

PART I
AREAS OF PHILOSOPHY

1

Epistemology

A. C. GRAYLING

For most of the modern period of philosophy, from Descartes to the present, epistemology has been the central philosophical discipline. It raises questions about the scope and limits of knowledge, its sources and justification, and it deals with sceptical arguments concerning our claims to knowledge and justified belief. This chapter firstly considers difficulties facing attempts to define knowledge and, secondly, explores influential responses to the challenge of scepticism. Epistemology is closely related to METAPHYSICS (chapter 2), which is the philosophical account of what kinds of entities there are. Epistemological questions are also crucial to most of the other areas of philosophy examined in this volume, from ETHICS (chapter 6) to PHILOSOPHY OF SCIENCE (chapter 9) and PHILOSOPHY OF MATHEMATICS (chapter 11) to PHILOSOPHY OF HISTORY (chapter 14). Chapters on individuals or groups of philosophers from DESCARTES (see chapter 26) to KANT (chapter 32) discuss classical epistemology, while several chapters about more recent philosophers also follow epistemological themes.

Introduction

Epistemology, which is also called the theory of knowledge, is the branch of philosophy concerned with enquiry into the nature, sources and validity of knowledge. Among the chief questions it attempts to answer are: What is knowledge? How do we get it? Can our means of getting it be defended against sceptical challenge? These questions are implicitly as old as philosophy, although their first explicit treatment is to be found in PLATO (c.427–347 BC) (see chapter 23), in particular in his *Theaetetus*. But it is primarily in the modern era, from the seventeenth century onwards – as a result of the work of DESCARTES (1596–1650) (chapter 26) and LOCKE (1632–1704) (chapter 29) in association with the rise of modern science – that epistemology has occupied centre-stage in philosophy.

One obvious step towards answering epistemology's first question is to attempt a definition. The standard preliminary definition has it that knowledge is *justified true belief*. This definition looks plausible because, at the very least, it seems that to know something one must believe it, that the belief must be true, and that one's reason for

believing it must be satisfactory in the light of some criterion – for one could not be said to know something if one’s reasons for believing it were arbitrary or haphazard. So each of the three parts of the definition appears to express a necessary condition for knowledge, and the claim is that, taken together, they are sufficient.

There are, however, serious difficulties with this idea, particularly about the nature of the justification required for true belief to amount to knowledge. Competing proposals have been offered to meet the difficulties, either by adding further conditions or by finding a better statement of the definition as it stands. The first part of the following discussion considers these proposals.

In parallel with the debate about how to define knowledge is another about how knowledge is acquired. In the history of epistemology there have been two chief schools of thought about what constitutes the chief means to knowledge. One is the ‘rationalist’ school (see chapters 26 and 27), which holds that reason plays this role. The other is the ‘empiricist’ (see chapters 29, 30 and 31), which holds that it is experience, principally the use of the senses aided when necessary by scientific instruments, which does so.

The paradigm of knowledge for rationalists is mathematics and logic, where necessary truths are arrived at by intuition and rational inference. Questions about the nature of reason, the justification of inference, and the nature of truth, especially necessary truth, accordingly press to be answered.

The empiricists’ paradigm is natural science, where observation and experiment are crucial to enquiry. The history of science in the modern era lends support to empiricism’s case; but precisely for that reason philosophical questions about perception, observation, evidence and experiment have acquired great importance.

But for both traditions in epistemology the central concern is whether we can trust the routes to knowledge they respectively nominate. Sceptical arguments suggest that we cannot simply assume them to be trustworthy; indeed, they suggest that work is required to show that they are. The effort to respond to scepticism therefore provides a sharp way of understanding what is crucial in epistemology. Section 2 below is accordingly concerned with an analysis of scepticism and some responses to it.

There are other debates in epistemology about, among other things, memory, judgement, introspection, reasoning, the ‘*a priori*–*a posteriori*’ distinction, scientific method and the methodological differences, if any, between the natural and the social sciences; however, the questions considered here are basic to them all.

1 Knowledge

1.1 *Defining knowledge*

There are different ways in which one might be said to have knowledge. One can know people or places, in the sense of being acquainted with them. That is what is meant when one says, ‘My father knew Lloyd George’. One can know how to do something, in the sense of having an ability or skill. That is what is meant when one says, ‘I know how to play chess’. And one can know that something is the case, as when one says, ‘I know that Everest is the highest mountain’. This last is sometimes called ‘propositional knowledge’, and it is the kind epistemologists most wish to understand.

The definition of knowledge already mentioned – knowledge as justified true belief – is intended to be an analysis of knowledge in the propositional sense. The definition is arrived at by asking what conditions have to be satisfied if we are correctly to describe someone as knowing something. In giving the definition we state what we hope are the necessary and sufficient conditions for the truth of the claim ‘S knows that p ’, where ‘S’ is the epistemic subject – the putative knower – and ‘ p ’ a proposition.

The definition carries an air of plausibility, at least as applied to empirical knowledge, because it seems to meet the minimum we can be expected to need from so consequential a concept. It seems right to expect that if S knows that p , then p must at least be true. It seems right to expect that S must not merely wonder whether or hope that p is the case, but must have a positive epistemic attitude to it: S must believe that it is true. And if S believes some true proposition while having no grounds, or incorrect grounds, or merely arbitrary or fanciful grounds, for doing so, we would not say that S knows p ; which means that S must have grounds for believing p which in some sense properly justify doing so.

Of these proposed conditions for knowledge, it is the third that gives most trouble. The reason is simply illustrated by counter-examples. These take the form of cases in which S believes a true proposition for what are in fact the wrong reasons, although they are from his or her own point of view persuasive. For instance, suppose S has two friends, T and U. The latter is travelling abroad, but S has no idea where. As for T, S saw him buying and thereafter driving about in a Rolls Royce, and therefore believes that he owns one. Now, from any proposition p one can validly infer the disjunction ‘ p or q ’. So S has grounds for believing ‘T owns a Rolls or U is in Paris’, even though, *ex hypothesi*, he has no idea of U’s location. But suppose T in fact does not own the Rolls – he bought it for someone else, on whose behalf he also drives it. And further suppose that U is indeed, by chance, in Paris. Then S believes, with justification, a true proposition: but we should not want to call his belief knowledge.

Examples like this are strained, but they do their work; they show that more needs to be said about justification before we can claim to have an adequate account of knowledge.

1.2 Justification

A preliminary question concerns whether having justification for believing some p entails p ’s truth, for, if so, counter-examples of the kind just mentioned get no purchase and we need not seek ways of blocking them. There is indeed a view, called ‘infallibilism’, which offers just such a resource. It states that if it is true that S knows p , then S cannot be mistaken in believing p , and therefore his justification for believing p guarantees its truth. The claim is, in short, that one cannot be justified in believing a false proposition.

This view is rejected by ‘fallibilists’, who claim that one can indeed have justification for believing some p although it is false. Their counter to infallibilism turns on identifying a mistake in its supporting argument. The mistake is that whereas the truth of ‘S knows that p ’ indeed rules out the possibility that S is in error, this is far from saying that S is so placed that he cannot possibly be wrong about p . It is right to say: (1) ‘it is impossible for S to be wrong about p if he knows p ’, but it is not invariably right to say

(2) 'if S knows p , then *it is impossible* for him to be wrong about p '. The mistake turns on thinking that the correct wide scope reading (1) of 'it is impossible' licenses the narrow scope reading (2) which constitutes infallibilism.

An infallibilist account makes the definition of knowledge look simple: S knows p if his belief in it is infallibly justified. But this definition renders the notion of knowledge too restrictive, for it says that S can justifiably believe p only when the possibility of p 's falsity is excluded. Yet it appears to be a commonplace of epistemic experience that one can have the very best evidence for believing something and yet be wrong (as the account of scepticism given below is at pains to show), which is to say that fallibilism seems the only account of justification adequate to the facts of epistemic life. We need therefore to see whether fallibilist theories of justification can give us an adequate account of knowledge.

The problem for fallibilist accounts is precisely the one illustrated by the Rolls Royce example above, and others similar to it (so-called 'Gettier examples', introduced in Gettier 1963), namely, that one's justification for believing p does not connect with the truth of p in the right way, and perhaps not at all. What is required is an account that will suitably connect S's justification both with his belief that p and with p 's truth.

What is needed is a clear picture of 'justified belief'. If one can identify what justifies a belief, one has gone all or most of the way to saying what justification is; and en route one will have displayed the right connection between justification, on the one hand, and belief and truth on the other. In this connection there are several standard species of theory.

Foundationalism

One class of theories of justification employs the metaphor of an edifice. Most of our ordinary beliefs require support from others; we justify a given belief by appealing to another or others on which it rests. But if the chain of justifying beliefs were to regress without terminating in a belief that is in some way independently secure, thereby providing a foundation for the others, we would seem to lack justification for any belief in the chain. It appears necessary therefore that there should be beliefs which do not need justification, or which are in some way self-justifying, to serve as an epistemic underpinning.

On this view a justified belief is one which either is, or is supported by, a foundational belief. The next steps therefore are to make clear the notion of a 'foundation' and to explain how foundational beliefs 'support' non-foundational ones. Some way of understanding foundationalism without reliance on constructional metaphors is needed.

It is not enough barely to state that a foundational belief is a belief that requires no justification, for there must be a reason why this is the case. What makes a belief independent or self-standing in the required way? It is standardly claimed that such beliefs justify themselves, or are self-evident, or are indefeasible or incorrigible. These are not the same things. A belief might be self-justifying without being self-evident (it might take hard work to see that it justifies itself). Indefeasibility means that no further evidence or other, competing, beliefs, can render a given belief insecure. Yet this is a property that the belief might have independently of whether or not it is self-justifying. And so on. But what these characterizations are intended to convey is the idea that a certain immunity from doubt, error or revision attaches to the beliefs in question.

It might even be unnecessary or mistaken to think that it is *belief* that provides the foundations for the edifice of knowledge: some other state might do so. Perceptual states have been offered as candidates, because they appear to be suitably incorrigible – if one seems to see a red patch, say, then one cannot be wrong that one seems to see a red patch. And it appears plausible to say that one's belief that p needs no further justification or foundation than that things appear to one as p describes them to be.

These suggestions bristle with difficulties. Examples of self-evident or self-justifying beliefs tend to be drawn from logic and mathematics – they are of the ' x is x ' or 'one plus one equals two' variety, which critics are quick to point out give little help in grounding contingent beliefs. Perceptual states likewise turn out to be unlikely candidates for foundations, on the grounds that perception involves the application of beliefs which themselves stand in need of justification – among them beliefs about the nature of things and the laws they obey. What is most robustly contested is the 'myth of the given', the idea that there are firm, primitive and original data which experience supplies to our minds, antecedent to and untainted by judgement, furnishing the wherewithal to secure the rest of our beliefs.

There is a difficulty also about how justification is transmitted from foundational beliefs to dependent beliefs. It is too strong a claim to say that the latter are deducible from them. Most if not all contingent beliefs are not entailed by the beliefs that support them; the evidence I have that I am now sitting at my desk is about as strong as empirical evidence can be, yet given the standard sceptical considerations (such as, for example, the possibility that I am now dreaming) it does not entail that I am sitting here.

If the relation is not a deductive one, what is it? Other candidate relations – inductive or criterial – are by their nature defeasible, and therefore, unless somehow supplemented, insufficient to the task of transmitting justification from the foundations to other beliefs. The supplementation would have to consist of guarantees that the circumstances that defeat non-deductive justification do not in fact obtain. But if such guarantees – understood, to avoid circularity, as not being part of the putative foundations themselves – were available to protect non-deductive grounds, then appeal to a notion of foundations looks simply otiose.

1.3 Coherence

Dissatisfaction with foundationalism has led some epistemologists to prefer saying that a belief is justified if it coheres with those in an already accepted set. The immediate task is to specify what coherence is, and to find a way of dealing in a non-circular way with the problem of how the already accepted beliefs came to be so.

Hard on the heels of this task comes a number of questions. Is coherence a negative criterion (that is, a belief lacks justification if it fails to cohere with the set) or a positive one (that is, a belief is justified when it coheres with the set)? And is it to be understood strongly (by which coherence is sufficient for justification) or weakly (by which coherence is one among other justifying features)?

The concept of coherence has its theoretical basis in the notion of a system, understood as a set whose elements stand in mutual relations of both consistency and (some kind of) interdependence. Consistency is of course a minimum requirement, and goes without saying. Dependence is more difficult to specify suitably. It would be far too

strong – for it would give rise to assertive redundancy – to require that dependence means mutual entailment among beliefs (this is what some have required, citing geometry as the closest example). A more diffuse notion has it that a set of beliefs is coherent if any one of them follows from all the rest, and if no subset of them is logically independent of the remainder. But this is vague, and anyway seems to require that the set be known to be complete before one can judge whether a given belief coheres with it.

A remedy might be to say that a belief coheres with an antecedent set if it can be inferred from it, or from some significant subset within it, as being the best explanation in the case. To this someone might object that not all justifications take the form of explanations. An alternative might be to say that a belief is justified if it survives comparison with competitors for acceptance among the antecedent set. But here an objector might ask how this can be sufficient, since by itself this does not show why the belief merits acceptance over equally cohering rivals. Indeed, any theory of justification has to ensure as much for candidate beliefs, so there is nothing about the proposal that distinctively supports the coherence theory. And these thoughts leave unexamined the question of the ‘antecedent set’ and its justification, which cannot be a matter of coherence, for with what is it to cohere in its turn?

1.4 *Internalism and externalism*

Both the foundationalist and coherence theories are sometimes described as ‘internalist’ because they describe justification as consisting in internal relations among beliefs, either – as in the former case – from a vertical relation of support between supposedly basic beliefs and others dependent upon them, or – as in the latter – from the mutual support of beliefs in an appropriately understood system.

Generally characterized, internalist theories assert or assume that a belief cannot be justified for an epistemic subject *S* unless *S* has access to what provides the justification, either in fact or in principle. These theories generally involve the stronger ‘in fact’ requirement because *S*’s being justified in believing *p* is standardly cashed in terms of his having reasons for taking *p* to be true, where *having reasons* is to be understood in an occurrent sense.

Here an objection immediately suggests itself. Any *S* has only finite access to what might justify or undermine his beliefs, and that access is confined to his particular viewpoint. It seems that full justification for his beliefs would rarely be available, because his experience would be restricted to what is nearby in space and time, and he would be entitled to hold only those beliefs which his limited experience licensed.

A related objection is that internalism seems inconsistent with the fact that many people appear to have knowledge despite not being sophisticated enough to recognize that thus-and-so is a reason for believing *p* – that is the case, for example, with children.

A more general objection still is that relations between beliefs, whether of the foundationalist or coherence type, might obtain without the beliefs in question being true of anything beyond themselves. One could imagine a coherent fairy tale, say, which in no point corresponds to some external reality, but in which beliefs are justified nevertheless by their mutual relations.

This uneasy reflection prompts the thought that there should be a constraint on theories of justification, in the form of a demand that there should be some suitable connection between belief possession and external factors – that is, something other than the beliefs and their mutual relations – which determines their epistemic value. This accordingly prompts the idea of an alternative: externalism.

1.5 *Reliability, causality and truth-tracking*

Externalism is the view that what makes S justified in believing p might not be anything to which S has cognitive access. It might be that the facts in the world are as S believes them to be, and that indeed they caused S to believe them to be so by stimulating his or her sensory receptors in the right kind of way. S need not be aware that this is how his or her belief was formed. So S could be justified in believing p without it.

One main kind of externalist theory is reliabilism, the thesis – or cluster of theses – having it that a belief is justified if it is reliably connected with the truth. According to one influential variant, the connection in question is supplied by reliable belief-forming *processes*, ones which have a high success rate in producing true beliefs. An example of a reliable process might be normal perception in normal conditions.

Much apparent plausibility attaches to theories based on the notion of external linkage, especially of causal linkage, between a belief and what it is about. An example of such a theory is Alvin Goldman's (1986) account of knowledge as 'appropriately caused true belief', where 'appropriate causation' takes a number of forms, sharing the property that they are processes which are both 'globally' and 'locally' reliable – the former meaning that the process has a high success rate in producing true beliefs, the latter that the process would not have produced the belief in question in some 'relevant counterfactual situation' where the belief is false. Goldman's view is accordingly a paradigm of a reliabilist theory.

An elegant second-cousin of this view is offered by Robert Nozick (1981). To the conditions

- (1) p is true

and

- (2) S believes p

Nozick adds

- (3) if p were not true, S would not believe p

and

- (4) if p were true, S would believe it.

Conditions (3) and (4) are intended to block Gettier-type counter-examples to the justified true belief analysis by annexing S's belief that p firmly to p 's truth. S's belief

that p is connected to the world (to the situation described by p) by a relation Nozick calls 'tracking': S 's belief *tracks* the truth that p . He adds refinements in an attempt to deflect the counter-examples that philosophers are always ingenious and fertile at devising.

If these theories seem plausible it is because they accord with our pre-theoretical views. But as one can readily see, there are plenty of things to object to in them, and a copious literature does so. Their most serious flaw, however, is that they are question-begging. They do not address the question of how S is to be confident that a given belief is justified; instead they help themselves to two weighty realist assumptions, one about the domain over which belief ranges and the other about how the domain and S are connected, so that they can assert that S is justified in believing a given p even if what justifies him lies outside his own epistemic competence. Whatever else one thinks of these suggestions, they do not enlighten S , and therefore do not engage the same problem that internalist theories address.

But worst of all – so an austere critic might say – the large assumptions to which these theories help themselves are precisely those that epistemology should be examining. Externalist and causal theories, in whatever guise and combination, are better done by empirical psychology where the standard assumptions about the external world and S 's connections with it are premised. Philosophy, surely, is where these premises themselves come in for scrutiny.

1.6 Knowledge, belief and justification again

Consider this argument: 'If anyone knows some p , then he or she can be certain that p . But no one can be certain of anything. Therefore no one knows anything.' This argument (advanced in this form by Unger 1975) is instructive. It repeats Descartes's mistake of thinking that the psychological state of feeling certain – which someone can be in with respect to falsehoods, such as the fact that I can feel *certain* that Arkle will win the Derby next week, and be wrong – is what we are seeking in epistemology. But it also exemplifies the tendency in discussions of *knowledge* as such to make the definition of knowledge so highly restrictive that little or nothing passes muster. Should one care if a suggested definition of knowledge is such that, as the argument just quoted tells us, no one can know anything? Just so long as one has many well-justified beliefs which work well in practice, can one not be quite content to know nothing? For my part, I think one can.

This suggests that in so far as the points sketched in preceding paragraphs have interest, it is in connection with the *justification of beliefs* and not the *definition of knowledge* that they do so. Justification is an important matter, not least because in the areas of application in epistemology where the really serious interest should lie – in questions about the PHILOSOPHY OF SCIENCE (chapter 9), the PHILOSOPHY OF HISTORY (chapter 14) or the concepts of evidence and proof in LAW (see chapter 13) – justification is the crucial problem. That is where epistemologists should be getting down to work. By comparison, efforts to define 'knowledge' are trivial and occupy too much effort in epistemology. The disagreeable propensity of the debate generated by Gettier's counter-examples – anticipated beautifully in Russell's review of James (Russell 1910: 95) – to proceed on a chessboard of '-isms', as exemplified above, is a symptom.

The general problem with justification is that the procedures we adopt, across all walks of epistemic life, appear highly permeable to difficulties posed by scepticism. The problem of justification is therefore in large part the problem of scepticism; which is precisely why discussion of scepticism is central to epistemology.

2 Scepticism

Introduction

The study and employment of sceptical arguments might in one sense be said to define epistemology. A chief epistemological aim is to determine how we can be sure that our means to knowledge (here 'knowledge' does duty for 'justified belief') are satisfactory. A sharp way to show what is required is to look carefully at sceptical challenges to our epistemic efforts, challenges which suggest ways in which they can go awry. If we are able not just to identify but to meet these challenges, a primary epistemological aim will have been realized.

Scepticism is often described as the thesis that nothing is – or, more strongly, can be – known. But this is a bad characterization, because if we know nothing, then we do not know that we know nothing, and so the claim is trivially self-defeating. It is more telling to characterize scepticism in the way just suggested. It is a challenge directed against knowledge claims, with the form and nature of the challenge varying according to the field of epistemic activity in question. In general, scepticism takes the form of a request for the justification of those knowledge claims, together with a statement of the reasons motivating that request. Standardly, the reasons are that certain considerations suggest that the proposed justification might be insufficient. To conceive of scepticism like this is to see it as being more philosophically troubling and important than if it is described as a positive thesis asserting our ignorance or incapacity for knowledge.

2.1 Early scepticism

Some among the thinkers of antiquity – Pyrrho of Elis (c.360–c.270 BC) and his school, and Plato's successors in his Academy – expressed disappointment at the fact that centuries of enquiry by their philosophical predecessors seemed to have borne little fruit either in cosmology or ethics (this latter was broadly construed to include politics). Their disappointment prompted them to sceptical views. The Pyrrhonians argued that because enquiry is arduous and interminable, one should give up trying to judge what is true and false or right and wrong; for only thus will we achieve peace of mind.

A less radical form of scepticism overtook Plato's successors in the Academy. They agreed with Pyrrho that certainty must elude us, but they tempered their view by accepting that the practical demands of life must be met. They did not think it a workable option to 'suspend judgement' as Pyrrho recommended, and therefore argued that we should accept those propositions or theories which are more PROBABLE (pp. 308–11) than their competitors. The views of these thinkers, known as Academic sceptics, are recorded in the work of Sextus Empiricus (c.150–c.225).

In the later Renaissance – or, which is the same thing, in early modern times – with religious certainties under attack and new ideas abroad, some of the sceptical arguments of the Academics and Pyrrhonians acquired a special significance, notably as a result of the use to which René Descartes put them in showing that they are powerful tools for investigating the nature and sources of knowledge.

In Descartes's day the same person could be both astronomer and astrologer, chemist and alchemist, or physician and magician. It was hard to disentangle knowledge from nonsense; it was even harder to disentangle those methods of enquiry which might yield genuine knowledge from those that could only deepen ignorance. So there was an urgent need for some sharp, clean epistemological theorizing. In his *Meditations* (1986) Descartes accordingly identified epistemology as an essential preliminary to physics and mathematics, and attempted to establish the grounds of certainty as a propaedeutic to science. Descartes's first step in that task was to adapt and apply some of the traditional arguments of scepticism. (I shall comment on his use of scepticism again later.)

The Anatomy of Scepticism

Sceptical arguments exploit certain contingent facts about our ways of acquiring, testing, remembering and reasoning about our beliefs. Any problem that infects the acquisition and employment of beliefs about a given subject matter, and in particular any problem that infects our confidence that we hold those beliefs justifiably, threatens our hold on that subject matter.

The contingent facts in question relate to the nature of perception, the normal human vulnerability to error and the existence of states of mind – for example, dreaming and delusion – which can be subjectively indistinguishable from those that we normally take to be appropriate for acquiring justified beliefs. By appealing to these considerations the sceptic aims to show that there are significant questions to be answered about the degree of confidence that we are entitled to repose in our standard epistemic practices.

Sceptical considerations pose problems for epistemologists of both the rationalist and the empiricist camps. This division into competing schools of thought about knowledge is rough but useful, giving a shorthand way of marking the difference between those who hold that reason is the chief means to knowledge, and those who accord that role to experience. Rationalists emphasize reason because in their view the objects of knowledge are propositions that are eternally, immutably and necessarily true – the examples they offer are the propositions of mathematics and logic – and these, they say, can only be acquired by reasoning. Empiricists hold that substantive and genuine knowledge of the world can only be learned through experience, by means of the senses and their extension via such instruments as telescopes and microscopes. The rationalist need not deny that empirical awareness is an important, even an ineliminable, aid to reason, nor need the empiricist deny that reason is an important, even an ineliminable, aid to experience; but both will insist that the chief means to knowledge is respectively one or the other.

The refinements of debate about these matters merit detailed examination for which this is not the place. For present purposes, the point to note is that scepticism is a problem for both schools of thought. For both, possibilities of error and delusion pose a challenge. For the empiricist in particular, to these must be added distinctive problems about perception.

2.2 *Error, delusion and dreams*

One characteristic pattern of sceptical argument is drawn from a set of considerations about error, delusions and dreams. Consider the error argument first. We are fallible creatures; we sometimes make mistakes. If, however, we are ever to be able to claim to know (that is, at least to be justified in believing) some proposition p , we must be able to exclude the possibility that at the time of claiming to know p we are in error. But since we typically, or at least frequently, are not aware of our errors as we make them, and might therefore unwittingly be in error as we claim to know p , we are not justified in making that claim.

The same applies when a person is the subject of a delusion, illusion or hallucination. Sometimes people undergoing one or other of these states do not know that they are doing so, and take themselves to be having veridical experiences. Clearly, although they think they are in a state which lends itself to their being justified in claiming to know p , they are not in such a state. Therefore they are not justified in claiming to know p . So in order for anyone to claim knowledge of some p , they must be able to exclude the possibility that they are the subject of such states.

This pattern of argument is at its most familiar in the argument from dreaming employed by Descartes. One way of setting it out is as follows. When I sleep I sometimes dream, and when I dream I sometimes – indeed, often – do not know that I am dreaming. So I can have experiences that appear to be veridical waking experiences on the basis of which I take myself to be justified in claiming to know such and such. But because I am dreaming, I do not in fact know such and such; I merely dream that I do. Might I not be dreaming now? If I cannot exclude the possibility that I am now, at this moment, dreaming, I am unable to claim knowledge of the things I at this moment take myself to know. For example, it seems to me that I am sitting at a desk next to a window admitting a view of trees and lawns. But because I might be dreaming that this is so, I cannot claim to know it.

In these arguments the possibility of error, delusion or dreaming acts as what might be called a ‘defeater’ to knowledge claims. The pattern is: if one knows p , then nothing is acting to subvert one’s justification to claim knowledge of p . But one can seem to oneself fully entitled to claim to know some p , and in fact lack that entitlement, as the foregoing considerations show. So our claims to knowledge are in need of better grounds than we standardly take ourselves to have. We must find a way of defeating the defeaters.

2.3 *Perception*

Both rationalist and empiricist views about the sources of knowledge are threatened by the arguments just sketched. Arguments that pose particular problems for empiricism are suggested by the nature and limitations of perception, the best current account of which tells us something like the following story.

Light reflects from the surfaces of objects in the physical environment and passes into the eyes, where it irritates the cells of the retinas in such a way as to trigger impulses in the optic nerves. The optic nerves convey these impulses to the region of the cerebral cortex that processes visual data, where they stimulate certain sorts

of activity. As a result, in ways still mysterious to science and philosophy, coloured 'motion pictures' arise in the subject's consciousness, representing the world outside his or her head. This remarkable transaction is repeated *mutatis mutandis* in the other sensory modalities of hearing, smell, taste and touch, giving rise to perceptions of harmonies and melodies, perfumes and piquancies, smoothness, softness, warmth – and so forth.

This model can be used to furnish another sceptical application of the defeater argument. The complex causal story thus told is one which – so the sceptic can point out – might be interrupted in problematic ways at any point along its length. The experiences which we say result from the interaction of our senses and the world might occur in us for other reasons. They might occur when, as noted above, we dream, hallucinate or suffer delusions; or, to be fanciful, they might be produced in us by a god, or by a scientist who has connected our brains to a computer. From the point of view of the experiencing subject, there might be no way of telling the difference. So, says the sceptic, unless we find means of excluding these possibilities, we are not entitled to claim knowledge of what we standardly take ourselves to know.

2.4 *Perceptual relativities*

These same considerations about perception can prompt sceptical challenge by a different route. A little reflection of the kind taught us by Locke, BERKELEY (1685–1753) (chapter 30) and other earlier contributors to the debate shows that some of these properties we seem to perceive in objects are not 'in the objects themselves' but are in fact creatures of the perceptual relation. The qualities of objects – their colour, taste, smell, sound and texture – vary according to the condition of the perceiver or the conditions under which they are perceived. The standard examples are legion: grass is green in daylight, black at night; tepid water feels warm to a cold hand, cool to a hot hand; objects look large from close by, small from far away; and so on.

These perceptual relativities are cited by the sceptic to raise questions not just about whether perception is a trustworthy source of information about the world, but whether the world can be said to exist independently of perception at all. For what if the properties by whose means we detect the presence of objects cannot be described apart from their being objects of perception? Consider the old conundrum whether a sound is made by the tree that falls in the forest when no sentient being is present to hear it do so. The answer, on a standard theory of perception current in contemporary science, is that the tree falls in complete silence. For if there is no ear to hear, there is no sound; there are only at best the conditions – vibrating airwaves – which would cause sound to be heard if there were normally functioning eardrums, aural nerves and the rest to be stimulated by them.

These considerations suggest a sceptical picture in which perceivers are in something like the following predicament. Imagine a man wearing a visorless helmet which so encloses his head that he cannot see, hear, taste or smell anything outside it. Imagine that a camera, a microphone and other sensors are affixed to the top of the helmet, transmitting pictures and other information to its interior. And suppose finally that it is impossible for the wearer to remove the helmet to compare this information with whatever is outside, so that he cannot check whether it faithfully represents the

exterior world. Somehow the wearer has to rely on the intrinsic character of the information available inside the helmet to judge its reliability. He knows that the information sometimes comes from sources other than the exterior world, as in dreams and delusions; he has deduced that the equipment affixed to the helmet works upon the incoming data and changes it, for example adding colours, scents and sounds to its picture of what intrinsically has none of these properties (at very least, in those forms); he knows that his beliefs about what lies outside the helmet rest on the inferences he draws from the information available inside it, and that his inferences are only as good as his fallible, error-prone capacities allow them to be. Given all this, asks the sceptic, have we not a job of work to do to justify our claims to knowledge?

2.5 *Methodological and problematic scepticism*

Before considering these arguments and canvassing some ways of responding to them, it is important to note two things. One is that sceptical arguments are not best dealt with by attempts at piecemeal – that is, one-by-one – refutation. The second is that there is a vitally important distinction to be drawn between two ways in which scepticism can be employed in epistemology. It is important to note these matters because otherwise the *prima facie* implausibility of most sceptical arguments will mislead us into underestimating their significance. I take each point in turn.

Attempted refutation of sceptical arguments piecemeal is, arguably, futile for two good reasons. As suggested at the outset, sceptical arguments are at their strongest not when they seek to prove that we are ignorant about some subject matter but when they ask us to justify our knowledge claims. A challenge to justify is not a claim or a theory, and cannot be refuted; it can only be accepted or ignored. Since the sceptic offers reasons why justification is needed, the response might be to inspect those reasons to see whether the challenge needs to be met. This indeed is one good response to scepticism. Where the reasons are cogent, the next good response is to try to meet the challenge thus posed.

The second reason is that sceptical arguments taken together have the joint effect of showing that there is work to be done if we are to get a satisfactory account of knowledge – and scepticism indicates what is needed. If one could refute, or show to be ungrounded, one or another individual sceptical argument, others would be left in place still demanding that such an account be sought.

These points can be illustrated by considering Gilbert Ryle's (1900–76) attempt to refute the argument from error by using a 'polar concept' argument. There cannot be counterfeit coins, Ryle observed, unless there are genuine ones, nor crooked paths unless there are straight paths, nor tall men unless there are short men. Many concepts come in such polarities, a feature of which is that one cannot grasp either pole unless one grasps its opposite at the same time. Now, 'error' and 'getting it right' are conceptual polarities. If one understands the concept of error, one understands the concept of getting it right. But to understand this latter concept is to be able to apply it. So our very grasp of the concept of error implies that we sometimes get things right.

Ryle obviously assumed that the error sceptic is claiming that, for all we know, we might always be in error. Accordingly his argument – that if we understand the concept

of error, we must sometimes get things right – is aimed at refuting the intelligibility of claiming that we might always be wrong. But of course the error sceptic is not claiming this. He or she is simply asking how, given that we sometimes make mistakes, we can rule out the possibility of being in error on any given occasion of judgement – say, at this present moment.

But the sceptic need not concede the more general claims that Ryle makes, namely, that for any conceptual polarity, both poles must be understood, and – further and even more tendentially – to understand a concept is to know how to apply it, and for it to be applicable is for it actually to be applied (or to have been applied). This last move is question-begging enough, but so is the claim about conceptual polarities itself. For the sceptic can readily cite cases of conceptual polarities – ‘perfect–imperfect’, ‘mortal–immortal’, ‘FINITE–INFINITE’ (chapter 11) – where it is by no means clear that the more exotic poles apply to anything, or even that we really understand them. After all, taking a term and attaching a negative prefix to it does not guarantee that we have thereby grasped an intelligible concept.

These comments suggest that sceptical arguments, even if singly they appear implausible, jointly invite a serious response; which is what, in large measure, epistemology seeks to offer. But there is still the matter of the distinction between methodological and problematic scepticism to be explained, and here a brief recapitulation of Descartes’s use of sceptical arguments will be helpful.

Descartes’s Method of Doubt

Descartes’s aim was to find a basis for knowledge, which he did by looking for a starting point about which he could be certain. To find certainty he needed to rule out anything that could be doubted, however absurd that doubt, for only in this way would we be left with what is truly indubitable. In the first *Meditation* he embarks on this task by borrowing some sceptical arguments from the ancients. First he cites the fact that we can be misled in perception. But this is not a thoroughgoing enough scepticism, for even if we misperceive there is still much that we can know. So he next considers the possibility that on any occasion of claiming to know something, one might be dreaming. This sceptical thought catches more in its net, but is still insufficient, for even in dreams we can know such things as, for example, mathematical truths. So, to get as sweeping a consideration as possible, Descartes introduces the ‘evil demon’ idea. Here the supposition is that with respect to everything about which one could possibly be misled, an evil demon is indeed misleading one. Famously, what such a being cannot mislead one about is *cogito ergo sum* – when one thinks ‘I exist’, this proposition is true.

It is essential to note that Descartes’s use of these arguments is *purely methodological*. The rest of the *Meditations* is devoted to showing that we know a great deal, because the fact (as Descartes unsuccessfully tries to prove) that there is a good DEITY (see chapter 15) guarantees that, just so long as we use our faculties responsibly, whatever is perceived with clearness and distinctness to be true will indeed be true. This is because a good deity, unlike an evil one, would not wish us to embrace ignorance. Descartes was by no means a sceptic, nor did he think that sceptical arguments, least of all the one employed as a device to set aside as many beliefs as possible, were persuasive. The ‘method of doubt’ is merely a tool.

Descartes' successors, however, were far more impressed by the sceptical arguments he employed than his answer to them. For the tradition of epistemological thinking after his time, these and allied sceptical arguments were not mere methodological devices, but serious problems requiring solution. Hence the distinction I draw here between *methodological* and *problematic* scepticism.

It is clear that there are sceptical considerations that have merely methodological utility, and are not genuinely problematic, because they do not represent a stable and cogent challenge to our ordinary epistemological standards. Descartes's 'evil demon' is a case in point. Since the hypothesis that there is such a thing is as arbitrary and groundless as a hypothesis can get, it does not merit being taken seriously otherwise than as a ploy to make a point. But sceptical considerations about perception, error, delusion and dreams raise more interesting and troubling general issues, and accordingly merit examination.

Among the many things worth noting about Descartes's discussion are the following two. Firstly, as hinted earlier, his quest for certainty is arguably misconceived. Certainty is a psychological state one can be in independently of whether or not one believes truly. The falsity of a belief is no bar to one's feeling certain that it is otherwise. Descartes sought to specify ways of recognizing which of our beliefs are true, but he led himself into talk of certainty because – and this is the second point – he assumed that epistemology's task is to provide one with a way of knowing, from one's own subjective viewpoint, when one possesses knowledge. Accordingly, he starts with the private data of a single consciousness and attempts to move outside it, seeking guarantees for the process en route. Nearly all of Descartes's successors in epistemology, up to and including RUSSELL (1872–1970) (chapter 37) and Ayer (1910–89), accepted this perspective on their task. In this respect at least they are all therefore Cartesians. It is largely for this reason, as we shall see hinted below, that they found it hard to meet scepticism's challenge.

2.6 *Some responses to scepticism*

The sceptical challenge tells us that we suffer an epistemic plight, namely, that we can have the best possible evidence for believing some p , and yet be wrong. Stated succinctly and formally, scepticism is the observation that there is nothing contradictory in the conjunction of statements s embodying our best grounds for a given belief p with the falsity of p .

An informative representation of scepticism thus summarized is as follows. Sceptical arguments open a gap between, on the one hand, the grounds a putative knower has for some knowledge claim, and, on the other hand, the claim itself. Responses to scepticism generally take the form of attempts either to bridge this gap or to close it. The standard perceptual model, in which beliefs are formed by sensory interaction with the world, postulates a causal bridge across the gap; but that bridge is vulnerable to sceptical sabotage, so the causal story at least needs support. Descartes, as noted, identified the epistemological task as the need to specify a guarantee – call it X – which, added to our subjective grounds for belief, protects them against scepticism and thus elevates belief into knowledge. His candidate for X was the goodness of a deity; rejecting this candidate (while continuing to accept his view of the epistemological task)

obliges us to find an alternative. If an X cannot be found to support a bridge across the sceptical gap, the option is to try closing it – or more accurately, to show that there is no gap at all. Both the quest for X and the closing of the gap have constituted major epistemological endeavours against scepticism in modern philosophy. Some of these endeavours, in brief, are as follows.

Descartes's immediate successors were, as mentioned, unpersuaded by his attempt to bridge the gap by invoking a good divinity to serve as X. LOCKE (chapter 29), without much fanfare, employed a weaker version of the Cartesian expedient by saying that we can ignore sceptical threats to the causal story because 'the light that is set up in us shines bright enough for all our purposes'. From Locke's point of view it does not matter whether the inner light is set up by God or nature; the point is that there is something – X, the inner light which could be, perhaps, reason, empirical intuition or native trust in the reliability of the senses – that gives us grounds for accepting our ordinary knowledge-acquiring means as adequate.

Others, not content with such unsatisfactory moves, look for X elsewhere, and claim to find it in some version of foundationalism, the thought – sketched above – that our epistemic system has a basis in special beliefs that are in some way self-justifying or self-evident and which, in conjunction with the evidence we ordinarily employ in making knowledge claims, secures them against scepticism. As we saw earlier, a chief ground for rejecting such theories is alleged to be that none of them identifies satisfactory candidates for 'foundations'. But one stimulating way of making something like a foundationalist case is offered by Kant, whose attempt prompted others.

2.7 *Transcendental arguments*

KANT (1724–1804) (chapter 32) regarded failure to refute scepticism as a 'scandal' to philosophy, and offered his *Critique of Pure Reason* (1929) as a solution. His thesis is that our minds are so constituted that they impose a framework of interpretative concepts upon our sensory input, among them those of the objectivity and causal interconnectedness of what we perceive. Application of these concepts transforms mere passive receipt of sensory data into EXPERIENCE (pp. 726–33) properly so called. Our faculties are such that when raw data comes under the interpreting activity of our concepts, they have already had spatial and temporal form conferred on them by the nature of our sensory capacities; all our experience, considered as relating to what is outside us, is experience of a spatially structured world, and all our experience, considered as relating to its received character in our minds, is of a temporally structured world. Upon the spatio-temporal data thus brought before our minds we impose the categories, that is, the concepts that make experience possible by giving it its determinate character. And here is Kant's point: if the sceptic asks us to justify our claims to knowledge, we do so by setting out these facts about how experience is constituted.

Kant claimed HUME (1711–76) (chapter 31) as his inspiration for these ideas, because Hume had argued that although we cannot refute scepticism – reason was not, he claimed, up to the task – we should not be troubled, for human nature is so constituted that we simply cannot help having the beliefs that scepticism challenges us to justify. Those beliefs include, for example, that there is an external world, that causal relations hold between events in the world, that inductive reasoning is reliable, and so

forth. From this hint Kant elaborated his theory that the concepts that the sceptic asks us to justify are constitutive features of our capacity to have any experience at all.

The strategy, if not the details, of Kant's attack upon scepticism has prompted interest in more recent philosophy. The argument he employs is a transcendental argument, briefly characterizable as one which says that because A is a necessary condition for B, and, because B is the case, A must be the case also. An example of such an argument in action against scepticism is as follows.

A typical sceptical challenge concerns belief in the unperceived continued existence of objects. What justifies our holding this belief and premising so much upon it? The transcendental arguer answers that because we take ourselves to occupy a single unified world of spatio-temporal objects, and because on this view spatio-temporal objects have to exist unperceived in order to constitute the realm as single and unified, a belief in their unperceived continued existence is a condition of our thinking both about the world and our experience of it in this way. Since we do indeed think this way, the belief that the sceptic asks us to justify is thereby justified. A contemporary thinker who makes notable use of this style of argument is P. F. Strawson (b. 1919).

2.8 *Idealism and phenomenalism*

There is, in parallel to these Kantian ways of responding to the sceptical challenge, another approach, which denies the existence of a scepticism-generating gap. The chief figures in this camp are Berkeley and, more recently, the phenomenologists, who – allowing for differences among them, and for the fact that the two latter held these views only for part of their careers – include MILL (1806–73) (chapter 35), Russell and Ayer.

In Berkeley's view, scepticism arises from thinking that behind or beyond our sensory experiences there lies a *material* world. The word 'material' means 'made of matter', and 'matter' is a technical philosophical term supposed to denote an empirically undetectable substance believed by Berkeley's philosophical predecessors to underpin the sensorily detectable properties of things, such as their colours, shapes and textures. Berkeley rejected the concept of matter thus understood – it is a common misreading of him to take it that he thereby denied the existence of physical objects; he did no such thing – arguing that because physical objects are collections of sensible qualities, and because sensible qualities are ideas, and because ideas can only exist if perceived, the existence of objects therefore consists in their being perceived; if not by finite minds such as our own, then everywhere and at all times by an infinite mind. (We may note that Berkeley thought that his refutation of scepticism was at the same time a powerful new argument for the existence of God.)

Berkeley's habit of saying that things exist 'in the mind' has led uncritical readers to suppose he means that objects exist only in one's head, which is what a subjective idealist or solipsist might try to hold. Berkeley's idealism, whether or not it is otherwise defensible, is at least not quite so unstable a view. His 'in the mind' should be read as meaning 'with essential reference to experience or thought'.

For present purposes, the point is that Berkeley sought to rebut scepticism by denying the existence of a gap between experience and reality, on the grounds that experience

and reality are the same thing. (He had a theory of how, despite this, we can nevertheless imagine, dream and make mistakes.) The phenomenologists, with one very important difference, argued likewise. Their view, briefly stated, is that all our beliefs about the world are derived from what appears to us in experience. When we analyse appearances – the ‘phenomena’ – we see that they are built out of the basic data of sense: the smallest visible colour patches in our visual fields, the least sounds in our auditory fields. Out of these sense-data we ‘logically construct’ the chairs and tables, rocks and mountains constituting the familiar furniture of the everyday world.

An alternative but equivalent way of putting this point, the phenomenologists claim, is to say that statements about physical objects are merely convenient shorthand for longer and more complicated statements about how things seem to us in the usual employment of our sensory capacities. And to say that objects continue to exist unperceived is to say – in Mill’s phrase – that they are ‘permanent possibilities of sensation’, meaning that one would experience them if certain conditions were fulfilled.

Berkeley holds that things remain in existence when not perceived by finite minds because they are perceived by a deity. The phenomenologists argue that what it means to say that things exist unperceived is to say that certain counterfactual conditionals are true, namely, those asserting that the things in question would be perceived if some perceiver were suitably placed with respect to them. These conditionals are notoriously problematic, because it is not clear how to understand them. What, in particular, makes them true when they are (or seem quite obviously to be) true? The usual answers, in terms of possible worlds, laws, ideal regularities and similar exotica, do little to help. It is not clear that much of an advance is made over Berkeley’s ubiquitous deity by substituting barely true counterfactuals in its place. Berkeley’s view has the modest attraction that everything in the world is actual – anything that exists is perceived – whereas in the phenomenologist’s universe most of what exists does so as a possibility rather than an actuality, namely, as a possibility of perception.

One thing is clear, at least: that one does not get phenomenology simply by subtracting the theology from Berkeley’s theory. One has to do that and then, in the resulting metaphysical gap, substitute a commitment to the existence of barely true counterfactuals, with an accompanying commitment to the existence of possibilities. Both Berkeley’s theory and phenomenology thus exact high prices for closing the sceptical gap.

2.9 *Sceptical epistemology versus anti-Cartesianism*

Some epistemologists do not attempt to refute scepticism for the good reason that they think it true or at least irrefutable. Their views might be summarized as stating that scepticism is the inevitable result of epistemological reflection, so we should accept either that we are only ever going to have imperfectly justified beliefs, always subject to revision in the light of experience, or that we have to recognize that scepticism, despite being irrefutable, is not a practical option, and therefore we have to live as most people anyway do, namely, by simply ignoring it.

Some commentators on Hume interpret him as taking this latter view of the matter, and accordingly call it the ‘Humean’ response to scepticism. In Stroud (1984) and Strawson (1985), something like the Humean view is taken.

Others in the recent debate are more combative, among them DEWEY (1859–1952) (see chapter 36) and WITTGENSTEIN (1889–1951) (chapter 39). Despite substantial differences in other respects, these two thinkers hold an interesting view in common, which is that scepticism results from accepting the Cartesian starting-point among the private data of individual consciousness. If instead, they say, we begin with the public world – with considerations relating to facts about the essentially public character of human thought and language – a different picture emerges.

Dewey argued that the Cartesian model makes the epistemic subject a merely passive recipient of experiences, like someone sitting in the dark of a cinema watching the screen; but, he pointed out, ours is in fact a participant perspective – we are actors in the world, and our acquisition of knowledge is the result of our doings there.

Wittgenstein contested the very coherence of the Cartesian approach by arguing that PRIVATE LANGUAGE (pp. 817–20) is impossible. A private language in Wittgenstein's sense is one that is logically available only to one speaker, which is what a Cartesian subject would need in order to begin discoursing about his private inner experience. His argument is this: language is a rule-governed activity, and one only succeeds in speaking a language if one follows the rules for the use of its expressions. But a solitary would-be language-user would not be able to tell the difference between actually following the rules and merely believing that he is doing so; so the language he speaks cannot be logically private to himself; it must be shareable with others. Indeed, Wittgenstein argues that language can only be acquired in a public setting (he likens language-learning to the training of animals; to learn a language is to imitate the linguistic behaviour of one's teachers), which similarly weighs against the idea that the Cartesian project is even in principle possible.

The anti-sceptical possibilities of the private language argument seem not to have been wholly apparent to Wittgenstein himself. In draft notes on scepticism and knowledge written in the last months of his life – later published under the title *On Certainty* (1969) – he offers a response to scepticism, which marks a return to a more traditional approach, not unlike that offered by Hume and Kant. It is that there are some things we have to accept in order to get on with our ordinary ways of thinking and speaking. Such propositions as that there is an external world, or that the world came into existence a long time ago, are simply not open to doubt; it is not an option for us to question them. Nor therefore, says Wittgenstein, can we say that we know them, because knowledge and doubt are intimately related, in that there can only be knowledge where there can be doubt, and vice versa.

The propositions we cannot doubt constitute the 'scaffolding' of our ordinary thought and talk, or – Wittgenstein varies his metaphors – they are like the bed and banks of a river, down which the stream of ordinary discourse flows. In this sense the beliefs that scepticism attempts to challenge are not open to negotiation; which, says Wittgenstein, disposes of scepticism.

These thoughts are as suggestive as they are in the philosophies of Hume and Kant; but one of the problems with Wittgenstein's way of putting them is that he uses foundationalist concepts in describing the relation of 'grammatical' propositions to ordinary ones, but repudiates foundationalism as such, and seems to allow a version of relativism by doing so – the river's bed and banks, he says, might in time be worn away. But relativism is just scepticism in disguise – it is, indeed, arguably the most powerful and

troubling form of scepticism, for it is the view that knowledge and truth are relative to a point of view, a time, a place, a cultural or cognitive setting: and knowledge and truth thus understood are not knowledge and truth.

Concluding Remarks

There is much one would like to insist upon in trying correctly to describe the work that needs to be done in epistemology, for that is the necessary preliminary to making what progress we can. Here I shall simply underline a couple of remarks already made above.

Firstly, debates over the definition of 'knowledge' seem to me to be a side-show. The justification of claims in the natural sciences, the social sciences (not least history) and law is where the real work cries out to be done in epistemology. And this comment applies only to the empirical case: what of the epistemological questions that press in ETHICS (chapter 6) and the PHILOSOPHY OF MATHEMATICS (chapter 11)? There can be no guarantee – and indeed it is unlikely – that high generalities about justification and knowledge will apply univocally across all these fields. 'Justification' is a dummy concept that needs to be cashed out in terms particular to particular fields; so much should be obvious from the fact that unrestrictedly general accounts of justification prove hopelessly vulnerable to counter-example.

Secondly, little in current literature about scepticism makes one confident that its nature is properly understood. Scepticism defines one of the central problems in epistemology, namely, the need to show how justification of belief is possible. This is done by meeting the challenge to show that sceptical considerations do not after all defeat our best epistemic endeavours in this or that specified field. Implicit in this characterization are two important claims: firstly, that scepticism is best understood as a *challenge*, not as a *claim* that we do or can know nothing; and secondly, that the best way to respond to scepticism is not by attempting to refute it on an argument-by-argument basis, but by showing how we come by justification for what we believe. Somehow these two points, which were obvious to our predecessors, seem to have been lost to sight.

Further Reading

General

Some useful texts are R. Nozick (1981); J. Dancy and E. Sosa (1992); B. Williams (1978); K. Lehrer (1974); L. BonJour (1985); P. F. Strawson (1985); and G. Pappas and M. Swain (1978).

The classic texts in epistemology include Plato's *Meno* and *Theaetetus*, Descartes's *Meditations*, Locke's *Essay Concerning Human Understanding*, Berkeley's *Principles of Human Knowledge* and *Three Dialogues Between Hylas and Philonous*, Hume's *Treatise of Human Nature* and *Enquiry Concerning Human Understanding*, and Kant's *Critique of Pure Reason*. One of the best short elementary books remains Russell's *The Problems of Philosophy* (1912, much reprinted).

Knowledge

The debate about knowledge and justification commands a large literature, of which the following are good examples: W. Alston (1983); L. BonJour (1985); A. Brueckner (1988); R. Chisholm (1977); J. Dancy (1985); F. Dretske (1971); R. Feldman (1985); E. Gettier (1963); A. Goldman (1979, 1980, 1986); G. Harman (1973, 1984); K. Lehrer (1974); P. Moser (1985); R. Nozick (1981); J. Pollock (1979, 1984, 1986); R. Shope (1983); E. Sosa (1981).

Scepticism

The best general introduction to sceptical arguments remains Bertrand Russell's *The Problems of Philosophy*, but it is essential to see the arguments in a classic setting, and for this one must read René Descartes's *Meditations on First Philosophy* (translated by J. Cottingham, 1986), especially the *First Meditation*. Useful discussions of sceptical arguments and the sense-datum theory are to be found in A. J. Ayer (1956) and J. L. Austin (1961), although one should also look at Ayer's reply (1967). For discussion of perception see J. Dancy (1988); T. Crane (1992); R. Swartz (1965); F. Jackson (1977); and M. Perkins (1983).

For an attempt at being sceptical see P. Unger (1975). For responses to scepticism influenced by Kant see P. E. Strawson (1959, 1985) and A. C. Grayling (1985). Allied lines of thought occur in G. E. Moore (1959) and L. Wittgenstein (1969). More recent discussions are B. Stroud (1984) and M. Williams (1991). Scepticism, foundationalism and coherence theories of knowledge are discussed in K. Lehrer (1974) and in useful papers collected by G. S. Pappas and M. Swain (1978) and M. Clay and K. Lehrer (1989). For a discussion of the views variously taken by Dewey and Wittgenstein see R. Rorty (1979). A textbook which surveys the field and provides a useful bibliography is J. Dancy (1985). For the history of scepticism see M. Burnyeat (1983) and R. Popkin (1979).

References

Classical texts

- Berkeley, G. 1995 [1710]: *The Principles of Human Knowledge* (edited by H. Robinson). Published with *Three Dialogues*. Oxford: Oxford University Press.
- 1995 [1713]: *Three Dialogues Between Hylas and Philonous* (edited by H. Robinson). Published with *The Principles of Human Knowledge*. Oxford: Oxford University Press.
- Descartes, R. 1986 [1641]: *Meditations on First Philosophy* (translated by J. Cottingham). Cambridge: Cambridge University Press.
- Hume, D. 1978 [1739–40]: *A Treatise of Human Nature* (edited by P. H. Nidditch). Oxford: Clarendon Press.
- 1975 [1748]: *Enquiry Concerning Human Understanding* (edited by P. H. Nidditch). Oxford: Clarendon Press.
- Kant, I. 1929 [1781 and 1787]: *Critique of Pure Reason* (1st edn 1781, 2nd edn 1787) (translated by N. K. Smith). London: Macmillan.
- Locke, J. 1975 [1690]: *An Essay Concerning Human Understanding* (edited by P. H. Nidditch). Oxford: Clarendon Press.
- Plato 1961: *Meno and Theaetetus*. In *The Collected Dialogues of Plato* (edited by E. Hamilton and H. Cairns). Princeton, NJ: Princeton University Press.

Other writings

- Alston, W. 1983: What is Wrong with Immediate Knowledge? *Synthese*, 55.

- Austin, J. L. 1961: *Sense and Sensibilia*. Oxford: Oxford University Press.
- Ayer, A. J. 1956: *The Problem of Knowledge*. London: Macmillan.
- 1967: Has Austin Refuted the Sense-Datum Theory? *Synthèse*, 17.
- BonJour, L. 1985: *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Brueckner, A. 1988: Problems with Internalist Coherentism. *Philosophical Studies*, 54.
- Burnyeat, M. (ed.) 1983: *The Sceptical Tradition*. Berkeley: University of California Press.
- Chisholm, R. 1977: *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice-Hall.
- Clay, M. and Lehrer, K. (eds) 1989: *Knowledge and Scepticism*. Boulder, CO: Westview Press.
- Crane, T. (ed.) 1992: *The Contents of Experience*. Cambridge: Cambridge University Press.
- Dancy, J. 1985: *Introduction to Contemporary Epistemology*. Oxford: Blackwell.
- (ed.) 1988: *Perceptual Knowledge*. Oxford: Oxford University Press.
- Dancy, J. and Sosa, E. (eds) 1992: *A Companion to Epistemology*. Oxford: Blackwell.
- Dretske, F. 1971: Conclusive Reasons. *Australasian Journal of Philosophy*, 49.
- Feldman, R. 1985: Reliability and Justification. *The Monist*, 68.
- Gettier, E. 1963: Is Justified True Belief Knowledge? *Analysis*, 23.
- Goldman, A. 1979: What is Justified Belief? In G. Pappas (ed.) *Justification and Knowledge*. Dordrecht: Reidel.
- 1980: The Internalist Conception of Justification. In *Midwest Studies in Philosophy* 5, ed. P. A. French et al. Minneapolis: University of Minnesota Press.
- 1986: *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Grayling, A. C. 1985: *The Refutation of Scepticism*. London: Duckworth.
- Harman, G. 1973: *Thought*. Princeton, NJ: Princeton University Press.
- 1984: Positive versus Negative Undermining in Belief Revision. *Noûs*, 18.
- Jackson, F. 1977: *Perception*. Cambridge: Cambridge University Press.
- Lehrer, K. 1974: *Knowledge*. Oxford: Clarendon Press.
- Moore, G. E. 1959: *Philosophical Papers*. London: Allen and Unwin.
- Moser, P. 1985: *Empirical Justification*. Dordrecht: Reidel.
- Nozick, R. 1981: *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pappas, G. and Swain, M. (eds) 1978: *Knowledge and Justification*. Ithaca, NY: Cornell University Press.
- Perkins, M. 1983: *Sensing the World*. Indianapolis: Hackett Publishing.
- Pollock, J. 1979: A Plethora of Epistemological Theories. In G. Pappas (ed.) *Justification and Knowledge*. Dordrecht: Reidel.
- 1984: Reliability and Justified Belief. *Canadian Journal of Philosophy*, 14.
- 1986: *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman and Littlefield.
- Popkin, R. 1979: *The History of Scepticism*. Berkeley: University of California Press.
- Rorty, R. 1979: *Philosophy and the Mirror of Nature*. Princeton, NJ: Princeton University Press.
- Russell, B. 1910: *Philosophical Essays*. New York: Longman, Green.
- 1912: *The Problems of Philosophy*. London: Oxford University Press.
- Shope, R. 1983: *The Analysis of Knowing*. Princeton, NJ: Princeton University Press.
- Sosa, E. 1981: The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge. In *Midwest Studies in Philosophy* 5, ed. P. A. French et al. Minneapolis: University of Minnesota Press.
- Strawson, P. F. 1959: *Individuals*. London: Methuen.
- 1985: *Scepticism and Naturalism: Some Varieties*. London: Methuen.
- Stroud, B. 1984: *The Significance of Philosophical Scepticism*. Oxford: Clarendon Press.
- Swartz, R. (ed.) 1965: *Perceiving, Sensing and Knowing*. Berkeley: University of California Press.

Unger, P. 1975: *Ignorance*. Oxford: Clarendon Press.

Williams, B. 1978: *Descartes: The Project of Pure Enquiry*. Harmondsworth: Penguin Books.

Williams, M. 1991: *Unnatural Doubts*. Oxford: Blackwell.

Wittgenstein, L. 1969: *On Certainty*. Oxford: Blackwell.

Discussion Questions

- 1 How important is it to have a definition of knowledge?
- 2 If a proposition is false, can one be justified in believing it?
- 3 Why do 'Gettier examples' raise difficulties for a fallibilist account of knowledge?
- 4 Could states other than beliefs provide the foundations of knowledge?
- 5 If there are foundational beliefs, how are they related to dependent beliefs?
- 6 Is a belief justified if it coheres with an already accepted set of beliefs? Do we have an adequate account of the notion of coherence?
- 7 Does justification consist in internal relations among beliefs?
- 8 Can a belief be justified for someone who does not have cognitive access to what justifies the belief?
- 9 Is knowledge 'appropriately caused true belief'?
- 10 Does the notion of 'tracking' help to explain how beliefs are justified?
- 11 If there is no certainty, can there be knowledge?
- 12 Should we seek a single account of justified belief or different accounts tailored to the different areas in which epistemological questions may be asked?
- 13 What role does scepticism play in philosophy?
- 14 How can we disentangle those methods of enquiry that might yield genuine knowledge from those that can only deepen our ignorance?
- 15 Our knowledge claims are sometimes in error without our knowing it. Does this undermine justification for any such claims?
- 16 Could you be dreaming now?
- 17 Could the experiences I take to be of the world have some other origin, without my being able to tell that this is so?
- 18 Can the properties by which we detect the presence of objects be described apart from their being objects of perception?
- 19 Are we like the man in a visorless helmet who cannot check the information transmitted to him to see whether it faithfully represents the external world?
- 20 Can sceptical arguments that are singly implausible jointly require a serious response?
- 21 What is the importance of distinguishing methodological and problematic scepticism?
- 22 Can epistemology provide a way of knowing from a first-person subjective viewpoint? Is there any other viewpoint available?
- 23 If scepticism opens a gap between the grounds for a knowledge claim and the claim itself, is it better to bridge the gap or to close it?
- 24 Could the possibility of our having experience be unintelligible to us unless we held a certain belief, and yet that belief be false?
- 25 Are experience and reality the same thing?
- 26 Can we accept the role of counterfactuals in a phenomenalist account of physical objects?

27 What follows from beginning our account of knowledge with the public world rather than with the private data of individual consciousness?

28 Does the 'private language argument' show that the Cartesian project is impossible?

29 If there are propositions that are simply not open to doubt, how can we identify them? How are they related to propositions that we can doubt?