role to play in depression: such factors, the negative heuristic specifies, must be

understood as being mediated by neurotransmitters. The positive heuristic, in this case, concern, among other things, the general shape of mechanisms that might be appealed to. The idea, for example, that the mind has no physiological correlate and might float free of it will receive short shrift.

The history of science is, according to Lakatos, a history of the development of theories within a larger structure of competing research programmes. Research programmes are the units that progress or degenerate—according to how theories are developed in line with evidence—rather than the individual theories themselves.

Rational reconstructions and the history of science

Lakatos's approach is thus particularly interesting in the way it takes the simple logical point about the asymmetry of refutation and confirmation of universal claims by single instances, which Popper (1972) highlights, and shows that this does not by itself yield a plausible scientific methodology. In attempting to preserve Popper's emphasis on the negative use of evidence, Lakatos develops a model of a historically extended process in which groups of theories are assessed holistically, i.e. in relation to each other and to the data, which is itself regarded as fallible and to be assessed in the light of theoretical context. But as a result, the methodology cannot give decisive advice about when a theory should be accepted or rejected. In the main, it is better to support progressive rather than degenerating research programmes. However, exactly when it is appropriate to abandon a degenerating research programme is a matter of uncodifiable *judgement*. There are no settled criteria for this. A programme could be degenerating for any length of time before again making progress.

Lakatos's philosophical model gave rise in the 1970s to a project of writing rational reconstructions in the history of science. Effectively this was a way of viewing the history of science as exemplifying Lakatos's model of what is rational in the pursuit of science. Events that fitted the model of rational behaviour for scientists were thought to require no further explanation. Only events that did not fit that model called for further historical or sociological explanation. But as we will see below, this approach faces criticism.

Kuhn's account of the history of science

Thomas Kuhn's paradigms

Lakatos's (1970) methodology of scientific research programmes is an attempt in part to articulate a rational method of theory choice. Within a single research programme, theories are preferred if they meet the tests of sophisticated falsificationism while programmes as a whole are preferred if they are progressive.

But the view that there is any such clear rational method by which proceeds had already received sustained criticism from the historian of science Thomas Kuhn. Kuhn was originally a physicist who became interested in the history of science and wrote a very influential book, *The Structure of Scientific Revolutions* (1962), published shortly before Lakatos' definitive statement of

his own work described above. In it he sets out an account of normal scientific practice and occasional scientific revolution. Over the course of the next two decades a broadly Kuhnian view prevailed in the philosophy of science at least to the extent that conjoined history and philosophy of science departments carrying out Lakatosian 'rational reconstructions' has waned.

Kuhn argues that the kind of radical theory change emphasized in the starkest kind of falsificationism (if not so much in Lakatos's version) is the exception rather than the norm in the history of science. For most of their time, scientists engage not in the critical testing of current theories but in the gradual extension and application of them through 'puzzle solving'. Kuhn calls such activity 'normal science' and the taken-for-granted background of theory the dominant 'paradigm'. (In fact he uses this word in a number of different ways. One common effect of reading Kuhn is the subsequent indiscriminate use of the word 'paradigm'. Avoid it!)

In Chapter 14 we discussed Kuhn's account of the importance of *tacit* knowledge and agreement, which, according to Kuhn, partly constitute a shared paradigm. It is made up not only of explicit high-level theories but also implicit factors. These include metaphysical commitments about the kind of theoretical description of the world that is acceptable (in terms of particles or fields, for example) and implicit standards for what a satisfactory level of agreement is between theory and observation as well as tacit knowledge about how to solve theoretical puzzles. But it should be clear from this brief sketch that a Kuhnian paradigm is quite similar to a Lakatosian research programme with its hard-core of scientific and metaphysical claims.



Fig. 16.3 Thomas Kuhn

Dominant paradigms can, however, be overturned in scientific revolutions. But although Kuhn gives a helpful account of the sorts of factors responsible for such revolutions—a growing number of unsolved puzzles, dissatisfaction among leading scientists with the paradigm, the development of alternatives, and so forth—the revolution itself, he suggests, is not subject to a rational methodology. Furthermore, and perhaps more contentiously, he argues that there are no rational measures by which to judge that science is cumulative or progressive to use Lakatos's phrase. The range and nature of puzzles after a revolution is merely different from that before and both are determined by the different paradigms operative at the time. Kuhn famously, or perhaps notoriously, even suggests that after a revolution scientists are effectively living in a different world.

EXERCISE 6

(15 minutes)

Read the extract from:

Kuhn, T.S. (1970). Logic of discovery or psychology of research? In *Criticism and the Growth of Knowledge* (ed. I. Lakatos and A. Musgrave). Cambridge: Cambridge University Press, pp. 1–23 (extract 4–6)

Link with Reading 16.5

Outline the key differences between Kuhn's approach and those of Popper and Lakatos. Is the difference merely one of emphasis?

The role of testing for Kuhn and Popper

In an introduction just before this extract, Kuhn (rather misleadingly) emphasizes the agreement between his view of science and that of Popper. Here, though, he takes as a first instance of their disagreement, the role of empirical testing of hypotheses. Kuhn suggests that testing does play an important role in science but not the one emphasized by falsificationism (p. 4).

Thus, within what Kuhn calls 'normal science', researchers attempt to solve 'puzzles'. Hypotheses are proposed to explain specific natural phenomena and then tested. However, such tests are not designed, as Popper supposed, to put 'maximum strain' on the overall theory or paradigm. To the contrary, within normal science the background theories and assumptions are presupposed in the very way the problems are defined and in the constraints placed on what an adequate solution or explanation would be. Solving puzzles is the most common activity of working research scientists and it is the method through which a dominant theory or paradigm is extended. The working assumption is that puzzles can be solved *within* the current theoretical setting. Failure is a mark of the inadequacy of the practitioner rather than the theory.

This emphasis on puzzle solving enables Kuhn to offer a rival to Popper's account of what is distractive about scientific method. Popper demarcates science from non-science by falsifiability of theories. Astrology, then, for Popper, is not scientific because it's

theories are not falsifiable. For Kuhn, by contrast, astrology is not scientific because the rules that constitute the paradigm of astrology do not determine any puzzles that could be solved. 'The occurrence of failures could be explained, but particular failures did not give rise to research puzzles, for no man, however skilled, could make use of them in a constructive attempt to revise the astrological tradition.' (p. 9).

Learning from one's mistakes

Kuhn then goes on to question the Popperian idea that a central aspect of scientific development is learning from one's mistakes. Kuhn argues that this idea makes sense within the context of normal science in which a mistake is a failure to follow the rules laid down by accepted theory and practice. But Popper attempts to apply that notion to *revolutions* in science and Kuhn argues that no sense can be attached to the idea that an abandoned theory—such as Ptolemaic astronomy—is mistaken. That would only make sense if there were an inductive logic by the rules of which theories could be shown to be good by induction from the data. But as there are no such rules, no theory is simply mistaken.

In addition to attacking the role of learning from mistakes, Kuhn also presses the point that has been discussed already in the context of Lakatos's work. Despite the attractions of naive falsificationism, the refutation of theories is not the straightforward business presupposed by that approach 'Rather than a logic', Kuhn suggests, 'Sir Karl has provided an ideology; rather than methodological rules, he has supplied procedural maxims' (p. 15).

How, then, if there is no such inductive logic, does science work? Kuhn's answer is to deploy an idea based on a different sense of the word 'paradigm'. Rather than thinking of science as turning on logically explicit theories, Kuhn suggests that it turns on implicit generalizations based on particular instances or 'paradigms'. This view, which emphasizes the importance of open-ended or open-textured concepts, also suggests that the response to anomalies is less determined than a view of scientific theories as strict universal generalizations would imply. Indeed, Kuhn goes so far as to suggest that some of the rhetoric of the Popperian account is an apt expression of the methodological imperatives that make up the tacit values of a paradigm. Such values, he concludes, should be understood through social psychology rather than through a logical methodology of theory change of the kind proposed by Popper.

Overview of Kuhn on theory and evidence

Kuhn's overall picture of science is clearly distinct in emphasis from that of Popper or even Lakatos. He suggests that the structure of science is one of stable normal research work interrupted by occasional radical revolutions. It is not a matter of constant revolution as described by falsificationism. He does, however, share with sophisticated forms of falsificationism a view that science comprises a complex structure of theory and observation against a background of taken for granted assumptions about the shape of any plausible theory. Empirical testing is thus not a

matter of mere data gathering but involves much background theoretical presupposition. Without such a background, theory testing would be impossible. But for most of the time, the background itself is simply not called into question. Furthermore, much of this background is also tacit. Practitioners learn tacit skills in the practical manipulation of instruments, for example. They also learn to make assumptions about what factors do and do not matter in particular empirical contexts. In Kuhn's account, furthermore, as in Lakatos's, there is an important balance between loyalty to existing theory and a preparedness on occasion for radical revision when the data are in conflict with theory.

Kuhn's account of how science works is based on much historical analysis. Thus the burden of proof is carried by historical evidence rather than a priori reasoning. But Kuhn's views are shared and further developed in the work of contemporary sociologists of science. It is to the sociology of science, and the further debate it has stimulated about the role of rational methodology in the selection and testing of theory, that we turn in the last part of this session.

The sociology of science

The short extract (linked with Exercise 7 below) is part of a chapter from an influential textbook on the sociology of science written by David Bloor (1976). Although a sociologist, Bloor has also written philosophical papers on Wittgenstein. Unsurprisingly he argues for a reading of Wittgenstein that emphasizes the importance of *social* relations for understanding language and cognition (Bloor, 1997).

The main objective of Bloor's (1976) chapter is to set out a programme for the sociological investigation of science. It is thus concerned with rival models for the explanation of scientific development. But it also relates to the concerns of this chapter and session because it articulates a rival view of scientific rationality. Recall that Popper claimed for his 'solution' to the problem of induction the merit that it preserved a rational method for science. However, that claim was weakened in the modifications that Lakatos made to falsificationism, which were required to make it a plausible account of good but actual scientific practice. Bloor implicitly undermines the idea that there is *any* substantial universal model of scientific rationality. That is why science requires piecemeal and local sociological investigation.

EXERCISE 7

(15 minutes)

Read the extract from:

Bloor, D. (1976). *Knowledge and Social Imagery*. London: Routledge, (extract pp. 1–3).

Link with Reading 16.6

- ◆ What arguments does Bloor direct against Lakatos's ideas about the role of a philosophical model of scientific rationality?
- ◆ What relationship between a sociology of knowledge and a philosophy of science is suggested in Bloor's work?

Bloor's view of the role of sociology

Bloor sets out the challenge for sociology to explain the content of scientific knowledge. He argues that this is a natural extension of the area of competence of sociology and that it is merely a failure of nerve that has prevented sociologists from taking up this challenge in the past.

It is worth briefly reflecting on what is at issue here: Bloor thinks that there can be a sociological explanation not merely of the institutional context of scientific research, nor merely of the cultural impact of knowledge claims, but of those very claims themselves, of their very content.

This suggests (and this suggestion is borne out by subsequent sociological work) that a characteristically realist explanatory strategy is also being ruled out along with the more explicit attack on this reading on a form of a priori rationalism. The realist strategy is to explain the content of a knowledge claim in part at least by saying that it is true or that it is a fact. So in answer to the question: why did nineteenth century scientists accept such and such a claim, the realist replies that they accepted it because it was true. Bloor rejects this reply. He proposes instead a four-point characterization of what he calls the Strong Programme for the sociology of science, which is roughly:

1. It would be causal, that is, concerned with the conditions that bring about belief or states of knowledge.
2. It would be impartial with respect to truth and falsity, rationality or irrationality, success or failure. Both sides will require explanation.
3. It would be symmetrical in its style of explanation.
4. It would be reflexive. (cf. pp. 4–5)

Bloor goes on to contrast this view with the view of the relation of sociological or historical explanation and rationality that Lakatos develops. For Lakatos (1970), there is a sharp divide between 'internal' and 'external' history. The internal history of science is a demonstration of how scientific development accords with a philosophical model of scientific rationality—a model such as Lakatos's own methodological falsificationism. Sociological analysis is only needed as part of external history: the history of irrationality and failure. The underlying idea is that rationality is its own explanation. What needs concrete or substantial explanation is only deviation from this course.

Bloor characterizes Lakatos's view as a *teleological* view by contrast with his own *causal* view. The idea is that a teleological conception relies on the assumption that the rationality of an action or belief draws scientists on towards it. Such terminology needs some care, however, as most philosophers who subscribe to something like Lakatos's view here would also subscribe to a causal view in the philosophy of mind. And, again, most philosophers would agree that rationality plays a key and constitutive role in the explanation of behaviour. Bloor plays this down, but in fact nearly all sociological accounts help themselves to something like the everyday folk psychological pattern of rational action

explanation. However, Bloor wants to distance himself from the idea, implicit in Lakatos, that nothing else needs to be said to explain action if it can also be said to be rational. This point is made even more clearly in the writing of another sociologist: Barry Barnes.

The Barnes view

Barry Barnes, like Bloor, a sociologist of science, wrote, at about the same time as Bloor, in support of more or less the same views (Barnes, 1974). Barnes, however, is more explicit about the connection between the problem of induction, the assumption among philosophers of science that there is a model of scientific rationality, and the rival sociological conception of scientific knowledge.

Barnes, like Bloor, contrasts causal analysis and rational explanation of action; however, in his case the underlying point of the contrast is clearer. The key problem of invoking the rationality of a belief or action in order to explain it, is not so much that such an explanation is false as that it begs the question. Barnes illustrates this point with the poison oracle among a tribe called the Azande. A chicken is given poison and a question is asked. If the chicken dies, the answer to the question is 'yes'. The irrationality of the oracle may seem obvious to twentieth century Western society. But Barnes asks us to consider how we would explain the actions of a member of that society who rejects the status of the oracle and subscribes to something more like Western thinking. Would that person's actions require no further explanation because his or her beliefs were (by twentieth century Western standards) rational? Barnes replies that precisely their deviance from (local) societal norms would, all the more, *require* explanation.

In other words, the problem with invoking a teleological conception of rationality is fleshing out what it involves. In a particular context, reasoning in a particular way that diverges markedly from our twentieth century views, may be rational. It may make sense to follow patterns of thought which we, initially at least, find alien. So the right conclusion to draw is not that sociological explanation flies in the face of rationality (as some sociologists of science sometimes appear to suggest), but rather that the content of acting or thinking rationally has to be unpacked in each context.

Bloor, Barnes, and delusion

As an aside, it is worth noting that neither Bloor nor Barnes, although sociologists, had, in the 1970s, anything more to say about the various forms of irrationality presented by psychopathology, than philosophers (see the quote from Anthony Quinton at the start of Chapter 2). Yet here, as in so many other contexts in philosophy, psychopathology, in all its diverse forms, presents, as Austin (see Chapter 5) and more recently Wilkes' (1988) have argued, a peculiarly sharp test of theory. Barnes (1974) might, perhaps, have pointed to delusions as an example of (extreme) departures from local norms of rationality. But as we saw in Chapter 3, it is very far from clear by precisely what standards such a departure is to be judged. Certainly, there is as

yet no causal story of the kind Bloor and Barnes envisaged, or indeed of any other kind, by which delusional and non-delusional beliefs are to be distinguished. Yet a distinction is made. Similarly, considerations will arise in the next section, when we consider in more detail the metaphysical claim that knowledge is justified true belief: the existence of delusions which are true and yet (in some deep but thus far unexplained sense) unjustified beliefs (see Chapter 3), has simply not been recognized, still less engaged with, still less explained, by philosophers interested in the nature of knowledge.

Content and context and some preliminary conclusions for evidence-based medicine

In this session we have seen that Popper (1972) and Lakatos (1970) put forward a model of what a universal scientific rationality would amount to. Kuhn, as a historian as well as a philosopher of science, disagrees with Popper and Lakatos and argues that there is no such context-independent rational account of science to be had. Barnes (1974) and Bloor (1976) with their 'Strong Programme' go further and propose a causal, rather than a rational, account of science that seeks to explain, even-handedly, successful and unsuccessful outcomes of scientific research.

The debate does not show the model of scientific rationality has nothing to offer. Nor is it that a more detailed 'causal' explanation should eschew rational explanation. It is rather that the problem with any substantial but imposed model of scientific rationality is that it can fail to deliver the right results in context. There may be good reasons, for example, not to reject a theory in the face of refutation if it has accommodated similar anomalies in the past and is otherwise reliable. This is the sort of thing that Lakatos attempts to accommodate. But because he attempts to describe a rational method in a context-free manner it opens his model up to the charge that it is not substantial: its advice verges on 'Do the right thing!'. The sociologists' arguments are best interpreted as saying that the content of what is rational in any particular context has to be investigated piecemeal.

This is not to undermine the rationality of the best scientific practice. To some extent, scientific practice just is our model of rational decision making. But the attempt to codify what is rational in a universal a-historical and context-free manner has not proved successful. That provides an important lesson for EBM's more modest attempt to codify how best to learn from experience. Careful attention to the details of the traditional model of science suggests that it cannot give a detailed account of scientific practice. This does not undermine science, however, but show that we have a tendency to oversimplify its workings. The traditional model has to be made complicated especially in areas of science, which are themselves complex. Thus we should be wary of oversimplifying scientific practice in accounts of EBM.

Consideration of the philosophy of science has shown that whatever role induction from evidence does have in science—and if nothing else falsificationism shows the importance of refutation as well as confirmation—there is much more to the

role of evidence and theory testing than a naively inductivist view would have us believe. Scientific research does not involve the harvesting of particular observations from which grow generalizations and eventually theories. As Chapter 12 already argued, evidence presupposes much theory. But as the historically informed accounts of science have plausibly shown in this session, the way that evidence is used is also subject to background theoretical considerations. The relation between theory and evidence is one of interdependence: both top-down and bottom-up.

Still, while Barnes (1974) explicitly discusses inductive inference as suitable for sociological investigation (as did Harry Collins discussed in Chapter 14) neither the sociology of knowledge, nor Kuhn's historical overview, nor falsificationism, provide a resolution of Hume's problem. One feature of the sociologists' rejection of a teleological construal of rationality is that they eschew a justificatory component in favour of description. Thus they provide no response to whether or why inductive inference is justified. For this reason, the next session will return to the problem of induction and responses to it made within philosophical epistemology. This will shed light on the nature of knowledge itself and its social dimension.

Reflection on the session and self-test questions

Write down your own reflections on the materials in this session drawing out any points that are particularly significant for you. Then write brief notes about the following:

1. What kind of response to Hume's problem of induction is suggested by the various readings in this session?
2. How does Popper address the problem?
3. What are the weaknesses of a naïve falsificationist model of scientific rationality?
4. What lessons are suggested by Kuhn's model of historical and other sociological approaches to scientific method described here?

Session 3 Epistemological responses to the problem of induction

Epistemology versus philosophy of science

In this session we will consider three responses to the problem of induction offered from within the perspective of epistemology rather than the more methodologically based philosophy of science. One difference in focus is that the approaches in this session attempt to resolve or dissolve the problem of induction rather than attempting merely to live with its sceptical consequences.

The discussion of philosophy of science responses in the previous session helped emphasize that the relation between theories

and evidence is more complex than one might at first have thought. By considering Popper (1972), then Lakatos (1970), Kuhn (1970), and the Strong Programme (Barnes, 1974; Bloor, 1976), one theme emerged whatever the general conclusions one draws for the nature of scientific methodology. The relation of theories and evidence has to be understood in a broader context of scientific research whether one thinks of this as comprising research programmes or paradigms or historically situated forms of rationality. So equally citation of evidence in medicine turns on a similar background of higher and lower level scientific and metaphysical assumptions whether or not these are made explicit.

This session will consider lessons learnt from epistemology. One lesson in particular will be that knowledge can be transmitted from one individual to another without the second person being able to provide an argument to ensure that it is watertight. This is a lesson that we often forget when thinking about knowledge in a philosophical context—when we come over 'all philosophical'—and yet it informs our everyday use of the term. By examining the extent to which individuals are actively responsible for what they know we will shed light on what can reasonably be expected of EBM.

Again the route to these general views will start with a response to Hume's problem of induction. A good statement of a recent short response to Hume is given by a recent professor of philosophy: D.H. Mellor. His inaugural address aimed to defuse any lingering problem of induction. It is published as 'The warrant of induction' in his *Matters of Metaphysics* (1991).

Mellor's diagnosis of the problem of induction

EXERCISE 8

(15 minutes)

Read the two extracts from:

Mellor, D. H. (1991). The warrant of induction. *Matters of Metaphysics*. London: Routledge, (Extract sections 1 and 6)

Link with Reading 16.7

- ◆ What is Mellor's solution to the problem of induction?
- ◆ What view of knowledge does it presuppose?
- ◆ How does it differ from the views examined so far?
- ◆ Does it really solve the problem?

There can be knowledge by induction!

The most obvious contrast between Mellor's approach and those discussed in the previous session is that Mellor assumes that we can and do have positive knowledge. He says: 'This lecture will last for less than twenty-four hours. I know that and so do you.' (p. 254) The grounding of this knowledge is, according to Mellor, the directly observational or experiential knowledge we have that none, or almost none, of the previous lectures we have



heard has lasted that long. We go on to form the further inductive belief—on the basis of the observational properties that identify an event as a lecture—that it will also have the as yet unobserved property of being shorter than a certain time. (This is reminiscent of Hume’s talk of the connection between sensible powers and secret powers.)

The central question is *how* such past observations can warrant or justify claims about so far unobserved states of affairs. Mellor points out that such warrant is both weak enough to be overridden by observation—present observation can overturn a previous induction—and yet strong enough to trust one’s life to (he gives the example of believing that the building will not shortly fall in). But how is this possible? The answer, Mellor suggests, lies in the Cambridge philosopher Frank Ramsey’s (1903–30) suggestion that ‘our conviction [in inductive arguments] is reasonable because the world is so constituted that inductive arguments lead on the whole to true opinions’ (p. 254).

The structure of the paper

Mellor’s paper then has the following structure. Having first introduced the problem and a thumbnail sketch of the solution, it goes on to explore an analogy with the warrant for beliefs that direct perception provides (so as to set an appropriate standard for induction to meet). Mellor suggests here a form of reliabilism: observation is a reliable method of arriving at true beliefs. We are causally disposed by observation to form beliefs that have a high probability of being true. Mellor then goes on to explain the objective construal of chance implicit in this claim and to reject the assumption that one needs to know that one is warranted in one’s belief to be so warranted. Finally, he returns to the proposed solution to the problem of induction.

An analogy with observation

The analogy with observation is supposed to work like this. Observation is fallible. We can think that we see sparrows when we don’t. To remind us of this, Mellor adopts the potentially confusing form of words that to ‘observe’ does not imply the truth of the object of that verb, although he does not himself use inverted commas. (In everyday life observing or seeing is a *success* concept—like knowing—if one sees that *x* is the case, then it is; and if one sees a *y*, there is a *y* to be seen.) Given this form of words, the fact that one ‘observes’ a sparrow does not entail, or logically imply, that there is a sparrow there. (To repeat, in everyday life this is an entailment, but Mellor does not want to assume that we see a sparrow when we seem to see a sparrow.) So if it is not a logical truth that when we ‘observe’ a sparrow there is a sparrow, what warrant or justification does ‘observing’ provide?

Mellor suggests the following approach. Think of the judgement that there is a sparrow present when we ‘observe’ one as an inferential disposition: a disposition to make that inference. Now dispositions are well known items in our causal picture of the world. The standing property of mass has further dispositional

properties such as being disposed to accelerate at a certain rate when acted upon by a specific force. Masses embody a causal link between forces and accelerations. Mellor suggests that just such a causal link underpins the warrant that ‘observations’ provide for perceptually based beliefs. The presence of sparrows causes the *belief* that sparrows are present.

In fact, he suggests that this causal link is probabilistic. The probability of there being sparrows when one ‘observes’ them is less than 1 but sufficiently close to it still to warrant the belief. This requires that Mellor construes probabilities or chances as objective indeterminacies, rather than as epistemic measures, to avoid the charge of circularity. Otherwise he would be guilty of explaining the warrant of our beliefs in terms of the warrant of our beliefs. On his preferred theory, chances can also themselves have causal effects. One such example is the pattern of radiation that results from the chances of decay of radioactive substances. Another, Mellor suggests, is the formation of beliefs as the causal effects of a probabilistic observational mechanism.

Knowledge and reliability

With these suggestions in place, Mellor (1991) is able to suggest more broadly that what it is for a belief to have a warrant is for it to have a high probability of being true. ‘What better measure could there be of the prospects of truth which observation gives the beliefs it produces, and hence of how strongly it warrants them?’ (p. 261).

This is, in effect, a challenge for epistemology. Mellor suggests that the purpose of warranting or justifying beliefs is to provide good prospects for their truth. Nothing more is required.

You don’t need to know that you know. The reliabilist account of knowledge . . .

This runs counter to a different, more traditional, assumption that a further necessary condition for justification or warrant is that it is self-intimating. That is, if one’s belief is warranted, then one also knows that it is warranted. If one knows then one knows that one knows. Mellor raises two objections to this requirement. First, it leads to scepticism because the requirement escalates what one needs to know: one must know that one knows that one knows . . . ad infinitum. Secondly, given his own suggestion for the nature of warrants, the extra requirement would not increase the probability of truth of a belief already warranted.

It is worth pausing at this point. Mellor puts forward a view that has recently become much more popular in philosophy but can seem counter-intuitive. It is possible to know something even if one does not know that one knows it. Providing the source of one’s beliefs is as a matter of fact reliable—again, whether or not one knows this—then one can acquire knowledge.

Thus a psychiatrist following the principles of EBM can make an appropriate search of sources of evidence. Providing that the resources invoked are as a matter of fact reliable then he or she can acquire knowledge whether or not he or she knows anything



more about those sources of evidence. What matters is that they are reliable, presumably through the careful attention of other research workers.

The solution to the problem of induction?

Mellor is thus now in a position to set out his solution to the problem of induction. Forming beliefs by induction is itself an inferential disposition whose inputs are past observations and whose output is the inductive belief. How can these past observations warrant the subsequent belief when that belief is not true simply in virtue of those former beliefs? The proposal is that it can providing that the chance of the output belief being true given the input beliefs is sufficiently close to 1. Providing that there are laws linking observable properties then an inductive inference mechanism linking beliefs about one to the other will have a high probability of yielding true beliefs and thus it will be warranted. (By contrast counter-induction—see below—will not reliably yield true beliefs.) If on the other hand, there are no laws linking two observable properties then no inferential habit will reliably yield true beliefs (neither induction nor counter-induction will be reliable). Thus in a law-like world, induction unlike counter-induction is warranted. That is Mellor's solution to the problem of induction.

So to return to Mellor's initial example. Providing that we live in a world where the laws that govern social interaction, individual psychology and so on, which are applicable to the duration of lectures are stable then the fact that lectures have been shorter than 24 hours in the past warrants the belief that they will be in the future. The inductive leap from the past to the future is underpinned by the *de facto* regularity of nature.

One might now ask: but how do we know that the world behaves in this regular way? Is not that the real problem of induction? However, Mellor's point is that we do not need to know that in order to know that the lecture will last less than 24 hours. As long as the world is regular—whether or not we have taken a view on that—then we can have knowledge by induction.

Does it work?

Is this a satisfactory solution? It is worth returning to Mellor's first argument against the requirement that justifications or warrants should be self-intimating. As we saw in Session 1, the fires of Humean scepticism are fanned by just the requirement Mellor here rejects. If one needs to know that one knows in order to know, then the fact that one cannot know that induction is a reliable method of justifying beliefs implies that one cannot know anything by induction. Even if induction were reliable, the extra requirement rules out ground-level knowledge claims if one cannot—at the meta-level—vouch for induction. Mellor effectively blocks this swift route to scepticism. Even if we did not *know* that induction were reliable, providing that it were reliable, it would warrant our beliefs (even though unbeknown to us). And thus it may seem that scepticism is undermined.

Externalism

The denial that one needs to know that one knows marks a commitment to what is called in epistemology 'externalism'. Beware! This label is used in a variety of different philosophical contexts to mean different things. We will come across it again in Part V to refer to accounts of how meaning or intentionality come about. Here, it refers to the perspective from which knowledge is theorized about. Rather than thinking about the justifications that are *available to* a knowing subject, epistemological externalism characterizes the justificatory component of knowledge from a third person perspective. The article by McDowell discussed below is also written from this perspective, although it does not share Mellor's commitment to reliabilism. One of the motivations for epistemological externalism is that internalism seems to lead inexorably to inductive scepticism because of the regress of justifications for one's beliefs that one would need to know. But does externalism answer inductive scepticism?

Well one response is that it does not because the sceptic naturally ascends to a higher level with the question: 'Granted that you need not know that you know in order to know, but do you have inductive knowledge?' Mellor's solution is to say that providing induction is reliable, then we do have 'ground level' knowledge claims about the sun rising tomorrow, or bread nourishing, and we need not know that we know for this to be true. But Mellor admits that if induction is *not* reliable then we do *not* have ground level knowledge claims. So the sceptic can press the question: Which is it? Do we have ground level knowledge or not?

Mellor does not really answer this question. He can point out that even in worlds in which induction is not reliable, other strategies, including counter-induction—assuming that connections held in the past will *not* hold in the future—will also fail. So we may always be best off pragmatically by using induction. But that is not to say that we do have inductive knowledge. If we want a positive answer to this question, we want to know whether we know: to know that we know or know that we do not know. So although we need not first answer this question for us to have ground level knowledge, we do swiftly want to answer it when beset by inductive sceptical doubts. And it is not clear that reliabilism yields a positive answer here.

This is a potential problem for *any* epistemologically externalist account of knowledge. Although not knowing that one knows to the power n does not imply that one does not know that one knows to the power $n-1$, externalism by itself does not return a positive answer to the question of one's knowledge to the power n . And responding to inductive scepticism may make that the question we wish to be answered.

John McDowell offers a somewhat different response that aims to combine the virtues of an externalist approach while retaining the idea that knowledge has something to do with giving reasons. Looking to his account will provide a better way of thinking about knowledge and the extent to which one is responsible for one's own knowledge.

McDowell's diagnosis of the problem of induction

A different kind of externalism tied to the space of reasons

We have already come across McDowell's book *Mind and World* (1994) in earlier chapters (Chapters 12 and 15). Both those discussions emphasized the importance of the rational structure of reasons that McDowell, following the American philosopher Wilfrid Sellars (1912–89), calls the 'space of reasons' and which stands distinct from the 'realm of law'.

McDowell has also developed a sketch of knowledge that shares that emphasis on the role of reasons and that sheds light both on induction and the general context of learning from experience that applies to EBM.

McDowell's starting assumption

The key paper outlining McDowell's approach to knowledge is 'Knowledge and the internal' from *Meaning, Knowledge, and Reality* (1998 pp 395–413). It begins with an important, although gnomic comment. McDowell says that he will assume that knowledge is a 'sort of standing in the space of reasons' but will explore how that insight can be distorted in some philosophical accounts. It turns out that a key source of distortion is specific sceptical arguments and in particular, Descartes' argument from illusion (discussed below). Such arguments lead to the mistaken supposition that we ought to be able to achieve the standing in the space of reasons that amounts to knowledge 'without needing the world to do us any favours'. McDowell aims to give us a better picture of the nature of knowledge by diagnosing and rejecting Cartesian assumptions that are often made.

The distorting effect of the argument from illusion

Descartes' argument from illusion highlights the fact that sometimes when we take things in the world to be thus and so on the grounds of how they look, they are not thus and so. Sticks can look bent in water even when they are not. Thus if things are indeed as they seem then the world has done us a favour. Descartes used the ever present possibility of illusion to cast doubt on our ability to acquire knowledge. (This was a first step in his seminal epistemological enterprise that eventually sought to regain knowledge through a religious turn.). Even if we are not deceived in particular cases, is that not merely a matter of luck? And if so, does that not undermine the right to call any such perceptual beliefs 'knowledge' even if, as a matter of fact, they are true?

Responding to this worry we might try to build an account of knowledge restricted to epistemological states whose flawlessness we can ensure without this apparently problematic dependency on favours from the world. We might, in other words, commit ourselves no further than that it *looks or seems* to me as if things are thus and so. McDowell calls this philosophical move 'the interiorization of the space of reasons'.

The key assumption is that theorizing about knowledge should be carried out through description of states that do not depend on the world. Descartes and many epistemologists who have followed him hoped that the retreat to how things seem was only a temporary move. They hoped that there would be some way to regain the idea of knowledge claims about the world rather than just how it seemed. However, McDowell points out, this hope has generally proved futile.

If one adds to talk of how the world seems some further reasons or justifications of the fact that in specific kinds of circumstances such 'seemings' are reliable, then one will be able to say that the world is thus and so in a way that does not depend on the world doing us a favour. McDowell argues, however, that unless reason can equip us with policies that are utterly risk free, such a reconstruction will never add up to knowledge of the world: 'if one's method falls short of total freedom from risk of error, the appearance plus the appropriate circumstances for activating the method cannot ensure that things are as one takes them to be.' (p. 399).

Four responses to the argument from illusion

McDowell suggests that there are four responses to this problem. He himself advocates that we stop thinking about knowledge via the argument from illusion. But there are also three others:

1. Scepticism: we could simply conclude that the conditions for knowledge cannot be met for perception.
2. A form of a priori-ism: we could assume that there simply must be a priori risk-free policies for judging when appearances are reliable; however, this remains implausible.
3. A hybrid approach: we could retain the justificatory approach to knowledge implicit in talk of the 'space of reasons' still construed in an interiorized way but add a further component to it. On this approach knowledge has two ingredients: an epistemic standing or justification construed in the interiorized way as wholly under the subject's control plus another external condition not ensured by that: that the world does one a favour in being arranged as it seems to be. This further condition is truth. And the interiorized condition, the justification, concerns the 'reliability in a policy or habit of basing belief on appearance' (p. 400).

A contrast with Reliabilist Externalism

McDowell points out that, in the third approach, describing the reliability of one's policies as an internal matter goes against a standard view of reliabilism, as set out, for example, in the discussion of Mellor (1991) above. But McDowell rejects Mellor's sort of position because it fails to address the rational interconnected nature of our knowledge claims.

According to a full blown externalist approach, knowledge has nothing to do with positions in the space of reasons: knowledge is a state of the knower linked to the state of affairs known in such a way that the knower's being in that state is a reliable

indicator that the state of affairs obtains. In the purest form of this approach, it is at most a matter of superficial idiom that we do not attribute knowledge to properly functioning thermometers. (p. 401)

So McDowell regards the third hybrid option as more attractive than pure externalist reliabilism because it at least recognizes that we can have reasons for knowledge claims that fit into a rational structure. But just because reliability features in the pure externalist account as an external constraint, it does not do so in the hybrid conception because here it is supposed to capture the idea of a knowing subject raising questions about their policies of forming beliefs on the basis of appearances.

Two key problems with a hybrid approach

When so construed, however, internal reasons and the assessment of reliability of one's policies become divorced from the external constraint of truth. This leads to a problem. If the world actually being thus and so is not directly available to the resources of reason—if such facts are external to our reason—how can knowing subjects have the resources to assess the reliability of methods of basing beliefs on appearances? Some such facts would have to be taken for granted in order to test whether a method really was reliable, but that contradicts the assumption that having good reason is an internal matter and truth is external.

There is also a second objection to the hybrid conception. One of the points of the distinction between the concept of knowledge (as *justified* true belief) and mere true belief is that ascribing knowledge to someone is supposed to rule out the idea that they have a true belief as a mere matter of luck or accident. But if two people can achieve equally good standings in the space of reasons while only one of them is actually right about a matter, that extra external condition is merely a matter of luck or accident. In other words, the hybrid conception does not support the key distinction behind the concept of knowledge.

McDowell's alternative to the hybrid approach

McDowell suggests that the hybrid conception is often taken to be obvious. But he suggests that this stems merely from an apparent lack of other options. To counter this he suggests an alternative: one that rejects the argument from illusion as a good basis for thinking about knowledge. McDowell suggests that achieving a good standing in the space of reasons is itself dependent of the kindness of the world. This contrasts with the hybrid position, influenced by the argument from illusion. Even being justified requires a worldly contribution:

But that the world does someone the necessary favour, on a given occasion, of being the way it appears to be is not extra to the person's standing in the space of reasons. Her coming to have an epistemically satisfactory standing in the space of reasons is not what the interiorized conception would require for it to count as her own unaided achievement; but then once she has achieved such a standing, she needs no extra help from the world to count as knowing. (p. 406)

And again,

Of course we are fallible in our judgements as to the shape of reasons as we find it, or—what comes to the same thing—as to the shape of the world as we find it. That is to say that we are vulnerable to the world's playing us false; and when the world does not play us false we are indebted to it. But that is something we must simply learn to live with, rather than recoiling into the fantasy of a sphere in which our control is total. (pp. 407–408)

How can the space of reasons be 'external' to us?

There are two further clues as to McDowell's suggestion here. One is hinted at but not developed in a footnote. McDowell compares the idea of being justified with the idea of responsibility for one's actions, what one does. In this latter sphere, that of intervening in an objective world, one obviously does not have complete control over what happens. But in this case we do not think that what we actually achieve is somehow less than our final actual actions (which also require a favour from the world such as that our arms are not tied down or the door we are trying to open is not jammed shut). We are (for reasons he does not attempt to explain) less tempted in the case of action to postulate an internal realm (perhaps of pure willing) where we do have complete control over what happens. Similarly, he suggests, we should regard reason as requiring a favour from the world, while nevertheless not undermining the idea that this still comprises *our* epistemic standing.

Secondly, McDowell draws on an idea he has developed in the philosophy of thought and language and suggested in previous readings from *Mind and World*. One of the points of announcing that he would investigate having a standing in the space of reasons was to mark out the idea that thoughts stand in a network of rational relations. (This is the contrast with Mellor's (1991) externalism, which ignores this crucial feature of thought.) But what stand in these relations are thoughts about various different aspects of the world. Thoughts have (or simply are) 'contents' of the form: there is a door in front of me, that is a red patch, yellow is lighter than black. But if thought is construed as an internal contrast to the external world, if 'we set it off so radically from the objective world, we lose our right to think of moves within the space we are picturing as content-involving'. (p. 409)

So McDowell thinks that the underlying objection to the hybrid picture of our epistemological standing is more fundamental than just a question of epistemology: of the nature knowledge and justification construed narrowly. It is instead a broader or deeper problem of thinking of our mental states as world-involving at all.

In the case of beliefs based on perceptions, McDowell suggests that scepticism of the sort generated by the argument from illusion can be resisted by rejecting its grounding assumption: that in perceptual experience we always take in less than full facts and are restricted merely to appearances. In other papers, McDowell describes the philosophical account of perception that he rejects as a 'highest common factor' theory of perception. On such a

theory, what the argument from illusion reveals is the *common* factor between illusory and veridical experience: an appearance that stops short of the facts. He argues instead for a 'disjunctive conception'. In veridical experience we take in the full facts: we see that there is a door before us, or whatever. Only in illusory experience do we take in a mere appearance, something that stops short of the facts.

Induction again

McDowell goes on to extend his argument from scepticism applied to perception to scepticism applied to induction. The idea that he wishes to reject in this context (analogous to rejecting that perceptions can be characterized as taking in appearances that stop short of facts, and thus do not require worldly favours) is that an epistemic position can be characterized prior to the application of risky inductive principles. If such a position can coherently be described then the further move to form inductive inferences may appear problematic and in need of further justification. Indeed this is what Hume presupposes when he says that he will ask what evidence there is to ground inferences about what lies 'beyond the present testimony of our senses, or the record of our memory'.

However, McDowell argues that if one takes seriously such talk of the *testimony* of our senses, then what the senses deliver must involve a characterization of the world, rather than simply a thin description of, for example, a 'wash of chromatic sensation'. But if so:

there cannot be a predicament in which one is receiving testimony from one's senses but has not yet taken any inductive steps. To stay with the experience of colour... colour experience's being testimony of the senses depends on the subject's already knowing a good deal about, for instance, the effect of different sorts of illumination on colour appearances... (p. 411)

So the supposed predicament of the inductive sceptic is a fiction... Hume's formulation can seem to describe a predicament only if one does not think through the idea that its subject already has the testimony of the senses and this means that scepticism about induction can seem gripping only in combination with a straightforwardly interiorizing epistemology for perception. (p. 412).

The solution summarized

McDowell argues that it turns out that to have perceptual knowledge requires that one already has inductive knowledge. So scepticism about induction is not as localized as it might have seemed. It threatens to undermine perceptual knowledge as well. Furthermore, there is something unstable about thinking that one can talk about having reasons for beliefs if at the same time one construes those reasons as internal mental states that are cut off from the world. Coupled with the realization that a highest common factor account of perceptual experience can never account for the fact that knowledge concerns external worldly states, this helps undermine the attraction of a form of scepticism here.

In the case of perception, if one does not assume that perception stops short of taking in facts, then the question: 'how do we

know that the appearances are reliable?' will not be pressing. Perception is reliable because it is a (fallible) but direct openness to the world. Similarly, in the case of induction, if one does not assume that one could characterize a state of trusting the testimony of the senses without induction, then one will not be drawn to ask whether that extra step is justified. McDowell's point here is that there is no such middle ground and thus no alternative to accepting induction alongside the reliability of perception.

Does it work?

But does this really block inductive scepticism? Like most forms of philosophical scepticism it seems that here again, once inductive scepticism is unleashed there is no way to *refute* it. The sceptic can continue to press the following point. While it may be necessary to *think* of induction as reliable in order to characterize the deliverances of the senses as world involving, why believe that it is actually reliable? Although we may have to think that tables and chairs will continue to exist into the next instant when we report seeing that they are there now, what assurance is there that they actually will?

On the other hand, such residual scepticism is strangely unmotivated. Without the thought that induction is a further and risky move beyond the secure base of perception (or that it is a further risk to construe the output of perception as anything more than mere appearances), what grounds are there for calling that 'further move' into question? Hume explicitly contrasts induction and perception when in fact they are interconnected and interdependent. McDowell's account does not make either an infallible method of arriving at beliefs but suggests that they stand or fall together.

The social articulation of the space of reasons and evidence-based medicine

McDowell's paper 'Knowledge and the internal' (1998 pp 395–413) was originally published with a reply from his colleague at Pittsburgh Robert Brandom called 'Knowledge and the social articulation of the space of reasons'. This contained in part an account of Brandom's particular views of the nature of meaning, which he argues depends on social practices. However, whether or not one follows that line of argument, McDowell's approach to knowledge dovetails particularly easily with a social approach.

One of the ways in which one's epistemic good standing can depend on a worldly favour is exemplified in asking for directions. Under favourable circumstances this is a way of getting knowledge of, for example, where the nearest station is. But typically when we ask for directions we cannot construct an argument from what we seem to hear, or whatever might be immune to Descartes argument from illusion, to a conclusion that the station has to be where we subsequently think it is. Reading Descartes, or when we 'come over philosophical', we may conclude that if we cannot do this then we cannot *know* where the station is. However, that is not how we normally use the concept of knowledge.

Asking for directions provides a model of, for example, inquiring about treatment options: a key aspect of EBM. To acquire medical knowledge we have to take responsible steps to acquire a good standing in the space of reasons. But we do not need to be able personally to vouch for every step in the process of validating medical evidence. McDowell thus begins to sketch out a model for thinking of the balance between personal and collective responsibility in EBM. Much of what we rely on is not explicitly articulated as evidence in the narrow sense but rather rubs off on us through a medical apprenticeship. But on the other hand, this need not make it any the less a matter of knowledge.

Note that two themes inter-relate here. One is the idea that there is a tacit dimension to knowledge even in the idea of being justified. This connects back to the more general discussion of tacit knowledge in Chapter 14. The suggestion that one can acquire justified beliefs by asking other reliable sources even if one cannot establish an argument oneself to show that the belief must be true is an application of idea of a tacit dimension to medical knowledge as well as practice. The second aspect is the social dimension. Knowledge can rub off on others. Taken together this is a powerful alternative view to the traditional Cartesian focus on an individual acquiring knowledge by him- or herself and being able to establish its credentials unaided. That Cartesian picture is very intuitive but provides an unrealistic picture of real knowledge.

Wittgenstein's picture of the inherited background to knowledge claims

Wittgenstein on knowledge versus certainty

The context for understanding EBM by thinking of McDowell's sketch of the nature of empirical knowledge generally and his attempt to dissolve inductive scepticism can be complemented by considering a more general account of the 'space of reasons' in which we frame knowledge claims. The importance of an inherited belief structure and the relation between knowledge and certainty were also explored in a fragmentary way by Wittgenstein in the last months of his life, subsequently published as *On Certainty* (1979). The work comprises a chronological arrangement of remarks which, unlike his *Philosophical Investigations* (1953), were not sorted into a rational order. They were working notes rather than remarks prepared for publication. The reading comprises a representative sample of the view developed throughout the book.

EXERCISE 9

(30 minutes)

Read the extracts from

Wittgenstein, L. (1969). *On Certainty*. Oxford: Basil Blackwell

Link with Reading 16.8

- ◆ To what extent is Wittgenstein's account foundational?

Moore's attack on scepticism

Wittgenstein's remarks were prompted (possibly indirectly) by his fellow Cambridge philosopher G. E. Moore's defence of common sense realism against scepticism. Moore's argument was remarkably simple. He held one hand (his own) and claimed that he knew that it was a hand. As a hand, it was a material object. Thus he knew that there was at least one material object in the world. And thus philosophical arguments against the reality of the material world were refuted.

Wittgenstein rejects that argument by suggesting that a sceptic (or idealist as such scepticism is here characterized) will say that 'he was not dealing with the practical doubt that was being dismissed, but... a further doubt behind that one' (§19). Moore mistakenly treats sceptical doubts as though they were practical doubts requiring practical justifications. Instead, Wittgenstein suggests a different response to scepticism, which turns on the point that: 'a doubt about existence only works in a language-game. Hence that we should first have to ask: what would such a doubt be like?, and don't understand this straight off?' (§24). Thus Wittgenstein's suggestion for how to respond to scepticism is to argue that any doubts, whether practical or sceptical, requires a context to give it its meaning.

Wittgenstein on ordinary doubts

In fact the main theme of *On Certainty* is not the problem of scepticism but charting the context of *ordinary* knowledge claims and expressions of doubt. Our knowledge claims and our doubts form a *system*. Without this context they would not have any clear meaning.

All testing, all confirmation and disconfirmation of a hypothesis takes place already within a system. And this system is not a more or less arbitrary and doubtful point of departure for all our arguments: no, it belongs to the essence of what we call an argument. The system is not so much the point of departure, as the element in which arguments have their life. (§105)

Wittgenstein also calls this system a 'picture of the world' that members of a community largely share. But we do not as individuals arrive at such a world picture by satisfying ourselves of its correctness. 'No: it is the inherited background against which I distinguish between true and false.' (§94). Thus the suggestion is that in order to test or check a claim just such a background is required to provide the ground rules for empirical inquiry. Thus the background itself cannot as a whole be checked for its truth.

Five other features of this background are important:

1. It comprises a motley of different sorts of 'claims'. Moore's example that 'this!' is a hand is one. Others include the claim that the earth existed long before my birth; that my name is such and such; that I have never been to the Moon. So if these claims are taken to express a kind of foundation for empirical inquiry, it is not a traditional view of foundations within philosophy. Traditionally, these have been construed as an

homogeneous class of claims about my own mental states, experiences, or appearances in my visual field.

2. There is some difficulty in characterizing our epistemic attitudes to the 'claims' that comprise the background. They are typically *not* claims we make. Contra Moore we do not *know* that 'this is a hand' because it not the sort of thing we could doubt or provide grounds for. Neither are they beliefs or assumptions. Rather they are certainties expressed by our *actions*. Thus for example, we do not *assume* that we have feet when we stand up, but our animal certainty here is expressed in our lack of tentativeness in standing. 'In the beginning was the deed' (§402).
3. As the point above suggests, Wittgenstein separates the concepts of knowledge (and doubt etc.) from that of certainty. He argues that there is symmetry between knowledge and doubt. We can only claim to know what it would also make sense to doubt. This is because both belong to a 'language-game' or linguistic practice of asking for and giving reasons. Certainty, by contrast, characterizes the necessary background for that practice. It is the mark of what lies at the edge of empirical inquiry and is not called into question. Certainties 'lie apart from the route travelled by enquiry' (§88).
4. Just as the background comprises a motley of different sorts of matters that are taken as certainties, so too the limits of reasonable empirical inquiry is taught by example. This is another case where the rationality comprising empirical inquiry resists codification as a theory. We do not have a theory about what can and cannot be doubted. 'We do not learn the practice of making empirical judgements by learning rules. We are taught *judgements* and their connexion with other judgements. A *totality* of judgements is made plausible to us.' (§140).
5. While the systematic background or picture of the world is a prerequisite for empirical testing, elements of it can be called into question. 'The mythology may change back into a state of flux, the river-bed of thoughts may shift. But I distinguish between the movement of the waters on the river bed and the shift of the bed itself; though there is not a sharp division of the one from the other.' (§97).

One application of Wittgenstein's account of knowledge and certainty in psychopathology has been the idea that delusions share many of the features of Moore propositions (their central role in reasoning; the fact they stand fast; their imperviousness to argument etc.). (See, for example, Campbell, 2001.)

Wittgenstein on induction

More so than many of his works, the notes published in *On Certainty* are more descriptive than argumentative. Wittgenstein offers the account as a proposed description of our epistemic practices without providing a watertight case that it *must* be true. (There are some arguments: such as the argument that the giving

of reasons must terminate somewhere if inquiry is to be possible at all.) But it remains a plausible account to combine with McDowell's attempt to defuse scepticism discussed above. What, then, does Wittgenstein say about induction? Here are some (about half) of his explicit comments:

Have we in some way learnt a universal law of induction, and do we trust it here too?—But why should we have learnt one *universal* law first and not the special one straight away? (§133)

But do we not simply follow the principle that what has happened will happen again (or something like it)? What does it mean to follow this principle? Do we really introduce it into our reasoning? Or is it merely a natural law which our inferring apparently follows? This latter it may be. It is not an item in our considerations. (§145)

The squirrel does not infer by induction that it is going to need stores next winter as well. And no more do we need a law of induction to justify our actions or our predictions. (§287)

I might also put it like this: the 'law of induction' can no more be *grounded* than certain particular propositions concerning the material of experience. (§499)

Wittgenstein's suggestion is, in other words, that searching for a justification for a law of induction marks a misunderstanding of the role that induction plays in our epistemic practices. Reasoning from past experiences is taught by example in particular cases. We are not first taught a universal law as a piece of knowledge (for which reasons might thus be offered). Rather we are trained in drawing inferences from past experiences as a taken for granted, unquestioned and certain practice.

Furthermore, given practical certainty expressed in our reasoning from past experiences, nothing *could* be offered to ground the law of induction as it is as certain as anything that might be offered in support of it. 'When one says that such and such a proposition can't be proved, of course that does not mean it can't be derived from other propositions; any proposition can be derived from other ones. But they may be no more certain than it is itself.' (§1).

Endgame?

If Wittgenstein's account is correct then despite its fundamentally different aim and feel to the philosophy of science accounts in the previous session, it does reinforce a general moral. Empirical testing must take place against a background that is not simultaneously called into doubt. Of course, elements of that background can be called into question when the 'river-bed of thought' shifts. But any testing requires that some things are not called into question and are instead held certain. This view of science sometimes summarized in Otto Neurath's (1932) metaphor: a boat whose planks have to be replaced one by one even while it is afloat.

What Wittgenstein's discussion especially reinforces in Lakatos' (1970) and more especially Kuhn's (1970), Barnes' (1974), and Bloor's (1976) accounts is the practical and piecemeal character of what is held certain. This includes the way in which we learn from experience: both in the case of direct perception (to recapitulate

the subject of Chapter 12) but also, the subject of this chapter, in the case of drawing general conclusions from finite past experience. What does and does not count as good evidence is not grounded on a general context-free method—as Hume suggests in setting up the challenge to justify a general method of induction—but in particular judgements we learn to make. This general claim will also apply to the evidence appealed to in EBM, as we will see in the next session.

Reflection on the session and self-test questions

Write down your own reflections on the materials in this session drawing out any points that are particularly significant for you. Then write brief notes about the following:

1. What is the main difference between philosophy of science responses to the problem of induction and epistemological responses?
2. How does Mellor attempt to defuse the problem?
3. How does McDowell's response differ from Mellor's?
4. What does Wittgenstein add to our understanding of knowledge?

Session 4 Evidence-based medicine and clinical trials

This chapter has so far examined conflicting responses to the problem of induction within two different philosophical traditions: the methodologically orientated philosophy of science and the more analytically orientated tradition of philosophical epistemology.

One key lesson from the philosophy of science was that there is at least debate about whether evidence is deployed in science to confirm or to refute theories. But more importantly, however evidence is used, any realistic account of scientific research has to stress a top-down influence of theory on evidence as well as a bottom-up dependence. Evidence is gathered in the context of broader theoretical structures—perhaps Lakatosian research programmes or Kuhnian paradigms—and these have a profound influence on the interpretation of evidence.

The key lesson from the branch of philosophy called epistemology was that the apparent profound difficulties with induction outlined by Hume stem from a particular intuitive but nevertheless misleading understanding of knowledge. More contemporary accounts suggest that knowledge is not a matter for one individual to ensure through a process of argument from scratch, as Descartes famously assumed. Instead to have knowledge depends on in part on factors outside an individual's control. An important example of this is the acquisition of knowledge

second hand through the testimony of others. This suggests that EBM should also not be assumed to be an individualistic project.

This final session will return to the concrete issue with which the chapter began: the growth of EBM in psychiatry.

Evidence-based medicine, science, and the philosophy of science

The key aim of EBM is to optimize learning from past experience so as to draw conclusions about future treatment options. It thus aspires to the practical codification of good inductive strategies. These centre on the use of research based on clinical trials. In this session we will look at how the philosophy discussed so far in this chapter and the topics discussed in previous chapters touch on the kind of evidence available for EBM.

The purpose of this discussion is not to criticize the application of EBM approaches to psychiatry in general. This reiterates the moral of this part of the book, which has examined what is involved in being a science. In so doing it has criticized intuitive but oversimple models of science, but it has not been antiscientific. The moral has been rather different: that the form of rationality exemplified by good scientific practice cannot be codified in simple models or mechanical principles. This is not to say that there is no place for general principles such as: 'favour simple explanations'; 'do not multiply entities needlessly'; 'favour generality'; 'pay attention to the evidence'; 'do not abandon a theory too soon in the face of anomalies'; etc. But the conflicting demands of these principles need to be weighed up together. They do not amount to an abstract model of scientific method.

In a similar way, then, the purpose of this session is not to criticize the very idea of EBM but to bring out the fact that behind simple formulations of it lurk principled complexities. These are especially important in the case of the application of EBM to psychiatry.

Trials and treatments: Mill's methods

Look again at a standard EBM hierarchy of evidence taken from Geddes and Harrison (1997).

- 1a Evidence from a meta-analysis of RCTs
- 1b Evidence from at least one RCT
- 2a Evidence from at least one controlled study without randomization
- 2b Evidence from at least one other quasi-experimental study
- 3 Evidence from non-experimental descriptive studies, such as comparative studies, correlation studies and case-control studies
- 4 Evidence from expert committee reports, or opinions and/or clinical experience of respected authorities.

EBM places considerable weight on research findings arrived at through clinical trials, in particular, through *randomized* clinical trials and meta-analyses of these trials. These are at the top of the list. At the bottom is a more traditional view of medicine: respect for authority. The list suggests that meta-analysis of RCTs

provides a better way of learning from experience than consulting an authority. And this prompts a question: How do we know that the one is better than the other? More specifically why should we think that such trials are a reliable method of discovering causal connections between, for example, treatments and results.

There seem to be two possible answers to this question.

- ◆ *Either*, there is an a priori argument to show what method should be at the top of the list.
- ◆ *Or*, it is a matter of empirical discovery. But if so, how is the evidence for the issue marshalled?

To get a feel for this question we will consider a brief introduction to clinical trials set out in a typical textbook: S.J. Pocock, *Clinical Trials: a practical approach* (1983).

Clinical trials as experiments

Pocock states baldly that a clinical trial is a planned experiment (on patients with a medical condition). The further constraints that Pocock places on good clinical trials refine this thought, but the underlying connection between clinical trials and experiments remains. So a clinical trial is a particular type of experiment designed to provide evidence for some form of clinical hypothesis or theory.

The methodology of clinical trials

In a chapter of the reading called 'Controlled clinical trials and the scientific method' (p. 4) Pocock sets out a summary of the scientific rationale for a standard kind of clinical trial. These are comparative trials in which a (normally new) treatment is compared with other existing treatments in accordance to a specific and rigorous research method. (Pocock labels these 'phase III trials'.)

Pocock specifies a number of features of such trials:

1. They should be comparative. The experiences of patients on the treatment under trial are compared with a control group: the experiences of patients on other treatments (possibly including no treatment).
2. They should be randomized. This is supposed to prevent conclusions being drawn about the effects of drug treatments, for example, which are really the effects of some other uncontrolled factor present in the sample.
3. They should, wherever possible, be double blind trials. Neither the patients nor the clinicians testing results should know whether they belong to the test group or control group.

Aside from these general features, Pocock claims that clinical trials should proceed through a predetermined series of steps 'if the principles of scientific method are to be followed' (p. 5). Those steps are:

1. Define the purpose of the trial: state specific hypotheses.
2. Design the trial: a written protocol.

3. Conduct the trial: a good organization.
4. Analyse the data: descriptive statistics, tests of hypotheses.
5. Draw conclusions: publish results.

Data gathering and theory

It is significant that neither the third step: 'conduct the trial' nor the fourth step 'analyse the data' make any reference to the further complexities involved in the idea of experimental data: direct theory free reports of how the world is found to be. As Chapter 12 argued, the very making of observations and forming of observation statements is itself charged with theory.

What is the foundation for randomized-controlled trials?

But there is a different question about the claim that these are just the steps that a clinical trial should follow if it is to be scientific. Consider, for example, whether the claim is *analytic*. Does it follow from the very meaning of the phrase 'scientific method'? And if it is not, what sort of claim is it? Is it a description of scientific method as it is usually carried out, or a prescription on how it *should* be carried out? If the latter, what argument or process of reflection leads to just these steps?

What should be obvious is that these steps could not have been arrived at purely by a priori reasoning. While, for example, anyone with experience of carrying out scientific research will be able to think of advantages in having a *written* protocol for experimental design, these reasons are founded in an understanding of human memory and communication. One can imagine other circumstances in which such written record would be unnecessary or even useless. Think of a culture with a strong aural tradition who were also plagued by dyslexia. So the steps identified are plausible rules for good conduct distilled from practical experience. But it would be rash to construe them as *the* Scientific Method, where that is construed as timeless, universal, and context-free.

(Recall in Chapter 13, Bernard Williams's invocation of scientific method as a way of making sense of his Absolute Conception of the world. McDowell argued that such an attempt was bound to fail because the very idea of the scientific method was either too vacuous (if it simply meant the best way to arrive at empirical truths) or itself too parochial to ground an absolute conception of the world, a view expressive of no particular perspective. If these rules governing clinical trials are construed as the scientific method, it is a very *local* construal of that method.)

This suggests that the author has based the list of steps on experience (his and others') of successful clinical trials rather than a priori reasoning. But if so this raises an interesting question akin to Hume's problem of induction. *How* should one learn from such experience? How can experience of the results of scientific methods themselves properly inform scientific method? This leads to an obvious suggestion: one should compare different methods of gathering data to determine the most reliable method. In other

words, one should carry out a randomized control trial to determine the efficacy of just this sort of RCT. Clearly, however, this will beg a question because it will involve a prior assumption that RCTs are a reliable method of gathering data.

This line of thinking repeats an idea found in earlier readings in this chapter, especially those by Lakatos (1970), Kuhn (1970), and Wittgenstein (1969). Theory testing—in this case hypotheses about treatment efficacy or the aetiology of illnesses—can only take place against a background of other theories, hypothesis, and assumptions, which are not themselves simultaneously called into question. This background includes theories about the correct working of instruments. It involves methodological principles about how best and how often to take readings. But it also includes methodological assumptions about the efficacy of clinical trials. Now of course, this is not to say that the methodology of clinical trials in general could not itself be subject to testing. But it *is not* subject to testing in the course of routine medical trials. And testing it would require that other assumptions were held certain. This suggests that clinical trials are not so much the foundation of medical science but part of a fallible holistic and interdependent system that has no eternally fixed foundations.

We suggested earlier that there might be two approaches to the question: how should we determine what should lie at the top of the EBM evidence hierarchy? One answer is that we can find out from experience what the best of learning from experience is, which leads to a kind of circularity. The other approach is by developing a priori arguments. To see how such an argument might be begun at least consider the following intellectual tools drawn from the history of philosophy.

EXERCISE 10

(20 minutes)

Look at the statements of Mill's 'method of agreement' and 'method of difference' taken from J.S. Mill (1879) *A System of Logic*, chapter VIII, pp. 448–471:

The method of agreement

If two or more instances of the phenomenon under investigation have only one circumstance in common, the circumstance in which alone all the instances agree, is the cause (or effect) of the given phenomenon. (p. 451)

The method of difference

If an instance in which the phenomenon under investigation occurs, and an instance in which it does not occur, have every circumstance in common save one, that one occurring only in the former; the circumstances in which alone the two instances differ, is the effect, or the cause, or an indispensable part of the cause, of the phenomenon. (p. 452)

- ◆ What do they achieve?
- ◆ To your knowledge, how do they relate to clinical testing?
- ◆ What is their status and foundation (e.g. a priori or a posteriori) as methodological prescriptions?

The method of agreement and the cholera outbreak of 1854

Mill outlines five experimental methods, which he suggests underpin experimental inquiry. These are:

1. The method of agreement
2. The method of difference
3. The joint method of agreement and difference
4. The method of residues
5. The method of concomitant variations

It is worth reflecting on the applicability of these to clinical scientific research. We will pick out two. The method of agreement is stated: 'If two or more instances of the phenomenon under investigation have only one circumstance in common, the circumstance in which alone all the instances agree, is the cause (or effect) of the given phenomenon. (p. 451).

Now as a model of clinical research, the method of agreement is of limited use. One example that approximates to it was the work done by John Snow (1813–58) a member of the Royal College of Surgeons of England. After an outbreak of cholera in 1854, Snow investigated the spread of the disease by mapping its victims in London. He discovered that the key common element in each case was that the victims had drunk water from a particular well. After officials followed Snow's advice to remove the handle of the Broad Street Pump, the epidemic was contained.

But this is only an approximate application of the Method of Agreement. The method proper requires that the only condition that a variety of experimental cases have in common turns out to be the cause of the (controlled for) effect in question or, vice versa, is the effect of the cause in question. Clearly, the practical difficulties of picking a trial sample who were alike in only one respect would be enormous. But on further reflection this is really a principled problem because there is no limit to the number of (e.g. relational) properties that an individual (a human or simply a thing) has. All human subjects, for example, share the property of being in the earth's gravitational field, of being in this part of the galaxy, of being bigger than a pea, etc.

The method of difference

The method of difference, however, is a better model for clinical research. It is defined thus:

If an instance in which the phenomenon under investigation occurs, and an instance in which it does not occur, have every circumstance in common save one, that one occurring only in the former; the circumstances in which alone the two instances differ, is the effect, or the cause, or an indispensable part of the cause, of the phenomenon. (p. 452)

Mill suggests that the conditions for this method are rarely spontaneously met in nature as they require that two circumstances are exactly the same except for one feature. But they can be met experimentally. Mill suggests introducing a new phenomenon to

a situation—perhaps by physically moving a substance—yields just this sort of pair of circumstances. Mill warns, however, that the method of bringing about a change might itself constitute an important difference from the initial circumstance.

The method of difference and clinical trials

The method of difference also forms a plausible model of much clinical research. As we have seen, clinical trials are generally contrastive. The effects of a drug on one trial group are compared with the effects of no such treatment (in fact generally a placebo) on another 'control' group. If the one group is treated in the same way as the other in all respects other than the administration of the drug, then the subsequent differences are the causal result of the drug treatment.

What the methods achieve

But what do these methods achieve? Mill describes them as methods of *eliminative* induction. This contrasts with the Humean idea of enumerative induction: building up generalizations from particular cases. Mill's methods are negative: they aim at eliminating conditions as irrelevant to cause–effect relations. They resemble falsificationism. Thus the method of agreement is a method of eliminating conditions from a list of conditions that are necessary for another condition, an effect, for example. (Any conditions that are *not* present when the effect in question occurs cannot be *necessary* conditions of that effect.) The method can also be used to determine sufficient conditions. In this case, one looks for those conditions that are *never* present when an effect is absent. (If a condition is present when an effect is absent then it cannot be sufficient for that effect.) Thus one slowly eliminates conditions from a list of potential sufficient conditions.

The method of difference provides a way of eliminating conditions that are not sufficient for an effect from some overall circumstance that is sufficient, i.e. when the effect is present. (Any factors that are present when an effect in question is not present cannot themselves be sufficient for that effect.) Now there may be many different sufficient conditions, but the method of difference is a way of paring down some complex combination, which is itself sufficient to find out which parts of it are not sufficient.

What the methods do not achieve

What should be clear is that these methods do not engage with Hume's problem. Mill assumes that there are law-like relations between conditions in nature such that some conditions are necessary and some sufficient for others. This assumption presupposes a positive response to Humean scepticism. But after the work of the last session, that is reasonable to make.

Mill's methods provide models for determining causal connections, models that can be arrived at *a priori*. But they are, however, overly simplified models of empirical research. They cannot provide an *a priori* argument for the EBM evidence hierarchy.

Think again of the criticism of the method of agreement above. It is impossible to ensure that different complex circumstances have nothing (except a condition under scrutiny) in common. Likewise in the case of the method of difference it is impossible to ensure that two different circumstances are identical in all but one respect. (If they are different, are they in different places, for example?) However, it is possible to ensure that nothing *causally relevant* is the same in the former case or different in the latter case. This suggests that Mill's methods cannot play a foundational role, a role that presupposes no other causal knowledge. Rather, in accordance with the argument drawn from the philosophy of science in the second session, the role of Mill's methods lies within a background of other fallible prior scientific beliefs about what causes what. They are a further constraint on our causal knowledge.

A further way of bringing this out is to consider Mill's use of letters to label different conditions (antecedent conditions A, B, C, etc. and succeeding conditions a, b, c, etc.). One of his critics, the philosopher of science William Whewell, wrote: 'Upon these methods, the obvious thing to remark is, that they take for granted the very thing which is most difficult to discover, the reduction of the phenomena to formulae such as are here presented to us.' (*Of Induction, With Especial Reference to Mr J Stuart Mill's System of Logic*, 1849, p. 44).

Mill simply assumes that we already possess, when searching for causal factors, a complete description that captures all the causally relevant factors. Given this, then elimination of irrelevant factors is comparatively easy. But from where did this vocabulary arise? As the readings in the second session above emphasized, it is a matter of much greater complexity than simply reading off inductive generalizations from the world.

Both of these points can be translated in to the context of mental health to raise yet further practical difficulties. Establishing conditions as the same in all but one respect (in the Method of Agreement) or different in all but one (Method of Difference) is difficult enough in the case of physical medicine, but establishing these cases for psychiatric conditions is even more difficult. To mention just one reason: psychiatric symptoms are rarely present in isolation but are combined in ways that are not exactly repeated between different patients. Weakening the requirement to that of causal relevance (as mentioned above) does not help. Likewise, the prospect of reducing a description of the state of individuals' mental lives to the letters illustratively used by Mill is grossly artificial.

The status of the evidence-based medicine evidence hierarchy

While some arguments can be given for the efficacy of clinical trials, experimental methods are themselves subject to scientific reasoning and experiment. They are thus subject to the same boot-strapping that has been an implicit theme of this part.

A proper understanding of scientific research and a better model of knowledge helps put the nature of EBM in context. We will

finish this chapter with some reflections on the connections between EBM and other matters in this part and the next.

Evidence-based medicine and the theory dependence of observation

First, there is the application of the theory dependence of observation to the gathering of data in clinical trials. Chapter 12 discussed arguments to the effect that no principled separation of theoretical and observational claims was possible. Observation is always theoretically charged. In the context of this session this can be related to Whewell's (1849) criticism of Mill's methods. A clinical trial is only possible against a background of assumptions of what is and is not causally relevant, of what the best descriptions of effects are and so on. Not every feature of a trial can be attended to, controlled for or even described. And what is described, is described according to prevailing theory. Thus, to repeat a claim made earlier in this chapter, trials cannot be seen as methods of simply harvesting data for subsequent interpretation but instead have their role in a context already charged with theory.

Evidence-based medicine and natural classifications

Given that no principled separation of theory and observation is possible, there is no hope for thinking of descriptive psychotherapy as a purely descriptive enterprise that can escape assumptions about the underlying mental structures. Some such assumptions must form at the very least an implicit element of psychiatric classification and thus be presupposed in those clinical trails that pertain to mental health. This in turn raises a further question about the nature of evidence here. As we saw in Chapter 13, some clinicians and philosophers have argued both that there is an evaluative element to psychiatric diagnosis, which is found in neither physical science descriptions nor physical medicine and, further, that this implies that such classifications are not natural. They do not carve nature at its joints. As we saw in Chapter 13, there are arguments against the second, metaphysical, claim. But on the other hand, there is reason to be pessimistic in the assessment that classifications can be said to be natural.

Evidence-based medicine and tacit knowledge

This theme was further reinforced by consideration of psychiatric diagnosis in Chapter 14. Discussion of the nature of both explanation and replication in the physical sciences suggested an important role for tacit knowledge. It seems plausible, however, that a science of the mind must contain further elements of tacit or implicit knowledge. There are arguments that in this case diagnosis involves an overall judgement of a client's state of mind, which resists breaking down into component symptoms. These can only be identified in the context of the overall judgement. If this is so, then diagnosis is uncodifiable. It suggests in turn that the evidence available in EBM will turn on an important element of clinical judgement even though this is not explicit in bald reports of clinical data.

Evidence-based medicine, reasons and causes or laws

Why this might be so was explored further in Chapter 15. In addition to tracking down the causal aetiology of mental ill-health, psychiatry is also concerned to track meaningful relations among the causes and symptoms of conditions. But given arguments that the logical space of reasons and causes or, less question-beggingly, laws are not isomorphic, this suggests that the two aims are complementary rather different aspects of an underlying science.

Conclusions

The previous chapters have suggested that despite the attractions of a certain simple picture of science, which was further developed and given philosophical credibility by the Logical Empiricists, in reality, scientific method is necessarily more complex and resists simple codification as a recipe. In the context of EBM, the same general moral applies. There is no simple account of how evidence can be marshalled. But an awareness of the necessary complexities provides a vaccination against the dangers of assuming that there is.

In fairness, in the reading at the start of this chapter (linked with Exercise 1), Geddes and Harrison (1997) emphasize some of the complexities in the process of applying the results of clinical trials in practical clinical contexts. Thus although they argue that, ideally, treatment options should be based on evidence of effectiveness from a source as high up the list of favoured forms of evidence as possible, they also point out that the evidence of clinical trials has also to be *relevant* to the case at hand. So a further assessment has to be made of whether the subjects involved in the trial are a close enough match for the actual patient. Subjects in trials may be all of the same sex and they may be less likely to have only a single medical condition. However, what should also be pointed out is that complexities do not simply affect the practical application of such clinical findings, but also the scientific process of arriving at findings in the first place.

Thus a proper understanding of EBM plays up rather than plays down the need for critical judgement in assessing evidence. Here as throughout Part III there is a key role for the subjects engaged in scientific psychiatry to exercise skilled judgement in a way that is only ever partly codifiable.

To evidence-based medicine add values-based medicine

EBM, then, properly understood, is not a practical counterpart of the philosophers' logical positivism. It is an attempt to respond positively to the ever (and exponentially) expanding knowledge base of decision-making in health care by a reflective and methodologically formalized, approach to distilling the best available explicit evidence from the research literature.

EBM has of course been widely attacked by practitioners. Such attacks, to the extent that EBM claims, or has claimed for it by others, hegemony in health-care decision-making, are fully justified. But that there are problems with current EBM methodologies would be a poor reason to abandon the whole thing—no more should the 'problems' with (all) current scientific theories, indeed

with the scientific method itself, suggest that we abandon science! Future EBMs like future sciences, informed by conceptual, methodological, statistical and other advances, and by growing experience, will be more sophisticated.

An even worse reason for abandoning the whole thing (EBM) would be that EBM itself, to the extent that it is concerned only with explicit knowledge, is an *incomplete* response to the growing complexities of the knowledge base of decision-making in health care. Our understanding of implicit knowledge, of the 'craft' knowledge of those with practical experience, is less well developed even than our understanding of explicit knowledge. However, developments in 'narrative-based-practice' (Greenhall & Hurwitz 1998), for example, and the potential resources of a wide range of empirical (qualitative), philosophical (e.g. phenomenological, hermeneutic) and literary (e.g. thematic analysis) methods (Fulford *et al.*, 2002), are all available to supplement, not supplant, the resources of EBM.

The worse reason of all for abandoning EBM, would be that, even to the extent that it combines implicit with explicit knowledge, it misses out altogether the values base (as distinct from the evidence base) of health-care decision-making. Value-judgements, as R.M. Hare's account of their meanings so translucently shows (Part I, Chapter xx) are made on the basis of criteria that are *factual* (or descriptive) in nature. Value-judgements have an action-guiding (or 'prescriptive', in Hare's (1952) account) element as well, of course. This is why decision-making in health care, as in any other area, is always values based as well as evidence based. Evidence alone is not enough to determine decisions. Decisions are determined by evidence weighted by values: in deciding between treatments, for example, we weigh the (*desired*) effects against (*unwanted*) side-effects, and both against *cost*; and in any real life scenario, the different elements of the decision, factual, evaluative and other, will often be intimately woven together in an (apparently) seamless tapestry (Fulford, 1989, chapter 12). However, the point is that, to the extent that a decision *is* values-based, the factual criteria for the value-judgements in question should reflect, like any other factual question in medicine, the best available evidence—evidence, then, which in the model sketched here, is derived in part from craft knowledge, but also, when available, from the resources of EBM.

There is no clearer statement of the interdependence of explicit evidence, implicit (or craft) knowledge, and values, than in the introduction to one of the classics of EBM itself, David Sackett and colleagues' (2000) highly regarded 'Evidence-Based Medicine: How to Practice and Teach EBM'. In their introduction (p. 1), Sackett *et al.* define EBM as follows:

Evidence based medicine is the integration of best research evidence with clinical expertise and patient values.

By *best research evidence* we mean clinically relevant research... New evidence from clinical research and treatments both invalidates previously accepted diagnostic tests and treatments and replaces them with new ones that are more powerful, more accurate, more efficacious and safer.

By *clinical expertise* we mean the ability to use our clinical skills and past experience to rapidly identify each patient's unique health state and diagnosis, their individual risks and benefits of potential interventions, and their personal values and expectations

By *patient values* we mean the unique preferences, concerns and expectations each patient brings to a clinical encounter and which must be integrated into clinical decisions if they are to serve the patient.

Much of their book, reflecting the focus of EBM itself, is concerned with the first of these, research evidence. But this focus only underlines the need for equivalently sophisticated approaches to the roles of clinical expertise (implicit knowledge) and of values in health-care decision-making. For it is only when 'these 3 elements are integrated', Sackett and colleagues continue, that 'clinicians and patients form a diagnostic and therapeutic alliance which optimizes clinical outcomes and quality of life.' We have considered some aspects of both explicit and implicit knowledge in this part of the book. It is to values that we now turn in Part IV.

Reflection on the session and self-test questions

Write down your own reflections on the materials in this session drawing out any points that are particularly significant for you. Then, thinking back over the chapter as a whole, write brief notes about the following:

1. How does the broader discussion of responses to induction impact on EBM?
2. What is the connection between Mill's methods and EBM?
3. How should the EBM hierarchy itself be assessed?

Reading guide

A thorough practical introduction to EBM is provided by Sackett (2000) *Evidence based medicine: how to practice and teach EBM*. Specifically in relation to psychiatry, see Geddes and Harrison (1997) and Geddes and Carney (2001).

Hume and the problem of induction

- Hume's philosophy is described in a number of introductions, including David Pears (1990) *Hume's System*, and Barry Stroud (1977) *Hume*.
- The problem of induction is thoroughly set out and a particular solution suggested in Howson (2000) *Hume's Problem: induction and the justification of belief*.

Philosophy of science response to the problem of induction

- ◆ The following all contain accessible introductions to falsificationism in its Popperian and Lakatosian versions and Kuhn's philosophy of science: Bird (1998) *The Philosophy of Science* (chapters 5 and 8); Chalmers (1999) *What is this thing called science?* (chapters 4–8); and Ladyman (2002) *Understanding philosophy of science* (chapters 3 and 4).
- ◆ Book length introductions include: Corvi (1997) *An Introduction to the Thought of Karl Popper*; Larvor (1998) *Lakatos: an introduction*; Bird (2001) *Thomas Kuhn*.
- ◆ The sociology of the natural sciences is further explored in: Barnes (1974) *Scientific Knowledge and Sociological Theory*; Bloor (1991) *Knowledge and Social Imagery*; Collins (1985) *Changing Order*; Latour and Woolgar (1992) *Laboratory Life: construction of scientific facts*; Latour (1987) *Science in Action: how to follow scientists and engineers through society*.

Epistemology

- ◆ Two good introductions to epistemology are: Williams (2001) *Problems of Knowledge: a critical introduction to epistemology*; Dancy (1985) *An Introduction to Contemporary Epistemology*.
- ◆ McDowell's account of knowledge is discussed in Thornton (2004) *John McDowell* (chapter 5).
- ◆ The best of several discussions of Wittgenstein's *On Certainty* is: McGinn (1989) *Sense and Certainty*; Geddes and Carney (2001) 'Recent advances in evidence-based psychiatry'; Geddes and Harrison (1997) 'Evidence-based psychiatry: closing the gap between research and practice.'

References

- Barnes, B. (1974). *Scientific Knowledge and Sociological Theory*. London: Routledge.
- Bird, A. (1998). *The Philosophy of Science*. London: Routledge.
- Bird, A. (2001). *Thomas Kuhn*. Chesham: Acumen.
- Bloor, D. (1976). *Knowledge and Social Imagery*. London: Routledge.
- Bloor, D. (1991). *Knowledge and Social Imagery*. Chicago: Chicago University Press.
- Bloor, D. (1997). *Wittgenstein, Rules and Institutions*. London: Routledge.
- Brandom, R. (1998). Knowledge and the social articulation of the space of reasons. *Philosophy and Phenomenological Research* 55: 895–908.
- Campbell, J. (2001). Rationality, meaning, and the analysis of delusion. *Philosophy, Psychiatry, & Psychology*, 8/2 (3): 89–100.

- Chalmers, A. (1999). *What is This Thing Called Science?* Buckingham: Open University Press.
- Collins, H.M. (1985). *Changing Order*. London: Sage.
- Corvi, R. (1997). *An Introduction to the Thought of Karl Popper*. London: Routledge.
- Dancy, J. (1985). *An Introduction to Contemporary Epistemology*. Oxford: Blackwell.
- Fulford *et al.* in "Many Voices: Human Values in Healthcare Ethics" pp 1–19.
- Fulford, K.W.M., Dickenson, D.L., Murray, P.H. (eds) in (2002). *Healthcare Ethics and Human Values: An introductory Text with Readings and Case Studies*. Oxford: Blackwell.
- Fulford, K.W.M. (1989, paperback 1995). *Moral Theory and Medical Practice*. Cambridge: Cambridge University Press.
- Geddes, J.R. and Carney, S.M. (2001). Recent advances in evidence-based psychiatry. *Canadian Journal of Psychiatry*, 46(5): 403–406.
- Geddes, J.R. and Harrison, P.J. (1997). Evidence-based psychiatry: closing the gap between research and practice. *British Journal of Psychiatry*, 171: 220–225.
- Greenhalgh, T. and Hurwitz, B. (1998). *Narrative Based Medicine: Dialogue and Discourse in Clinical Practice*. London: BMJ Books.
- Hare, R.M. (1952). *The Language of Morals*. Oxford: Oxford University Press.
- Howson, C. (2000). *Hume's Problem: induction and the justification of belief*. Oxford: Clarendon Press.
- Hume, D. (1975). *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. Oxford: Oxford University Press.
- Kuhn, T.S. (1962). *The Structure of Scientific Revolutions*. Chicago: Chicago University Press.
- Kuhn, T.S. (1970). Logic of discovery or psychology of research? In *Criticism and the Growth of Knowledge* (ed. I. Lakatos and A. Musgrave). Cambridge: Cambridge University Press, pp. 1–93.
- Ladyman, J. (2002). *Understanding Philosophy of Science*. London: Routledge.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In *Criticism and the Growth of Knowledge* (ed. I. Lakatos and A. Musgrave). Cambridge: Cambridge University Press, pp. 91–195.
- Larvor, B. (1998). *Lakatos: an introduction* London: Routledge.
- Latour, B. (1987). *Science in Action: how to follow scientists and engineers through society* Cambridge, MA: Harvard University Press.
- Latour, B. and Woolgar, S. (1992). *Laboratory Life: construction of scientific facts*. Princeton: Princeton University Press.



McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.

McDowell, J. (1998). *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.

McGinn, M. (1989). *Sense and Certainty*. Oxford: Blackwell.

Mellor, D.H. (1991). The warrant of induction. In *Matters of Metaphysics*. London: Routledge.

Mill, J.S. (1879). *A System of Logic*. London: Longman.

Neurath, O. (1932). Protokollsätze. *Erkenntnis*, 3: 204–214.

Pears, D. (1990). *Hume's System*. Oxford: Oxford University Press.

Pocock, S.J. (1983). *Clinical Trials: a practical approach*. Chichester: Wiley.

Popper, K. (1972). Conjectural knowledge: my solution to the problem of induction. *Objective Knowledge*. Oxford, pp. 1–31.

Sackett, D.L., Straus, S.E., Scott Richardson, W., Rosenberg, W., and Haynes, R.B. (2000). *Evidence-Based Medicine: how to*

practice and teach EBM, (2nd edn). London: Churchill Livingstone.

Stroud, B. (1977). *Hume*. London: Routledge and Kegan Paul.

Thornton, T. (2004). *John McDowell*. Chesham: Acumen.

Whewell, W. (1849). *Of Induction, With Especial Reference to Mr J Stuart Mill's System of Logic*. London: Parker.

Wilkes, K.V. (1988). *Real People: personal identity without thought experiments*. Oxford: Clarendon Press.

Williams, M. (2001). *Problems of Knowledge: a critical introduction to epistemology*. Oxford: Oxford University Press.

Williams, D.D.R and Garner, J. (2002). The case against 'the evidence': a different perspective in evidence-based medicine. *British Journal of Psychiatry*, 180: 8–12.

Wittgenstein, L. (1953). *Philosophical Investigations*. Oxford: Basil Blackwell.

Wittgenstein, L. (1979). *On Certainty*. Oxford: Basil Blackwell.



Conclusions to Part III

One central lesson of Part V as a whole has been this. Although scientific practice is, rightly, taken to be a paradigmatic form of rationality, it resists codification in any simple overarching and context-free account of scientific method. Such codifications (starting with a traditional model of scientific progress in Chapter 11, the more precise formulations offered by the Logical Empiricists introduced in Chapter 12, running through realist prescriptions for classification in Chapter 13, models of explanation in Chapter 14 and of causal explanation in Chapter 15, and to more recent accounts of the methodology of science in Chapter 16) are all useful in different ways for drawing attention to central aspects of scientific practice. Gathering data through observation, devising reliable and hopefully valid classifications, framing explanations that invoke natural laws, positing underlying causal entities, and testing theories against the data and against one another, are again all key aspects of scientific rationality.

Time and again, however, the traditional relatively simple and intuitive models of these different aspects of scientific rationality, have been shown to be inadequate. The traditional models have been shown to be inadequate *in principle* in all areas of science. However, they have been shown to be inadequate also *in practice* in those areas of science, like theoretical physics and psychiatry, in which the problems with which the science in question is concerned are as much conceptual as empirical in nature. Complex sciences such as these require more complex models.

To take two examples of this increased complexity: there is something powerful about the idea (1) that explanation involves a kind of logical argument for the explanandum (see Chapter 14), and (2) that the mark of scientific theory is that it is falsifiable (see Chapter 16). But while both ideas are most attractive when put at their simplest neither is completely plausible until refined and made more complex and less simple. That does not imply that there is nothing to the idea of comparing scientific explanations to deductive arguments or thinking of refutation as a central virtue of a scientific (by contrast with pseudo-scientific) claim. However, it does suggest that the underlying simple idea can only be an approximation of actual science and thus has to be considered to be one of many competing scientific 'virtues' rather than an element in an algorithmic method for generating uniquely scientific discoveries about the way the world is.

A further, and related, lesson from Part V as a whole is the irreducible role of individual judgement and science. Thus, the attempt to provide a logical algorithmic recipe for scientific progress is motivated in part by the (entirely proper) attempt to

minimize bias. Science aims to avoid any dependence on a 'cult of personality', an aim that is well exemplified by EMB's emphasis, as in its hierarchy of knowledge, on controlled trials rather than expert judgement or individual experience (Chapter 16). However, while replication, reliability, and so on are (again rightly) central aspects of science's claim to objectivity, validity, and truth, nevertheless the skills and abilities and, centrally, the individual good judgement of practicing scientists, clinicians, and others, refuse to be eliminated from a realistic account of how science actually works in practice.

Good judgement plays a central role in the practical aspects of scientific know-how where, indeed, one would expect craft skills and tacit knowledge to be at work. However, good judgement is also, and equally irreducibly, at work in epistemic contexts: in deciding, for example, whether an observation is sufficiently independent of a theory to count as supporting it (as in Chapter 12); in choosing between classifications for use in different contexts and different purposes (as in Chapter 13); in formulating an explanation in terms of law-like regularities (as in Chapter 14); in postulating underlying causal entities (as in Chapter 15); and, above all perhaps, when it comes to what counts as progress, in knowing (as in Chapter 16) whether to move on an established theory, i.e. because the research programme to which it is attached is 'degenerating', or to stick with the theory in question in the face of accumulating evidence against it.

These two lessons from Part V as a whole, the lesson of 'added complexity' and the lesson of 'irreducible judgement', it is important repeatedly to remind ourselves, have been drawn from a philosophy of science that has been focused primarily, not on the human sciences, such as sociology and psychology, but on the natural sciences, notably physics. The human sciences, the considerations of Part II suggest, add to the first of these two lessons, the lesson of added complexity, a whole extra layer of difficulty characterized, variously in such terms as meaning, reasons, empathy, and understanding. The specifically *medical* human sciences, the consideration of Part I suggest, add to the second of the two lessons, the lesson of irreducible judgement, a whole further layer of extra difficulty characterized, as in Part I, in terms of values.

Debate, as we have seen in this part, continues about the role of values in the natural sciences. In the next part of the book we explore the role of values in the specifically *clinical* judgements; these are an irreducible element of the model of science that is needed when we add to the complexities of a natural science at the cutting edge, like theoretical physics, the further complexities of the science that, we have argued in this part, is at the cutting edge of the human medical sciences, psychiatry.



OUP Copyright



OUP Copyright

