

Understanding Human
Knowledge

Philosophical Essays

BARRY STROUD

OXFORD
UNIVERSITY PRESS

Understanding Human Knowledge

Philosophical Essays

BARRY STROUD

OXFORD
UNIVERSITY PRESS

To my mother and to the memory of my father.

They did everything to make many things possible.

So I had only to make something actual.

INTRODUCTION

I collect here fourteen essays published over twice as many years in a number of different journals and books. In bringing them together in one place I hope to make them more conveniently and so more widely available. Each essay was written to stand on its own, independently of the others, and I have not altered the texts from their original appearance. But I hope that each might gain from being read in connection with others that deal in similar ways with closely related questions. Those written later develop or explore some of the earlier ideas in directions that were not foreseen at the time. I thank the original publishers of the essays for permission to reprint them here.

The essays centre on the task of understanding human knowledge, as it is pursued in philosophy. What has come to be called 'epistemology' is the attempt to explain how we know the things we know. It would seem that any successful explanation of our knowledge should carry the implication that we do in fact know at least many of the things we think we know. That conclusion would be the denial of philosophical scepticism, which says that we know nothing, or (within a certain range) that we know nothing of a certain sort. So it would seem that a serious philosophical interest in human knowledge, either in general or for this or that restricted range of subject-matter, cannot be separated from a concern with the truth or falsity of scepticism for the domain in question.

I say that is how it would seem, but in fact many philosophers declare that they have no interest in scepticism. That is understandable, since many philosophers have no interest in epistemology either, and there is no reason why they should. But many who claim to have no interest in philosophical scepticism also put forward philosophical theories of human knowledge, or of this or that region of it. Those theories are positive answers to some question, and it is presumably the same question to which scepticism gives a conflicting negative answer. Many of the essays in this volume focus on the importance, and the difficulty, of identifying that question or issue, and of arriving at an answer we can find

satisfactory. What exactly is the philosophical question? What is problematic, or to be explained, about human knowledge? What does a philosophical theory of (this or that domain of) knowledge seek?

Without some idea of an answer to these questions there will be no firm way to assess philosophical theories of knowledge, or to understand what they are meant to do, or even what they say. It is presumably not enough for them simply to say something or other that is true about human beings and how they know things. I think the issue can be identified more sharply by using the challenge of philosophical scepticism as the criterion and asking what it would take to show that scepticism in this or that area is not true. That would identify a task for the theory of knowledge, and a question for it to answer. That strategy would require an understanding of the challenge raised by philosophical scepticism, and of its source. But when that question, and the conditions of a successful positive answer to it, are more clearly understood, it comes to seem undeniable that scepticism is the right, in fact the only possible, answer: we could know nothing at all in the domain in question if knowing required what the sceptical challenge demands.

This leads to attempts to expose that epistemological question as meaningless, or as imposing impossibly high standards, or as in some other way illegitimate. If that could be shown, scepticism would be no threat, even if it seems like the only answer to that question. Many different diagnoses have been offered. In the view of J. L. Austin, which was shared by many linguistic philosophers in the 1950s and later, there is simply no intelligible epistemological enterprise. It is the result of nothing more than misuse or misunderstanding of a few familiar words, and of concentrating obsessively on a few half-studied 'facts' of perception. No serious philosophical issue is really at stake, and no substantive philosophical pay-off could be expected from it. The most that can be done is to 'dismantle the whole doctrine before it gets off the ground'.¹ That is a matter of 'unpicking, one by one, a mass of seductive (mainly verbal) fallacies, of exposing a wide variety of concealed motives—an operation which leaves us, in a sense, just where we began'.² Where we began is presumably wherever we were before we took a philosophical interest in the possibility of human knowledge.

¹ J. L. Austin, *Sense and Sensibilia* (Oxford, 1962), 142.

² *Ibid.* 4–5.

I think there is much more than that to be gained from a study of the sources of the philosophical problem of knowledge and its attendant scepticism. But at least this dismissive view is consistent in finding unintelligible or ill-formed not only scepticism and the question to which it is an answer, but also theories of knowledge that appear to give more positive answers to that same question. That is better than trying to show that the question is meaningless or of no interest and then giving what looks like a positive non-sceptical answer to it. That is what it seems many philosophers have continued to do. But when some of the 'concealed motives' of philosophical scepticism are brought to light and properly understood, I think it can be seen that many responses to it, and many diagnoses of its failings, are superficial or off target, and do not really get to the heart of it. Seeing how and why that is so can amount to what I regard as a real advance in philosophical understanding. In *The Significance of Philosophical Scepticism*³ I considered a number of promising-looking responses to see how far they could go.

I have included here only essays which for the most part do not take up ideas and themes most emphasized in that book. There are two or three exceptions, but almost all the rest were written after the book was finished. I have arranged them almost chronologically, except for placing first 'Scepticism and the Possibility of Knowledge', which appeared in the *Journal of Philosophy* in 1984. It can serve as an introduction to the problem of identifying the challenge of scepticism as I presented it in very general terms to a symposium at a meeting of the American Philosophical Association.

'Transcendental Arguments', which was published in the *Journal of Philosophy* in 1968, describes and begins to explore a very ambitious and potentially devastating response to scepticism. It would prove a priori, in the spirit of Kant, that we know many of the things we think we know, and that our having such knowledge is a condition of our being able to think anything at all. If that were so, no one could even think that scepticism is correct and be right. That would be a very reassuring conclusion for those who seek an understanding of how human knowledge is possible, and it would also help explain what can be felt to be the depth of scepticism: the

³ (Oxford, 1984).

question is about the very conditions of the possibility of human thought.

The essay casts doubt on the prospects of revealing the impossibility of scepticism by transcendental arguments which would establish conclusions about the way things are in the world, especially when they are severed from the idealism that their success appears to depend on in Kant. The essay now seems to me very compressed, and to try to do too many things at once to succeed in all of them, but I think the basic dilemma still stands out clearly enough to carry conviction. It is an assessment of a certain streamlined, or what Strawson called 'scaled-down', Kantian turn or strategy in philosophy, not of the philosophy of Kant itself, which would be a much more complex story. What can reasonably be expected from an epistemological strategy along some such lines is a recurring topic of many of the later essays.

'Doubts about the Legacy of Scepticism has not been printed before. It was my response to Thompson Clarke's 'The Legacy of Scepticism'⁴ in an American Philosophical Association symposium in 1972. In that paper Clarke illustrates the way in which use of the same form of words, even apparently with the same meanings, can be seen to have a different point, or different implications, in the ordinary affairs of everyday life from what it has in the philosophical investigation of knowledge. This is one thing that makes it so difficult to identify uniquely the elusive, peculiarly philosophical, question of knowledge. The distinction lies behind the difficulties I mention in 'Scepticism and the Possibility of Knowledge', and is one of the main themes of *The Significance of Philosophical Scepticism*. Part of this essay appears in slightly different form in the last chapter of that book. Despite the overlap, I publish it here because it expounds some of the central views of Clarke's important, rich, and difficult paper, and focuses attention on one of the crucial issues on which its success appears to depend.

'Taking Scepticism Seriously' appeared as a long review of Peter Unger's *Ignorance: A Case for Scepticism*⁵ in the *Journal of Philosophy* in 1977, but without that title. I there lament the fact that most philosophers at the time did not take scepticism seriously enough, but Unger certainly takes it seriously in that book, even to the point of believing it, or at least saying he believes it. He shows

⁴ *Journal of Philosophy* (1972).

⁵ (Oxford, 1975).

in interesting ways how his believing that he knows nothing is compatible with his speaking and acting as all the rest of us do in everyday life. He goes further and claims not only that he does not thereby violate or distort the meanings of the epistemic words we all use, such as “know” and “certain”, but that our understanding of the meanings of those words actually commits us to scepticism. I argue that something more is needed to get us all the way to that philosophical conclusion, and make some suggestions about what it is. This too is close to the account in *The Significance of Philosophical Scepticism*.

‘Reasonable Claims: Cavell and the Tradition’ has not been previously published in exactly this form. A shorter version, amounting to some two-thirds of what is here, appeared in the *Journal of Philosophy* in 1980 as part of an American Philosophical Association symposium on Stanley Cavell’s *The Claim of Reason*.⁶ Some of the ideas of this essay also appeared in different form in the last chapter of *The Significance of Philosophical Scepticism*. I include the longer version here because it gives a fuller appreciation of some of the most distinctive and most important aspects of Cavell’s response to scepticism, beyond what I mention in my book. I believe Cavell’s reply at the APA meeting was largely based on this longer version.

‘Transcendental Arguments and “Epistemological Naturalism”’ was published in *Philosophical Studies* in 1977. It is in part my reply to Jay Rosenberg’s ‘Transcendental Arguments Revisited’⁷ in an American Philosophical Association symposium under that name. The ‘epistemological naturalism’ in question is Rosenberg’s somewhat Peirce-like view of enquiry, and is not meant to include everything that could be given that label today. In this brief essay I focus a little more on the requirements of the Kantian project itself, and come back to the question of the special status of those conclusions or ‘principles’ that would be established by a successful transcendental argument—what in ‘Transcendental Arguments’ I call the ‘privileged class of propositions’. Exploration of the possibility of establishing some such conclusions, perhaps of different status and varying strengths, but all within a broadly Kantian or ‘transcendental’ strategy, runs through several of the later essays.

In ‘The Allure of Idealism’, which appeared in the *Proceedings*

⁶ (New York, 1979).

⁷ *Journal of Philosophy* (1975).

of the *Aristotelian Society: Supplementary Volume* in 1984, the Kantian project gets more detailed attention, and I try to bring out more fully how idealism is implicated in it and required for its success. I also expand a little on the profound disappointment, not to say unintelligibility, of transcendental idealism as a solution to whatever problem of knowledge of the objective world we might have felt at the beginning. The essay was presented in a Joint Session symposium with Jonathan Lear in which he went on from Kant to defend a very Kantian reading of the later works of Wittgenstein, even arguing for the presence of something very like Kant's 'transcendental deduction' in *Philosophical Investigations*.⁸ This was in effect to make a case for Wittgenstein as an idealist. I give some reasons for thinking the case cannot be supported. But Lear raises an important question about what exactly we can expect to get from Wittgenstein's writings, and especially from his drawing attention for philosophical purposes to what he calls facts of our 'natural history'. Some of my essays on Wittgenstein in *Meaning, Understanding, and Practice: Philosophical Essays*⁹ touch on this question.

'Understanding Human Knowledge in General' was published in M. Clay and K. Lehrer (eds.), *Knowledge and Skepticism* (Boulder, Colo. 1989). It returns to the question of the conditions of a fully satisfactory answer to the distinctively philosophical problem of knowledge, filling in with some detail the outline offered in 'Scepticism and the Possibility of Knowledge'. I give reasons for thinking that, because of the way we apparently want to understand human knowledge in philosophy, there is and could be no satisfactory answer to the question. This turns in part on the extreme generality of what we seek. Because we want to understand how *all* knowledge of a certain kind is possible, we cannot rely on any knowledge of that kind we might think we already possess in order to explain how we know anything at all of that kind. It seems as if a satisfactory explanation must then explain how we can get from knowledge of some other kind to knowledge of the kind in question. But no such step seems available.

Some theories of knowledge would deny that such a step is needed; someone who has knowledge of the kind in question does

⁸ See his 'The Disappearing "We"', *Proceedings of the Aristotelian Society: Supplementary Volume* (1984).

⁹ (Oxford, 2000).

not have to know that the step is justified in order to have that knowledge. I explain the point of that 'externalist' strategy and argue that, given the problem of knowledge with which we began, 'externalism' alone would still leave us in an unsatisfactory position. We would not really understand how we know the things we do, even if we do know them. The position we would be left in, and its unsatisfactoriness, are difficult to describe accurately, for reasons I try to explain, and perhaps even illustrate.

'Epistemological Reflection on Knowledge of the External World' was published in *Philosophy and Phenomenological Research* in 1996. It is my response to Michael Williams's *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*¹⁰ as presented in a symposium devoted to that book at Northwestern University. Contrary to Williams, it stresses the 'naturalness' or apparent irresistibility of the general reflections that lead to philosophical questions about knowledge, and so to the subsequent scepticism. Williams holds that those reflections depend from the very beginning on abstract philosophical assumptions and theories which distort or obscure the phenomena in ways that guarantee the unsatisfactory result. He and I agree that what is needed is a satisfying diagnosis of the source of the appeal and apparent inevitability of scepticism in philosophy. We differ about how to do it, or what rewards can be expected from it. I express doubts about the details of some of his diagnoses, in particular about the role of what he calls 'epistemological realism'. But if there are philosophical assumptions or theories at work in leading us to scepticism, I think understanding what seems to support them, and why it is so easy for us to fall into them, is likely to prove philosophically rewarding. I would put more weight on the demand for complete generality that is felt to be essential to the philosophical project, and on the (perhaps all too human) urge to get outside the human condition which it seems to be an expression of.

'Scepticism, "Externalism", and the Goal of Epistemology' was in part a reply to Ernest Sosa's 'Philosophical Scepticism and Epistemic Circularity' in a Joint Session symposium, and was published with his paper in the *Proceedings of the Aristotelian Society: Supplementary Volume* in 1994. It stresses the unavoidability of a concern with scepticism in any attempt to defend a philosophical

¹⁰ (Oxford, 1991).

theory of knowledge, and argues again for the idea that scepticism is probably the only right answer to the epistemological question. Sosa disagrees, and defends a version of 'externalism' by opposing the idea that it must leave us ultimately dissatisfied—something I had argued for in 'Understanding Human Knowledge in General'. But the reasons I gave for that verdict there are not really taken up and assessed by Sosa, so in effect I give them again. The point is elusive, and requires a certain difficult imaginative projection. I do not know how well I succeed in bringing out what I have in mind.

'Kantian Argument, Conceptual Capacities, and Invulnerability' appeared first in P. Parrini (ed.), *Kant and Contemporary Epistemology* (Dordrecht, 1994). It begins with something of a review of the problem of transcendental arguments as I saw it at that time, and of the development of P. F. Strawson's views on the subject. In 'Transcendental Arguments' I had taken his discussions of objective particulars and of persons in his *Individuals*¹¹ as giving perhaps the clearest and most persuasive arguments of the kind I thought could not succeed. In his *Skepticism and Naturalism: Some Varieties*¹² in 1985 he conceded that perhaps they could not go as far as he at one time had appeared to think. He thought we should settle for 'a certain sort of interdependence of conceptual capacities and beliefs', without supposing we can reach conclusions as to how things actually are.

The rest of my essay gives reasons against retreating so far so quickly, and giving up all hope of arriving by general argument at conclusions that enjoy some special status or other, and so have some special standing in our conception of the world and in our knowledge of it. The question is what is possible down that road: perhaps there is a way to disarm certain forms of philosophical scepticism without having to go as far as establishing truths about the independent world. I give the name "invulnerability" to that special status and begin to explore different forms it might take and different ways in which it might be shown to belong to some of our beliefs.

The depth and potential power of a broadly Kantian strategy lie in the fact that it begins with the very conditions of thought and experience. If scepticism is correct we have no knowledge, but even

¹¹ (London, 1959).

¹² (London, 1985).

to consider the question or to entertain the sceptical challenge we must at least have certain coherent thoughts. We must think of ourselves and others as having a set of beliefs about the world even to be in a position to contemplate drawing the sceptical conclusion that none of those beliefs is reasonable or amounts to something we know. In 'Radical Interpretation and Philosophical Scepticism' I examine Donald Davidson's account of belief and belief-attribution to see what anti-sceptical lessons can be drawn from it. The essay was first published in the Library of Living Philosophers volume, *The Philosophy of Donald Davidson*, edited by L. E. Hahn (LaSalle, Ill., 1999), but written several years before that.

On Davidson's view our beliefs about what the world is like in general and our beliefs about what thoughts and beliefs there are in the world must go hand in hand. This puts limits on how far we could find the contents of the beliefs we attribute to people to deviate in general from the way things are. My essay is an exploration of those limits, and where they would have to lie in order to provide an effective defence against the threat of philosophical scepticism. I argue that the actual truth of most of our beliefs cannot be derived solely from fulfilment of the conditions of our attributing them to ourselves and others, at least not without a doubtful transcendental argument that would appear to depend on some form of idealism. I try to indicate how an anti-sceptical argument could still succeed by demonstrating, not the truth, but only what I have called a certain kind of 'invulnerability' of the beliefs in question.

Discussion of the idea of 'invulnerability' is carried a few small steps further in 'The Goal of Transcendental Arguments'. That was my contribution to a conference which 'revisited' transcendental arguments once again, this time in Sheffield, England, and was published in Robert Stern (ed.), *Transcendental Arguments: Problems and Prospects* (Oxford, 2000). It first gives a quick review of some of the special features of transcendental arguments, overlapping to some extent with earlier essays. It then adds a distinction between two different sources of the kind of 'invulnerability' I have in mind, or two different ways of showing that certain beliefs have it. Being indispensable to any conception of an independent world at all would render a belief invulnerable to exposure as mere illusion, but beliefs can be invulnerable in a similar way even without

enjoying that kind of indispensability. Exactly what kind and degree of resistance to scepticism can thereby be secured is certainly a question for further investigation. In illustrating the idea in this essay with the example of the colours of things, I invoke the much fuller treatment of the whole strategy given in my *The Quest for Reality: Subjectivism and the Metaphysics of Colour*,¹³ where epistemological scepticism is not really the issue at all. I draw heavily in that book on ideas first presented in the essays 'Kantian Argument, Conceptual Capacities, and Invulnerability' and 'Radical Interpretation and Philosophical Scepticism'.

'The Synthetic A Priori in Strawson's Kantianism' has not been printed before. It was presented at a conference on Strawson and Kant in Reading, England in 1999. Kant's question of how synthetic a priori judgements are possible was in part a question of how philosophical results with the distinctive status of those he reached in the *Critique of Pure Reason* could be reliably arrived at, and metaphysics set on the secure path of a science. The question depends on a clear distinction between "analytic" and "synthetic" judgements, and on the assumption that if something is known to be necessarily true then it is known a priori. Kant's answer was that transcendental idealism is the only possible explanation of our knowledge of such truths. This paper asks whether there is a parallel question about the more 'austere' Kantian project pursued by Strawson with no appeal to transcendental idealism or, apparently, to a priori knowledge.

I try to show that conclusions with a very special, distinctive status can be reached if necessary connections can be discovered between the possession of certain concepts or conceptual capacities and others, and if certain specified conceptual capacities can be shown to be required for the possibility of any thought or experience at all. The distinctive status of those conclusions can be described without making use of any "analytic-synthetic" distinction, and without supposing that we know them a priori. There remains the question of how we can know of the kinds of necessary connections that the project requires. I suggest that when certain traditional reasons for thinking that necessity is uniquely problematic are overcome or set aside, what remains might be no

¹³ (New York, 1999).

more than a special case of the general question of how we can know anything. Issues touched on here overlap with those in earlier papers in this volume on transcendental arguments and the Kantian project, and with discussion of the idea of 'invulnerability to unmasking' in the two immediately preceding papers.

CONTENTS

<i>Introduction</i>	ix
1. Scepticism and the Possibility of Knowledge (1984)	1
2. Transcendental Arguments (1968)	9
3. Doubts about the Legacy of Scepticism (1972)	26
4. Taking Scepticism Seriously (1977)	38
5. Reasonable Claims: Cavell and the Tradition (1980)	51
6. Transcendental Arguments and 'Epistemological Naturalism' (1977)	71
7. The Allure of Idealism (1984)	83
8. Understanding Human Knowledge in General (1989)	99
9. Epistemological Reflection on Knowledge of the External World (1996)	122
10. Scepticism, 'Externalism', and the Goal of Epistemology (1994)	139
11. Kantian Argument, Conceptual Capacities, and Invulnerability (1994)	155
12. Radical Interpretation and Philosophical Scepticism (1999)	177
13. The Goal of Transcendental Arguments (1999)	203
14. The Synthetic A Priori in Strawson's Kantianism (1999)	224
<i>Index</i>	245

Scepticism and the Possibility of Knowledge

Scepticism in recent and current philosophy represents a certain threat or challenge in the theory of knowledge. What is that threat? How serious is it? How, if at all, can it be met? What are the consequences if it cannot be met?

I obviously do not have time to go into all these questions, or into any of them thoroughly. I can only sketch a point of view in the hope of provoking some discussion.

The first question is clearly the place to start. I believe the true nature of the sceptical threat is still not properly understood, nor are the consequences of it not being met. That is one reason we have tended to give inadequate answers to the other questions. It is still widely felt that scepticism is not really worth taking seriously, so it hardly matters whether the challenge can be met or not. That kind of reaction seems to me to rest on a philosophical misconception.

Many would dismiss scepticism and defend not taking it seriously on the grounds that it is not a doctrine or theory any sensible person would contemplate adopting as the truth about our position in the world. It seems to them frivolous or perverse to concentrate on a view that is not even a conceivable candidate in the competition for the true or best theory as to how things are. I would grant—indeed insist—that philosophical scepticism is not something we should seriously consider adopting or accepting (whatever that means). But does that mean that it is silly to worry about scepticism? I think it does not. A line of thinking can be of deep significance and great importance in philosophy even if we never contemplate accepting a ‘theory’ that claims to express it.

One reason that is so is that philosophy thrives on paradox, absurdity, dilemma, and difficulty. There are often what look like good arguments for surprising or outrageous conclusions. Taking

the paradoxical reasonings seriously and re-examining the assumptions they rest on can be important and fruitful when there is no question at all of our ever contemplating adopting a ‘theory’ or doctrine embodying the absurd conclusion.

The point is clearest in the case of antinomies—explicit contradictions. We know we cannot believe the conclusion; it couldn’t possibly be true. To take *The Liar*, or Russell’s paradox seriously is not to hold open even the remote option of believing that someone who says he is lying speaks both truly and falsely, or that there is a set that both is and is not a member of itself. Such ‘theories’ would be worse than outrageous as things to believe, but that in no way diminishes the need to take seriously the reasoning that leads to them.

The same is true even when the conclusion of the paradoxical or surprising reasoning falls short of explicit contradiction. The Eleatic doctrine that nothing moves, for example, need not be in any remote sense a live intellectual option for us in order for us to be rightly challenged, overwhelmed, perhaps even stumped, by Zeno’s argument that Achilles can never overtake the Tortoise. The mere idea of something’s being true at a time can seem to generate the absurd result that there is never any real alternative to what happens, that things are fated to happen as they do. We can be impelled to investigate that line of reasoning without thinking that otherwise we would have to adopt the ‘theory’ that we have no control over what we do or what happens to us. Again, it seems undeniable that adding one more molecule to a table would not turn it into a non-table, any more than pulling one hair from a bushy head would make it bald. The discomfort I feel in the thought that an exactly similar step can be taken again, and again, does not show that I in any way consider accepting a ‘theory’ according to which there could be a table the size and shape of the earth, or that a bushy head and a bald head are the same sort of thing.

Those modern philosophers most closely connected to the sceptical tradition and most impressed by sceptical reasoning—Descartes, Hume, and Russell, for example—do not hold that believing the conclusions of that reasoning is a real option for us. The ancient sceptics themselves seem not to have accepted, or to have contemplated accepting or declaring the truth of any ‘theory’ either. They were highly anti-theoretical philosophers, and their strictures would have extended to any theoretical pronouncements