fundamentals of philosophy

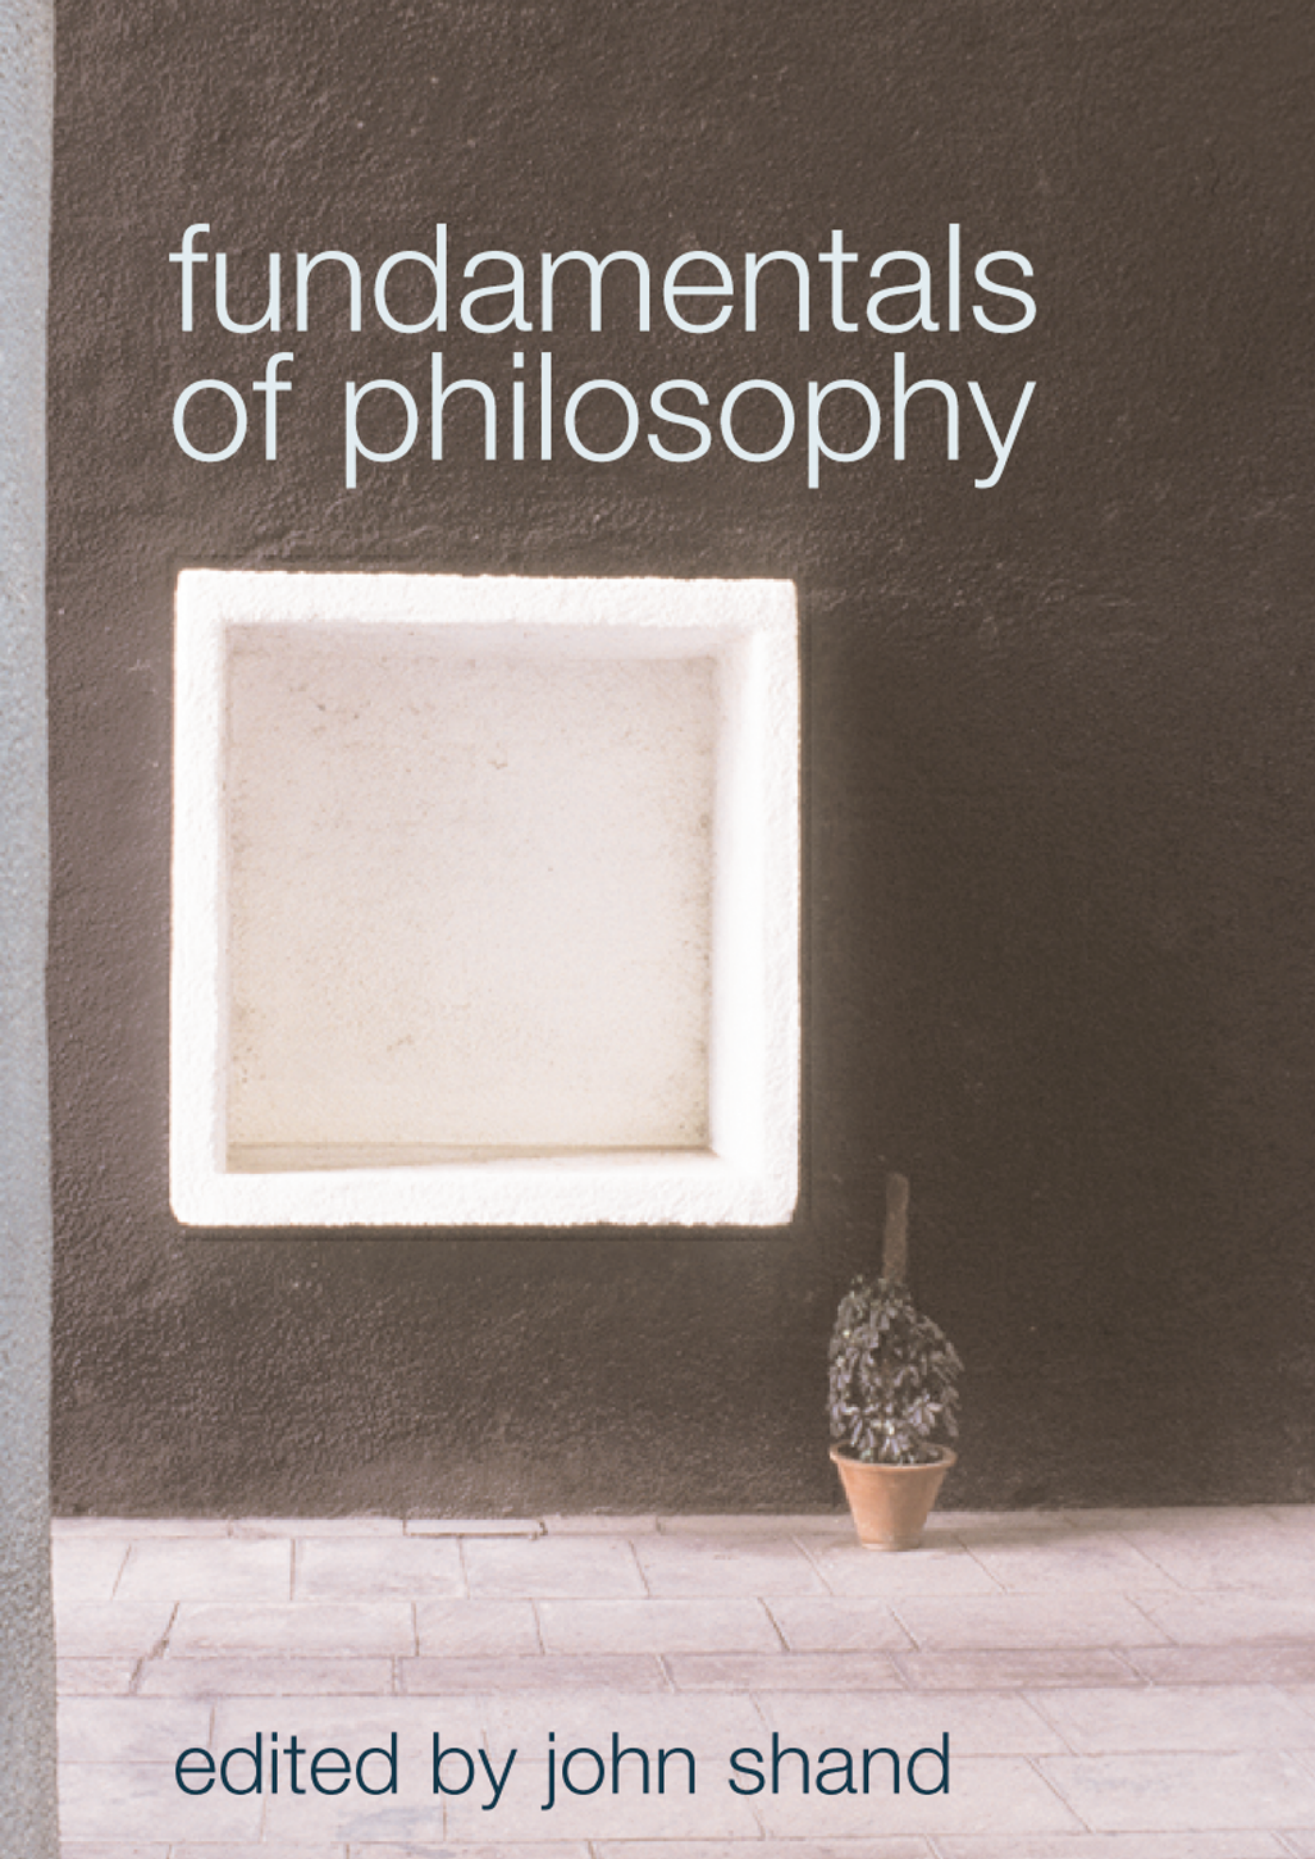

edited by john shand

FUNDAMENTALS OF PHILOSOPHY

Fundamentals of Philosophy is a comprehensive and accessible introduction to the major topics in philosophy and is designed to be used as a companion to any undergraduate philosophy course.

Based on the well-known series of the same name, this textbook brings together specially commissioned articles by leading philosophers. Each chapter provides an authoritative overview of topics commonly taught at undergraduate level, focusing on the major issues that typically arise when studying the subject. Discussions are up to date and written in an engaging manner so as to provide students with the core building-blocks of their degree course.

Helpful exercises are included at the end of each chapter, as well as bibliographies and annotated further reading sections.

Fundamentals of Philosophy is an ideal starting point for those coming to philosophy for the first time and will be a useful complement to the primary texts studied at undergraduate level. Ideally suited to novice philosophy students, it will also be of interest to those in related subjects across the humanities and social sciences.

John Shand is Associate Lecturer at The Open University. He is series editor of the *Fundamentals of Philosophy* series (Routledge), author of *Arguing Well* (Routledge 2000) and *Philosophy and Philosophers* (2002).

FUNDAMENTALS OF PHILOSOPHY

Edited by John Shand

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

First published 2003
by Routledge
11 New Fetter Lane, London EC4P 4EE

Simultaneously published in the USA and Canada
by Routledge
29 West 35th Street, New York, NY 10001

Routledge is an imprint of the Taylor and Francis Group

Typeset in Sabon and Frutiger by
Keystroke, Jacaranda Lodge, Wolverhampton
Printed and bound in Great Britain by
TJ International Ltd, Padstow, Cornwall

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication Data
Fundamentals of philosophy / [edited by] John Shand.
p. cm.
Includes bibliographical references and index.
1. Philosophy–Textbooks. I. Shand, John, 1956–

BD31.F86 2003
100–dc21
2002044529

ISBN 0–415–22709–7 (hbk)
ISBN 0–415–22710–0 (pbk)

TO SARAH WITH LOVE

CONTENTS

CONTRIBUTORS

Piers Benn is Lecturer in Medical Ethics at Imperial College, London. Previously he lectured in Philosophy at Leeds and St Andrews. He is author of *Ethics* (Routledge 1998) and much of his writing is in applied ethics. His interests within philosophy range widely and he has published in popular as well as scholarly outlets.

Alexander Bird is Lecturer in Philosophy at the University of Edinburgh, and has held visiting positions at the Universities of Caen, Siena and Cambridge, and at Dartmouth College. He was educated at Westminster School and the universities of Oxford, Munich and Cambridge. He is the author of *Philosophy of Science* (Routledge 1998) and *Thomas Kuhn* (2000).

Stephen Burwood was born in London in 1959 and now teaches philosophy at the University of Hull. He is co-author of *Philosophy of Mind* (Routledge 1998) and is currently working on books on the self and competing conceptions of human embodiment.

Richard Francks was born in Leicestershire in 1950. He has taught English in Spain, Japan and Scotland, and Philosophy at The Open University and University of York. He is currently Director of Undergraduate Studies in Philosophy at Leeds, and not very good at cricket.

Simon Glendinning is Lecturer in Philosophy at the University of Reading. He is the author of *On Being with Others: Heidegger–Derrida–Wittgenstein* (1998); and the editor of *The Edinburgh Encyclopedia of Continental Philosophy* (1999) and *Arguing with Derrida* (2001). He has published articles on perception, animal life, and the end of philosophy.

Alan Goldman is William R. Kenan, Jr Professor of Humanities and Professor of Philosophy, College of William and Mary, Virginia. He is the author of six books, including *Empirical Knowledge*, *Moral Knowledge*, *Aesthetic Value*, and *Practical Rules*. He loves to play tennis.

Michael Jubien is Professor of Philosophy at the University of California, Davis. He studied mathematics and philosophy at Dartmouth College and The Rockefeller University, and is the author of *Ontology, Modality, and the Fallacy of Reference* (1993) and *Contemporary Metaphysics* (1997).

Dudley Knowles is Senior Lecturer in Philosophy at the University of Glasgow. He is the author of *Political Philosophy* (Routledge 2001) and *Hegel and 'The Philosophy of Right'* (Routledge 2002).

Colin Lyas is Senior Lecturer in Philosophy at the University of Lancaster and the author of *Aesthetics* (Routledge 1997).

Alexander Miller is Senior Lecturer in Philosophy at Macquarie University, Australia. He is the author of *Philosophy of Language* (Routledge 1998) and co-editor (with Crispin Wright) of *Rule-Following and Meaning* (2002). He is currently working on a book on Michael Dummett and a book on metaethics.

Dermot Moran is Professor of Philosophy at University College Dublin and author of *John Scottus Eriugena: A Study of Idealism in the Middle Ages* (1989) and many articles on medieval philosophy. He is author of *Introduction to Phenomenology* (Routledge 2001) and Editor of *International Journal of Philosophical Studies*.

Greg Restall was born in Brisbane in 1969. He studied mathematics and philosophy at the University of Queensland. He is an Associate Professor in Philosophy at the University of Melbourne, and is the author of *An Introduction to Substructural Logics* (2000) and *Logic* (Routledge 2003).

John Shand studied philosophy at the University of Manchester and King's College, Cambridge. He is an Associate Lecturer in Philosophy at The Open University and is the author of *Philosophy and Philosophers: An Introduction to Western Philosophy* (2003) and *Arguing Well* (Routledge 2000).

Suzanne Stern-Gillet is currently Professor of Philosophy at the Bolton Institute. She holds degrees in Philosophy and in Classics from the Universities of Liège (Belgium) and Manchester (UK). She is the author of *Aristotle's Philosophy of Friendship* (1995). She is currently working on a history of ancient philosophy, *From Thales to Iamblichus* (Routledge).

W. Jay Wood received his PhD in philosophy from the University of Notre Dame, and has taught at Wheaton College since 1982. He is the author of *Epistemology: Becoming Intellectually Virtuous* (1998) and other articles on epistemology and the philosophy of religion.

PREFACE

Putting together this book has been interesting and rewarding. Most of all I should like to thank the authors of the individual chapters for their hard work, cooperation, care, and indeed the insightful skill with which they wrote each of their contributions. Individually and collectively I think they have done a great service for philosophy by presenting core aspects of the subject in such an accessible, thoughtful and well-written way. It should open doors for many. Their accomplishment is nothing like as easy as it may look.

I should like to thank Siobhan Pattinson at Routledge for being so easy to work with as she helped ferry the book through. I should also like to express as always my appreciation to my wife Judith whose eagle eyes and intelligence improved the text. Personal thanks go to my young daughter Sarah for being a delight. The book is dedicated to her.

John Shand
Manchester
2003

INTRODUCTION

John Shand

THE AIM OF THIS BOOK

This book is an accessible stimulating gateway to the central areas of philosophy. The chapters are carefully arranged to begin with what are usually regarded as the core areas of the subject and then extend out to other important subjects of less generality, not, one should emphasise, of less importance. The prime purpose of the chapters is not to give comprehensive coverage of each subject, but rather to open the door on the subject for the reader and encourage thought about all the ideas within. Someone once said to me that studying philosophy had 'opened doors'; if this book does that, it will have succeeded.

WHAT IS PHILOSOPHY?

Philosophy is a great intellectual adventure while at the same time what it discusses is one of the most important things we can do with our lives.

There is a standing joke among many professional philosophers that involves one of them being cornered at a party by someone, and on hearing that he is a philosopher, being asked, 'Well, what is philosophy then?' The joke in fact reflects the unease of many philosophers and the discomforting awareness of not being able to come out with a straight clear answer. Many philosophers resort to the list-method of answering, saying that it's about 'fundamental issues' such as 'truth', 'What can be known?', 'What is the nature of a good and bad action?', 'What is the nature of mind and how is it related to body?' The other way of dealing with the question is somewhat evasive and involves saying as little as possible, something like: 'well, the best way to understand what philosophy is is to do it'. Both these answers, neither of which is without truth, are likely to leave the original questioners rightly bewildered, dissatisfied and quickly heading off to get another drink – much to the relief of the philosopher.

I think it is incumbent on professional philosophers to tackle this question head-on. After all we do get paid. My immediate answer to the question, requiring a little refinement later on, is:

Philosophy is what happens when you start thinking for yourself.

A bit more may then be added. Once one frees oneself from the habits of received belief, those that one just happens to have acquired even about basic issues, and really starts to think about what one ought to believe, judged by reason (argument) and evidence, then one has started to do philosophy. The 'tradition' of relying instead on 'authorities' and 'holy text' is the usual state of affairs rather than the exception in history – for many it still is the natural way of going on. Moreover, thinking for oneself is not something easily taken on by mere momentary act of will, but rather something to be strengthened like a muscle through good mental habits. Philosophy is a way of life to be built up over years; philosophical thinking is a cast of mind that becomes part of a person's very nature.

Philosophy is often thought to be an unnecessary impractical luxury. A sort of futile, at best entertaining, addition to life after one has dealt with the practicalities. But this is a mistake.

Far from being unnecessary philosophy is unavoidable just as soon as people cease taking their received beliefs for granted and instead start thinking them through for themselves. The glory of philosophy – and certainly one of the original attractions for many drawn to it – is that nothing is out of bounds, not even the value of reason, or indeed (although this may seem paradoxical) the status of philosophy itself. No holds are barred. Only something like argument and debate without boundaries seems to be a constant. It's a wonderful freedom. Either one is a slave to the beliefs one happens to have acquired through the contingent circumstances of how and where one is brought up, or one is to some degree a philosopher. Philosophy is the bastion of free thought and of the exploration of ideas above all others.

What of the charge of it being an impractical luxury? This is a mistake too. This is because beliefs lead to actions (and inaction), and badly thought out ideas lead often to terrible actions. Our responsibility for what we believe, and what we leave ourselves open to being capable of believing, cannot be divorced from our responsibility for our actions. Ideas that in untesting times can even seem benign, in extreme circumstances lead to awful actions.

Philosophy sometimes addresses the question as to how one should live. It can be argued that keeping a philosophical stance itself is exactly how one should live – anything else is gullible slavery. Of course it's a matter of degree, but for the most part it's one-way to freedom of thought: after having it no-one wants slavery again.

It would be wrong to think that philosophy leaves one constantly in a state of vague doubt. One accepts one's beliefs on the basis of the best arguments. But one leaves the door ajar for further argument. In fact it is those who take on their beliefs as acts of will and faith that stand on a precarious escarpment from which they can be knocked by circumstance with the painful consequences of disappointment, emptiness and loss. The result may be catastrophic

because they fall, if they do, from such a great height and from a place they thought absolutely secure. After which, what? Philosophy does not set its hopes so high. It's prepared also to live bravely with that. Even if one changes one's beliefs in the light of new arguments, one can tell oneself that last time one held a view one did one's best to really get to the bottom of the matter. Philosophy breeds neither empty doubt nor an unattainable certainty.

As a way of life philosophy and philosophical thinking do not promise happiness, but they do, I think, enhance what is best in human beings. Philosophy embodies that which is noblest in our species.

THE HOUSE THAT PHILOSOPHERS BUILT

Philosophy is rather like a house built on stilts in a river. In the house one can do all sorts of things – construct things, move things about – but one is always aware that the structure is supported by pillars that are driven into something potentially and often actually shifting. Philosophy goes down repeatedly to see how things are going on around the foot of the pillars and indeed inspects the pillars themselves. Things may need changing down there. For philosophers this is not just the nature of philosophy, it is the true intellectual condition of mankind. It is philosophy that pays that condition close attention and take it seriously. This rather than ignoring it or solving it glibly.

THE AREAS OF PHILOSOPHY

The range of philosophy is large and basically unified. However, to clarify issues and build up expertise it divides its energies into areas of specialisation. There are two characteristics of these areas. One is those that have a subject matter that seems to underpin most of what we think and do. The others underpin more particular concerns we have. The areas feed on one another and are interrelated. Philosophy is not built like other subjects from unquestioned basic foundation upwards. It does not consist of easy bits we can all assume out of which the more complex bits are made. There is, as they say, no shallow end in philosophy – when one starts all the deep issues come into play straightaway.

As far as the subjects of the chapters in this book are concerned, philosophy can be divided into three groups.

Group I
- Logic
- Epistemology
- Metaphysics

Group II
- Ethics

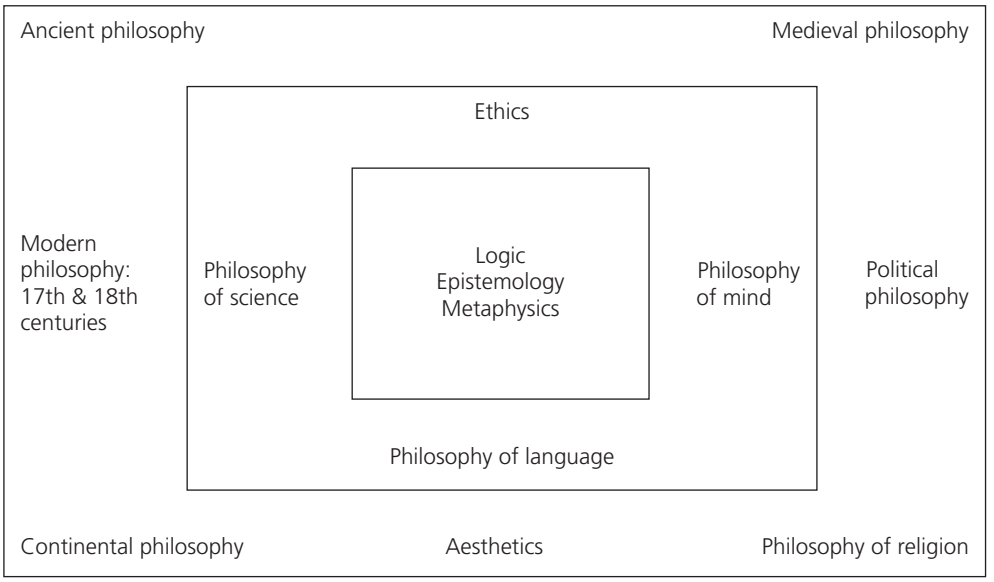

Figure I.1 Philosophy: the fundamentals

Philosophy of mind
Philosophy of language
Philosophy of science

Group III

Ancient philosophy
Medieval philosophy
Modern philosophy: seventeenth and eighteenth centuries
Political philosophy
Aesthetics
Continental philosophy
Philosophy of religion

The relation between these subdivisions of philosophy is not one of difficulty but one of generality, with lesser generality as one moves away from the centre. This is not to say the outer subjects are less important. Rather it is that those subjects in Group I underpin the problems considered in Group II, and have consequences for the conclusions one reaches in Group II – Group II finds itself referring back to Group I constantly. The subjects in Group III do not raise new fundamental philosophical considerations that are not dealt with in Groups I and II, but rather apply all the problems encountered in Groups I and II to specific areas. Here are some examples: Metaphysics may be concerned with what sort of entities fundamentally exist; aesthetics is concerned with thinking about in what way works of art exist; what sort of entities are they? Ethics examines what it is to say that we ought to do

something, for something to be right or wrong; political philosophy studies the right way to organise society, if it should be organised at all.

The historical chapters listed here, such as Ancient Philosophy and Medieval Philosophy of course deal with all the central problems of philosophy as they are treated by a period or school of thinkers.

THE PROBLEMS OF PHILOSOPHY

Here is a list of some of the most commonly addressed and most basic philosophical problems. Do not worry too much about how one would address these questions as a philosopher – just look through them and consider how you might answer them in an immediate intuitive way – my guess is that you will soon find yourself in deeper water than you may expect, philosophical water in fact. In fact do not feel pressure to find *an* answer, but think of various ways one may answer the questions and what reasons one has for those answers being correct. The answers, or merely how one should even start to approach them, are a good deal less straightforward than one may suppose.

What is the nature of philosophy?
Are there philosophical problems?
What is the correct method for solving philosophical problems?
When are inferences sound?
What is the nature of rationality?
What is truth?
What is it to know something?
What are we perceiving when we claim to be perceiving the world?
Can we know the external world exists?
What is reality?
What is it for something to exist?
What sorts of things exist?
What is a cause?
What is it for something to be morally good?
What is the good life?
Can ethical judgements be justified?
What is the nature of mind?
What is consciousness?
What is the self?
What is it for expressions in a language to have meaning?
What is it to understand the meaning of a word?
Can induction be justified?

What is a scientific law?
How should society best be organised?
What justifies the power of the state?
What are human rights?
What is a work of art?
Can we justify the evaluations we make of works of art?
What determines the meaning of a work of art?
What is it to justify the existence of God?
What is the nature of God?
How ought we to live?

TIMELESSNESS

It is not uncontentious to say that philosophical problems are timeless. To some it looks like an excuse for examining problems which in fact can have no answer because there is something wrong in considering them as 'problems' in the first place. However, the subject of philosophy certainly acts as if philosophical problems are timeless. Certain topics may be more of a central concern at a particular time, but that is mainly a function of fashion. The central topics and questions come round again and again. Rarely is it the case that a matter considered by philosophy is wholly dismissed, or the way it was once treated regarded as valueless. Quite the contrary. Philosophers find themselves going back to philosophers of the past at the least to use their ideas on certain topics as starting points, but often much more than that. A book that considers the nature of justice will naturally find itself looking to see what Plato had to say. The problems of induction and causation normally involve discussing Hume in depth. The starting point for considering the nature of mind is often Descartes.

It's far from clear that progress is made in philosophy as in some other subjects. In this sense philosophy is quite unlike science – a chemist would rarely find any value in checking to see what another chemist said about something a hundred years ago.

So one may wonder what is the point of philosophy in this case if it does not definitively solve problems. As suggested already philosophical problems arise when we start to think deeply about our most fundamental beliefs. When we do so we often find that we neither fully understand the content of those beliefs, nor have any clear justification for holding them. For a certain kind of mind this is perplexing and the problems will not go away through the acceptance of glib answers or in response to a dismissive frame of mind. We may not be able to present final solutions, nevertheless we can come to a conclusion that is a result of the best thinking on a certain matter.

I would conclude that philosophical problems are timeless by virtue of their profundity, generality and, as a consequence of that, the uncertainty surrounding the very methods by

which they may be best approached. The result is that the problems do not die, nor do the ways of attempting to solve them or at least deal with them.

One thing is pretty certain: the issue of whether philosophical problems are timeless is itself a philosophical problem.

BEYOND THE FACTUAL

Philosophy is not usually concerned with gathering facts. That can be left to other disciplines such as science, or history, or psychology or anthropology. The reason for this is twofold. First, philosophy usually deals with matters that have to be assumed in gathering the facts – questions about truth and the knowability of reality, for example. Any attempt to solve the philosophical problems by reference to the facts is therefore highly likely to be question begging. We cannot for example refer to the evidence gathered through perception about the world to solve the philosophical problem of what can be known, if anything, about the world through perception. Second, the facts are usually insufficient to deal with the philosophical problem. This is particularly obvious in ethics. It is generally argued that no reference to what people are like and what they actually do can answer the question of what people ought to do. This is not to say the facts are ignored, just that the facts are insufficient to allow us to come to conclusions about the matters with which philosophy deals.

THE SUBJECTS OF PHILOSOPHY

This section gives a thumbnail sketch of the subjects of philosophy discussed in this book. The book is not exhaustive of philosophy, but it can fairly be said that all the core areas are covered here.

Epistemology

The subject here is the nature of knowledge, and given that nature, what it can be truly said that we can know, as opposed to just having beliefs and opinions about. Can we counter views of sceptics who would claim that strictly speaking we cannot know as much as we claim to, or indeed anything at all?

Metaphysics

What sorts of things ultimately exist and how do they connect to each other and how things appear to us? Are all the things that appear to us real, or are they derived from something

more fundamental? And what do we say about the existence of things that do not in the usual sense 'exist' but to which we nevertheless refer, such as unicorns or numbers.

Logic

This is concerned with the nature and identification of good inferences: those circumstances in which one statement is said to follow from another. It seeks to understand and classify the cases where statements, if true, justify to whatever degree the truth of other statements.

Ethics

This is concerned with values (normative as opposed to factual matters) with respect to human actions. What is it for something we do to be counted good or bad? What is it to say we ought to do or not do something? It is not enough to talk of what we do, we need to address what we should do and what saying this means.

Ancient philosophy

This is the study of the philosophers of the Greek and Roman world. The usual concentration is on Greek philosophy from *c*.624BC, marking the birth of the Presocratic Thales, to 322BC as the death of Aristotle. The most important figures are undoubtedly Plato and Aristotle. Often this period is extended to include the Roman world. The significance of thought in the ancient world cannot be overestimated. Here we find almost everything, developed to varying degrees, that characterises the Western outlook. Indeed it represents a watershed in human history, where for the first time reason alone is applied across the board to the solving of the deepest problems rather than appeal to mere authority or an idea's longevity.

Medieval philosophy

This covers, we should note, the study of philosophers over a vast time of around one thousand years, extending from St Augustine of Hippo (AD354–430) and William of Ockham (*c*.1285–1349), and continuing beyond until at least the Renaissance. The connecting thread is the rise and dominance of Christianity which permeates the philosophy done during this period. The other most significant link throughout the period is the interpretation and adaptation of Aristotle's metaphysics.

Modern philosophy: the seventeenth and eighteenth centuries

It may seem strange to call philosophy done in the seventeenth and eighteenth centuries 'modern philosophy'. It indicates a period of astonishing fecundity in philosophical thought and a new way of doing philosophy that was a significant break from what had gone before. Moreover many of the ways that philosophy is presently done still derive from thought in this period. The central figures are Descartes, Spinoza, Leibniz, Locke, Berkeley, Hume.

Philosophy of mind

What kind of entity are we referring to when we talk about the 'mind'? How does talk of the mind relate to talk of what we normally call our bodies? Are the mind and the body one or is the mind non-physical? How can conscious awareness and understanding whereby we refer to things arise from inert matter? What do we mean by, and can it justify, saying that someone is the same person throughout his life?

Philosophy of language

What is it for an expression, spoken or written, to have meaning and the capacity to refer to things? What constitutes a person's understanding the meaning of a word, at which point they know how it ought to be used correctly?

Philosophy of science

What defines a law of nature? How does it differ from other claims about the world? How if at all are scientific theories justified by evidence? How can we know that our laws of nature describe features of the world that will persist next time we examine it?

Political philosophy

How ought society to be organised? What justifies the existence of the state that can rightly usurp power from people? How should the state be controlled? What justifies private property, if anything? How do people acquire rights that cannot be transgressed apart from exceptional circumstances, if at all?

Philosophy of arts

Can what a work of art is be defined? What do we mean when we say some work has a certain aesthetic quality, such as beauty? What determines the meaning of a work of art? What, if anything, justifies our valuing works of art differently?

Philosophy of religion

How good are the arguments justifying the existence of God? Are arguments for the existence of God required, or is faith enough? What is the nature of God and how does that relate to the sort of creatures we are?

Continental philosophy

It is controversial to claim that the group of philosophers often brought together under this title can be done so coherently, and the chapter here deals mainly with this matter. Negatively the title may indicate a divergence of methods and philosophical concerns between philosophers in Continental Europe and English-speaking philosophers in Britain, North America, New Zealand and Australia. Positively there is perhaps a thread that runs from the philosopher Immanuel Kant (1724–1804) to the present with such thinkers as Jacques Derrida, and this can be seen as various ways of responding to the philosophical outlook of trans-cendental idealism. The recent philosophers here are often marked by the most fundamental questioning of the nature, and indeed existence, of philosophy itself.

THE FUTURE OF PHILOSOPHY

Philosophy will go on just as long as some people hold the view that thinking things through for themselves is important. It is hard to say what philosophical concerns will be the centre of people's attention in the future. But it looks as if there will always be someone trying to struggle with the deepest questions and unwilling to take on trust the answers that happen to be around.

1

EPISTEMOLOGY

Alan Goldman

INTRODUCTION

Epistemology is concerned with the nature, scope and structure of knowledge. As epistemologists, we want to know first what knowledge is, and we want our analysis of the concept to guide us in determining the scope of knowledge, in deciding how much knowledge we have. In determining the scope of knowledge, the epistemologist will attempt to answer sceptical challenges to the sources that are usually assumed to produce knowledge, sources such as perception, memory, testimony of others, and various kinds of reasoning. The sceptic will question whether the ways things appear in perception or memory, for example, constitute good evidence for the ways we take them really to be, and whether various kinds of reasoning produce true beliefs from their data. In attempting to provide answers to the sceptic's questions, we should be able to reveal not only the scope of knowledge, but also its structure. We will see whether knowledge has a web-like structure, in which beliefs reflect their status as knowledge by connecting with other beliefs in a set, or whether knowledge has foundations, special beliefs which attain their status independent of connections with other beliefs, and with which other beliefs must cohere.

We will take up each of these topics in turn, beginning with the analysis of knowledge. A particular approach to epistemology will be endorsed and briefly defended here. But we will also note difficulties for this approach, and the major alternatives will be considered and criticised as well. At the end, the reader should have both a feel for the general field and an idea of how one theory might be developed and defended against alternatives and against sceptical objections.

THE NATURE OF KNOWLEDGE

Knowledge is the goal of belief. It is what belief aims to be, or, more precisely, what we aim at in believing. There may be some types of belief, for example religious, for which knowledge is seen to be impossible and belief itself sufficient (in its effects). But knowledge is always to be preferred to mere belief where it is possible; it is, other things being equal, the ideal form of belief. An analysis of knowledge must reflect this fact. What must knowledge be like to function properly as our cognitive goal? We want our beliefs to be true, but we want more of them as well. We want not just truth, but secure truth, truth that will be resistant to pressures against its acquisition or retention. If the truth of a belief is not firm in this way, then changes in the world or in the subject that are unrelated to the fact believed will likely alter the belief and render the resulting changed belief false. Beliefs acquired similarly in the future will be likely to be false as well, and we will not be able to tell as easily whether they are true or false. Thus, we want our beliefs to be *non-accidentally* true, so that they will not be subject to such whims of fortune. We want to remove luck from the acquisition and retention of true belief, just as we want to remove moral luck from the actions of agents. Acting in a morally right way by accident (when rightness is no part of an agent's intention) does not produce faith in or praise for the agent; similarly, believing the truth by accident does not produce faith in one's cognitive abilities or positive grades for the achievement.

It is relatively uncontroversial among epistemologists that knowledge involves true belief, and most would accept the claim that the truth of a belief must be non-accidental if it is to amount to knowledge. But controversy will arise over how to understand this crucial requirement. Certain kinds of luck or accident can enter into the acquisition of knowledge, while other kinds must be ruled out. And the absence of accident in certain senses will not guarantee that a true belief counts as knowledge. Regarding the first point, I might be just lucky to run into a friend of mine in Paris and hence to know he is there; but despite the fact that my running into him was accidental, I do know he is there. Regarding the second point, a perverse epistemologist might deliberately trick me into believing the truth when my belief is based on the wrong reasons or is unconnected in the right way with the fact I believe. He might trick me into believing that *someone* in my department owns a Ford by convincing me that he himself does, when he but not I know that only another member of my department owns a Ford. There is a sense here in which it is non-accidental that I believe a true proposition, but I still lack knowledge.

These two examples can help us to begin to sharpen the sense in which knowledge must be non-accidental. In the first example, *given the context* in which I acquire the belief, that in which I see my friend, it is non-accidental that I believe he is there.

And in the second example, while my perverse colleague deliberately sets up the context in which I acquire my belief, given that context, my belief that someone in my department owns a Ford is only accidentally true. Thus, we can say that a belief must be non-accidentally true in the context in which it is acquired in order to count as knowledge. Beyond this point, however, it will remain a matter of great controversy how to interpret the requirement of being non-accidental.

Ordinarily, when our beliefs are only accidentally true, they result from lucky guesses. A venerable but suspect tradition in epistemology seeks to eliminate lucky guesses by requiring that believers be *justified* in their beliefs. This concept of justification has its origin and natural home in ethics. In morally judging persons by their actions, we demand that they be justified in acting as they do and that they act as they do because of this justification. Similarly, in judging persons by their beliefs, we may demand that they be justified in believing as they do and not achieve truth by lucky guesses. But it remains questionable whether justification is either necessary for knowledge or sufficient when added to true belief.

Before attempting to answer these questions, it is necessary to clarify the concept of justification to which appeal is being made. While we often talk in non-philosophical contexts of agents being justified in acting as they do, 'justification' is a technical term of art in epistemology, rarely used in reference to beliefs outside the context of philosophical analysis and debate. And it is a concept about which epistemologists themselves have conflicting intuitions. The analogy with ethics suggests that justification is a matter of fulfilling one's obligations as these can be determined from an internal perspective, from the subject's own point of view. Moral agents are justified when acting in a subjectively right way given the information available to them. Similarly, believers might be said to be justified when they have fulfilled their epistemic obligations given the evidence available to them, for example, when they have critically assessed the available evidence.

But there are many problems with this internalist conception, based as it is on what subjects should believe from their own perspective. First, the analogy with ethics may be out of place, since we do not have the same degree of control over the acquisition of beliefs as we do over our actions. If we cannot help believing as we do, then talk of epistemic obligations is suspect, although we can still exercise control over the degree to which we gather evidence, seek to be impartial, and so on. Second, it must be clarified to what degree the justification for one's beliefs must be available and able to be articulated from one's own perspective. On the most extreme view, in order to be justified in a belief, one must be aware not only of the evidence for it, but of the justifying relation in which that evidence stands to the belief. But, given the motivation for this view, it seems that one's belief in that justifying relation must itself be justified, and that one's belief that it is justified must be justified, and so on. Even if that regress were to end somehow, it seems clear that

ordinary subjects are not aware of such complex sets of judgements and so could never fulfil this requirement.

A weaker internalism regarding justification would require only that evidence for one's beliefs be in principle recoverable from one's internal states. One question here is whether subjects must be able to articulate their evidence as such. This requirement would disallow the perceptual knowledge of children, for example, who cannot articulate the ways things appear to them as ways of appearing. Even without this requirement, there seem to be clear counterexamples to internalist concepts of justification as necessary for knowledge. (The internalist distinguishes between a person's being justified and there being some justification not in the person's possession, the latter being irrelevant.) A clairvoyant who could reliably foretell the future, an idiot savant who knows mathematical truths without knowing how he knows them, or a person with perfect pitch who can identify tones with almost perfect accuracy have beliefs that count as knowledge without having any apparent justification for those beliefs. Certainly they are not justified in their beliefs until they notice their repeated successes, but they have knowledge from the beginning. In more mundane cases, we all have knowledge when completely unaware of its source, when that source or the evidence for our beliefs is completely unrecoverable. I know that Columbus sailed in 1492, and I assume that I learned this from some elementary school teacher, but who that teacher was, or what her evidence for the date was, is, I also assume, completely unrecoverable by me. More generally, knowledge from the testimony of others requires neither that one knows the evidence for the proposition transmitted nor even that one have evidence of the reliability of those providing the testimony (what it does require will be discussed below).

Thus justification in the sense in which the concept is derived from ethics is not necessary for knowledge. It is more commonly accepted since Edmund Gettier's famous article that justification, when added to true belief, is not sufficient for knowledge (Gettier 1963). Many examples like the one cited earlier about the owner of the Ford exemplify justified, true belief that is not knowledge. They show that a person can be accidentally right in a belief that is not simply a lucky guess. Other examples that show the same thing include beliefs about the outcomes of lotteries, which falsify many otherwise plausible analyses of knowledge, and beliefs of those in sceptical worlds (also to be discussed later), such as brains in vats programmed to have experiences and beliefs, or victims of deceiving demons. A brain in a vat programmed to have the beliefs it does can occasionally be programmed to have a true belief grounded in its seeming perceptual experience about an object outside the vat, but that justified, true belief will not be knowledge. I can justifiably and truly believe that my ticket in this week's Florida lottery will not win, but I do not know it is a loser until another ticket is drawn.

Thus, justification in any intuitive sense is neither necessary nor sufficient, when added to true belief, for knowledge. Some philosophers have sought to beef up the notion so as to make it the sufficient additional condition for knowledge by requiring that justification be 'undefeated'. One's justification is said to be defeated when it depends on a false proposition, such as the proposition that my colleague owns a Ford in that earlier example (Lehrer 2000, p. 20). There are two fatal flaws in this position. One is that it takes justification to be necessary for knowledge, and we have seen that it is not. The other is that it cannot distinguish between examples in which one's claims to knowledge are threatened by misleading evidence one does not possess. Suppose in the Ford example that my colleague does own the car and gives me good evidence that he does, but that he has an enemy who spreads the false rumour that he is a pathological liar. If that enemy is also in my department and the chances were great that I would have heard his false rumour, then my claim to knowledge will be defeated. It will then be a matter of luck that, given the context of being in my department, I did not hear his testimony and so believe as I do. If, by contrast, my colleague's enemy is in some distant city, his attacks will be irrelevant to my knowledge. No way of unpacking the notion of 'depending on a false proposition' will distinguish correctly between these cases.

That knowledge is the goal of belief indicates yet again that the epistemologist's notion of justification is largely irrelevant. In a court of law, for example, where it is of utmost importance whether witnesses know that to which they testify, jurors must assess whether the evidence they present connects in the right way with the facts they allege. Jurors want to know whether the best explanation for the evidence presented by witnesses appeals to the facts as they represent them, or whether the explanation offered by the opposing attorney is just as plausible. They do not care whether the witnesses are justified in their beliefs, only again whether their beliefs hook up in the right way with the facts. Sceptical worlds also reveal that justification can be worthless, hence not a goal of belief, as firm truth is. One such sceptical world mentioned earlier is that of brains in vats programmed to have all the perceptual experiences that they have. Brains in vats are normally justified in their beliefs on the basis of such experience, but such justification is unrelated to truth and knowledge, not the sort of thing we seek for itself.

If justification is irrelevant to knowledge, we may wonder at the epistemologist's obsession with the notion. There are several explanations. One is that, while ordinary knowers need not be able to defend their claims to knowledge in order to have knowledge, it is one of the epistemologist's tasks in showing the scope of knowledge to defend it against sceptical challenges. In doing so, she will be justifying or showing the justification for various types of beliefs. Some epistemologists might confuse themselves for ordinary knowers, in thinking that ordinary knowers too

must justify their beliefs in the face of sceptical challenge. Another explanation for all the attention to this concept is the practice of some epistemologists of calling whatever must be added to true belief to produce knowledge 'justification'. This practice might be excused by the fact, noted earlier, that the term in epistemology is in any case a stipulative term of art. But, if this term refers only to an external relation between a belief and the fact believed, or to a process of acquiring belief that is outside the subject's awareness, then it will lose its normative force and any connection with the ethical concept of justification from which it supposedly derived. It will then lead only to confusion to refer to such additional conditions for knowledge as justification. Externalists might retain the concept by requiring only that there be some justification that perhaps no one has, but again this invites confusion in seeming to be, but not being, a normative concept.

Externalist accounts of knowledge do not require that the condition beyond true belief must be accessible to the subject. They take that condition to be either general reliability in the process that produces the belief or some connection between the particular belief and the fact to which it refers. We may consider reliabilism first (Goldman 1986). Can reliabilists capture the requirement that the truth of a belief that counts as knowledge must be non-accidental? If so, they must take it that when subjects use reliable processes, processes that produce a high proportion of true beliefs, it will not be accidental that they arrive at the truth. But reliabilists who require only general reliability in belief-forming processes would be mistaken in assuming this to be universally true. If a process is not 100 per cent reliable, then, even when it generates a true belief on a particular occasion, it may be only accidental or lucky that the belief is true. I may be not very reliable at identifying breeds of dogs by sight, except for golden retrievers, which I am generally reliable at identifying. But I may be not very good at identifying golden retrievers when they have a particular mark that I wrongly believe to indicate a different breed. I may then fail to notice that mark on a particular dog that I therefore identify correctly, albeit only by luck or accidentally.

This example reveals several problems, some insurmountable, in the account of knowledge that takes it to be true belief produced by a generally reliable process. First, at what level of generality should we describe the process that generates this true belief (Feldman and Conee 1998)? Intuitively, we take processes that generate beliefs to be those such as seeing middle-sized objects in daylight, inductively inferring on the basis of various kinds of samples, and so on. But the former, although used to generate the belief in this example, seems completely irrelevant to evaluating the belief. Whether I am generally reliable in identifying things that I see in daylight has little if anything to do with whether I acquire knowledge that this dog is a golden retriever. Given our judgement that I do not have such knowledge in this example, that I am only lucky to believe truly that this dog is a retriever, we

can choose as the relevant process the unreliable one of identifying dogs with the marks that tend to mislead me. This is a quite specific process, but, of course there is the yet more specific one of identifying retrievers with such marks without noticing the marks, which turns out to be reliable and to give the wrong answer in this case. By choosing the former process as the relevant one, we can make the reliability account appear to capture the example. In fact, we can probably do the same for any example, given that every instance of belief acquisition instantiates many different processes at different levels of specificity. But such ad hoc adjustments do nothing to support the reliabilist account. We need independent reason or intuition of the correct specification of the relevant process in particular cases, if not in general, in order for the account to be informative or illuminating.

Our example reveals the pressure to specify the relevant process more and more narrowly. But at the same time it shows that however narrowly we specify it in particular cases, as long as we leave some generality in its description, there will remain room for only accidentally true belief being produced by the process. This indicates clearly that what is important in evaluating a true belief as a claim to knowledge is not the reliability of any generalisable process that produced it, but the particular connection between that very belief and the fact believed. One might try to save the language of reliability by claiming that a process must be reliable in the particular conditions in which it operates on a particular occasion, but once more any looseness or generality at all will allow room for the type of accident that defeats a claim to knowledge. One might also demand perfect reliability, but then one would have to explain why we allow beliefs produced by perception and induction, both fallible processes, to count as knowledge. We do so when these methods connect particular beliefs to their referents in the proper way.

If the example discussed does not suffice, we can appeal to the lottery example once more to show the weakness of reliabilism as an analysis of knowledge. If one inductively infers that one's ticket will not win, we can make the reliability of this inductive process as high as we like short of 100 per cent by increasing the number of tickets. But one still does not know one's ticket will not win until another ticket is drawn. If one did know this, one would never buy a ticket. The problem is not the lack of high reliability or truth, but the lack of the proper connection between the drawing of another ticket and one's belief. Once one receives a report of the drawing of another ticket, then one knows, if the report is based on some witnessing of the event. One then knows even if the probability of error in such reports is the same as the initial probability that one's ticket would be drawn. Once more, it is not the probability or reliability of the process that counts, but the actual connection between belief and fact. Mere statistical inference about the future does not suffice in itself for knowledge, no matter how reliable, but one can have knowledge of the future if it is based on evidence that connects in the proper way with the future

events believed to be coming. If, for example, one discovers that the lottery is fixed, then one can come to know that one's ticket will not win.

Given the failure of reliabilism to rule out accidentality in true belief, one might again explain the popularity of the theory among epistemologists as the result of their confusing themselves with ordinary knowers. While the general reliability of belief-forming processes is irrelevant to the knowledge of ordinary knowers in particular cases, the epistemologist, who is interested in defending types of beliefs against sceptical challenges, does try to show that certain sources such as perception or induction are generally truth generating or reliable. The project of seeking to improve our epistemic practices must also seek to establish first which practices reliably produce truth. But the analysis of knowledge must focus instead on finding the right connection between belief and truth or fact.

The first attempt to specify the connection between belief and fact that renders the belief knowledge was the causal theory (Goldman 1967). This account holds that a true belief must be causally related to its referent in order to count as knowledge. The account captures such examples as the lottery, which, given the failure of so many other theories to do so, indicates that it's on the right track, but it proves to be too narrow. One can have knowledge of universal and mathematical propositions, for example, but universal and mathematical facts or truths do not seem to cause anything. It is also too weak in failing to rule out cases in which there is the usual causal connection between a perceptual belief, for example, and an object to which it refers, but in which the subject could not distinguish this object from relevant alternatives (Goldman, 1967). I might see a criminal commit a crime but not know that he is the culprit because I do not know that his twin brother, also in the vicinity, did not commit the act.

This sort of case is handled by what is perhaps the best-known attempt to specify the crucial connection between belief and fact, the counterfactual account (Nozick 1981, ch. 3). This holds that one knows a fact if one would not believe it if it were not the case, and if other changes in circumstances would leave one believing it. In terms of possible worlds, one knows a proposition if and only if in the closest possible world in which the proposition is false, one does not believe it, and in close possible worlds in which it remains true, one does believe it. (We measure closeness of possible worlds by how similar they are to the actual world.) This account captures the examples so far considered, but unlike the causal account that proves to be too weak, this one is too strong, disallowing genuine knowledge claims. Many of the most mundane facts that I know do not obtain only in very distant possible worlds. In worlds so unlike this one there may be no telling what I would believe there. And it does not matter what I would believe in such worlds. I know that my son is not a knight of the round table and that it is not ninety degrees below zero outside. There is no telling what I might believe if those propositions were false,

but this affects not the least my knowledge claims. Thus the first counterfactual condition is too strong. If the second requires retention of belief in all close possible worlds, then it is too strong also. An aging philosopher can still know this truth, although there are close worlds in which he cannot follow the argument that establishes it.

Thus, we need a specification of the relevant connection between fact and belief that counts as knowledge that requires that the connection hold only across close possible worlds, and only across most of them. Such a connection is what we would expect also from a naturalistic perspective, from the fact that the capacity for achieving knowledge is a likely product of natural selection. The capacity to achieve firm true belief is one that would be selected in slowly changing environments, so that true belief would be firm in situations close to actual, but not in distant possible worlds. An analysis that meets this condition and captures all of the examples so far discussed requires that the fact believed (the truth of the belief) enter into the best explanation for the belief's being held. The concept of explanation here can itself be explicated in terms of possible worlds. In this account A explains B if A raises the antecedent probability of B (given other factors, it will raise the probability to 1, where there is no indeterminism or chance involved), and there is no third factor that screens out this relation, that fully accounts for the difference. The last clause is required because evidence, for example, raises the probability of that for which it is evidence, but this relation is screened out by whatever explains both the evidence and that for which it is evidence. In intuitive terms, A explains B if, given A, we can see why B was to be expected. In terms of possible worlds, A explains B if the ratio of close worlds in which B is the case is higher in those in which A is the case than it is in the entire set of close worlds.

Let us review some of the examples that were problematic for the rival accounts of knowledge. When there is misleading evidence I am just lucky not to have noticed, then what explains my belief is the fact that I have not noticed this evidence. My believing the dog in that example to be a golden retriever is explained by my not having noticed the misleading mark. My believing as I do is not made significantly more probable by the fact believed, given all the close possible worlds in which I am aware of the misleading evidence. In the lottery example, the inductive evidence on the basis of which I believe that my ticket will lose does not explain its losing, since the probabilistic connection between that evidence and its losing is screened out by what does explain the latter, the drawing of another ticket. That drawing explains my losing but not my prior belief, which remains explanatorily unconnected to the fact to which it refers.

In regard to problems for the causal theory, the truth of universal propositions helps to explain our belief in them, or it helps to explain the inductive evidence that explains our beliefs. In the case in which I cannot distinguish the cause of my belief

from relevant alternatives in the vicinity, the explanation for my belief lies in the broader context and not in the specific cause, just as we do not explain the outbreak of a war by citing only the specific event that triggered it, when any number of equally likely events would have done so in the broader context of latent hostility. In such cases the specific cause does not significantly raise the probability of its effect across close worlds in which alternative causes are also present. To be able to rule out relevant alternatives in claiming knowledge is to be able to rule out alternative explanations for the evidence one has.

In regard to the cases that were problematic for the counterfactual account, what explains the fact that my son is not a knight of the round table, the fact that he lives in the present time and is a tennis player attending Yale, also explains my belief that he is not a Medieval knight. What explains the fact that it is not ninety degrees below zero outside, namely the fact that it is ninety above zero, also explains my belief that it is not subfreezing. Finally, in the aging philosopher example, his belief that the counterfactual analysis is too strong is connected with the evidence that it is too strong in many, although not all, close possible worlds.

In many of these examples, appeal is made to explanatory chains. It suffices for knowledge if what explains my true belief also explains or is explained by the fact to which the belief refers, as long as a certain constraint on these chains is met. Each link in such chains must make later ones more probable. This constraint defeats some purported counter-examples that will not be considered here (see Goldman 1988, pp. 46–50), but its relevance is also clear in the case of knowledge from testimony mentioned earlier in discussing the issue of justification. A person may be justified in believing the testimony of another without any evidence of the other's expertise or sincerity, as long as there is no evidence that the testimony is likely to be false. Testimony can create its own justification, just as perception can, whether or not the testifier is herself justified in believing her own testimony. But this again simply contrasts justified true belief with knowledge, since one cannot transmit knowledge one does not have. Knowledge from testimony requires an explanatory chain in which the truth of the testimonial evidence enters ultimately into the best explanation for its being given and believed. If I am completely gullible and believe absolutely anything I hear, then I do not gain knowledge from testimony, just as if I see everything as red, then I do not know a red object when I see one. But the last two points imply a third, that a completely gullible person anywhere in the testimonial chain destroys knowledge in the later links. For each link, the fact that the belief was more likely because true must make its transmission more likely to be believed at later links, the constraint mentioned earlier. This does not prevent children from gaining or transmitting testimonial knowledge, since they tend to believe their parents, for example, more than they believe their peers (Schmitt 1999, p. 372).

This completes our brief account of the nature of knowledge. As we shall now see, it will prove to be highly suggestive for the task of determining the scope and structure of knowledge.

THE SCOPE OF KNOWLEDGE

In the first section we utilised intuitions about when knowledge is had in order to derive an account of its nature. This might seem to beg the question against the sceptic by guaranteeing that our criteria for knowledge are met for the most part. But we are not in fact assuming that scepticism is false. This is because we allow that purported cases of knowledge to which we appeal in analysing its nature can turn out under sceptical attack to be not genuine. Indeed, sceptics themselves must adopt the same procedure of analysis – first using ordinary intuitions to derive criteria – and then give us reasons for doubting that these criteria are really satisfied. Otherwise, they risk basing their sceptical attacks on an assumed analysis that is too demanding and so out of touch with our concept of knowledge. In that case we would not need to take them seriously. Here we will take them seriously by dismissing all claims that their doubts are necessarily misplaced.

Scepticism challenges us because our beliefs about the properties of real things transcend the evidence we have for those beliefs. Such evidence consists in the ways those things appear to us. But objective properties of real objects are what they are independently of our beliefs about them and the ways they appear to us. Thus, our beliefs are underdetermined by our evidence. There will be alternative possible explanations for all the evidence we have. If everything can seem exactly as it does to us and yet nothing be as we believe it to be, then how can we know that it is as we believe it to be? If all our evidence is compatible with alternative explanations of it, then how can we rule out all but one, indeed any, of those explanations? If knowledge is belief best explained by its truth, then how can we know we have knowledge when different explanations are compatible with all the evidence we have for our beliefs? How can we know that the explanatory chains end in the facts as we take them to be?

Sceptics dramatize this problem by presenting us with alternative scenarios or sceptical worlds in which everything appears to us as it does now, i.e. our experience remains exactly the same, and yet nothing in the world is as we take it to be. Descartes challenged us to show that we are not dreaming all that we currently experience, or that we are not being deceived by some powerful demon who causes us to have the experiences we do. Or, to take the contemporary version, suppose that we are brains in vats programmed by super scientists or computers to have exactly the experiences we do. We believe this scenario to be possible, since we believe that our

experiences are immediately caused by happenings (neuronal firings) in our brains. How, then, could we know that it is not actual? If we cannot know that it is not actual, that we are not brains in vats, then it seems we cannot know that we have bodies surrounded by middle-sized objects with any of the properties we take them to have. Thus, the sceptic concludes, since his scenario is possible, we do not have any knowledge of real objects.

A recent trend among epistemologists who battle this sceptic is to grant that we do not know that we are not brains in vats, but then to argue that we do nevertheless retain ordinary knowledge of such things as the properties of middle-sized objects. This response is held to refute the brunt of the sceptic's argument while simultaneously showing the source of its plausibility, a goal now endorsed by anti-sceptical epistemologists as well. Both claims – that the sceptic's first premise must be granted, but his conclusion denied – are suggested by the counterfactual analysis of knowledge described in the first section. According to this account, we do not know we are not brains in vats because in the possible world in which we are, we do not believe we are (since everything appears as now). But this sceptic's world is assumed to be very distant from the actual world. It therefore does not affect the fact that in the closest possible worlds in which particular propositions now believed about ordinary objects are false, we do not believe them. Hence ordinary knowledge is retained despite the truth of the sceptic's premise and resultant plausibility of his argument.

There are nevertheless three crushing problems with this response to the sceptic. First, its dependence on the counterfactual account is itself problematic, since we saw earlier that this account is too strong, ruling out legitimate claims to knowledge. And, by its own lights, the response relies on the account in just the case in which it is most dubious, where our evaluation of a knowledge claim takes us to a distant possible world. Only in this way is the sceptic's premise endorsed. Second, the analysis implies that the sceptic's second, conditional premise (that if we do not know we are not brains in vats, then we do not know we are surrounded by middle-sized objects) is false, and it clearly seems to be true. If we do not know that we are not surrounded by a vat's clear liquid, how can we know that we are surrounded instead by tables and chairs? Third, and perhaps most important, in denying the sceptic's conclusion, the proponent of the counterfactual analysis simply assumes that the sceptic's world is a very distant one. But if, as the account admits, we cannot know that the sceptic's world is not actual, how could we possibly know that it is distant from the actual world? As an answer to the sceptic, this response simply begs the question. Even if accepted, it shows only that knowledge is possible, not that it is actual.

A yet more contemporary response, contextualism, builds upon the previous one by agreeing that we do not know that we do not occupy sceptical worlds even

though we do retain knowledge in ordinary contexts. Contextualists differ from counterfactualists in holding that the sceptic's second premise is true also, as is their conclusion in the context of their argument (DeRose 1995). In the context of sceptical doubt, sceptical worlds such as that of the vatted brains become relevant alternatives that cannot be ruled out. And if they cannot be ruled out, then we do not retain knowledge of mundane facts with which they compete. But in ordinary contexts free from sceptical doubts, the sceptic's distant worlds are irrelevant, and our beliefs must vary with the presence or absence of the facts to which they refer only in close possible worlds. When judges of knowledge claims raise sceptical doubts, they raise the standards for evaluating beliefs; and when beliefs have to be sensitive to facts in the distant worlds of the sceptic, they cannot pass this unusual test. Recognition of such varying standards in different contexts of evaluation allows the contextualist to say that the sceptic's argument is sound but irrelevant to our ordinary knowledge claims. How better to show the plausibility of the sceptical position while defending the ordinary knower?

Despite this attraction, contextualism fares no better in the end than counterfactualism. It may improve on the latter by allowing that if we cannot know we are not brains in vats, then we cannot know that we are surrounded by tables and chairs. Once more, in the context of the doubt that the antecedent expresses, this conditional is held to be true. But the other problems facing counterfactualism plague contextualism too, and additional ones as well. First, the position still relies on evaluating beliefs in what it holds to be distant possible worlds, and we have seen that this demand is too strong in any context. Some of the most mundane truths that are easiest to know are false only in very distant possible worlds, where there is no telling what we would believe. The counterfactual account makes these the most difficult facts to know. Second, in defending ordinary knowledge the position once more simply assumes that the sceptic's worlds are distant while admitting that we cannot know they are not actual. This does not satisfy the demand to answer the sceptic by showing that we have knowledge. Third, there is the implausibility of the claim that we can destroy knowledge we have by simply thinking of sceptical alternatives. One unwelcome implication of this claim is that philosophers, who frequently entertain sceptical hypotheses, have so much less knowledge than their more fortunate, if more naive, counterparts in the real world. While ignorance may be bliss in some contexts, pursuing a profession that so systematically substitutes it for knowledge is probably not what young philosophy undergraduates have in mind. Contextualists who may be content to know so much less than anyone else nevertheless had better not advertise their position.

Can we then defend knowledge by rejecting the sceptic's first premise? Can we claim to know that we are not brains in vats? Can we show that the evidence we have from experience is evidence for the world as we take it to be, and not for

the sceptic's worlds? How could we know or show this, when experience itself cannot differentiate between the world as we take it to be and the phenomenal world of a brain in a vat? One older answer favoured by some epistemologists is that we know this a priori, that its defence does not require any inductive argument since it *could* not be false. It is held that we *must* know it a priori precisely because experience in itself cannot distinguish these worlds and so cannot be the source of this knowledge. Defenders of this tradition give different but related explanations of how we have this a priori knowledge, of how we can know that the way something appears, for example, is *necessarily* evidence for how we take it to be.

Many of the arguments here begin from an account of how we learn to understand the terms in our language, how we learn to use them correctly or to interpret their use by others (Hamlyn 1970, ch. 3). If we learn to pick out tables, for example, by how they appear to us, how they look and feel, then it must be correct that whatever looks and feels to us continuously in those ways must be tables, or at least that such looks and feels are necessarily evidence for the presence of tables. In the language game in which we apply the term 'table' to tables, such ways of appearing are criteria for the correct use of this term. We therefore cannot all be mistaken in this use based on these experiences any more than we could all be mistaken in the way we play chess. Tables are whatever we call tables based on correct application of the term, and correct application is determined by the agreed upon criteria, in this case certain ways of appearing. Thus, these ways of appearing are necessarily evidence for tables and for the properties that define them to be tables. We can neither use the term correctly without accepting these criteria nor interpret its use by others without typically ascribing true beliefs about tables to them. Likewise, of course, for other middle-sized objects and their properties.

Is this argument sound? What it really establishes is only the way we must initially conceive of things. Once we develop the notion of objects whose properties are independent of our experiences and beliefs, once we develop theories of how these properties cause our experiences, and once we see that our experiences can mislead us as to the real properties that cause them, the possibility of wholesale error becomes intelligible. In fact this possibility is entailed by the notion of independence that defines the concept of realism about objects and their properties. That real properties are independent of the ways they appear and the beliefs they cause means that these appearances and beliefs can be misleading and false. Once we recognise the possibility of wholesale error on our part, we need not necessarily ascribe mostly true beliefs to others (although we will ordinarily do so). If, for example, we were to see some brains in vats and understood their situation, we would not ascribe to them mostly true beliefs about the objects around them. To interpret the language of others, we need to explain their utterances, but truth of the beliefs expressed need not necessarily enter into the majority of these explanations. Nor will we explain

the brains' utterances as true of phenomenal objects instead of false of real ones, since they will have the same concept of real objects as we do and will intend to refer to them and their properties. What we and the brains take to be evidence for the objective properties of real objects cannot dictate what those properties are. Our shared concept of chess may determine the nature of that game, but this is what distinguishes games from reality.

Thus, premises about how we learn and interpret our language do not show that the evidence we have for our beliefs about real objects must necessarily be evidence for their objective properties as we take them to be. Is this notion of criteria as necessary evidence short of entailment even coherent? If we do not do away with real objects and their properties by reducing them to experiences or appearances, can the latter nevertheless necessarily be evidence for the former? To say that appearances are necessarily evidence for real properties is to say that they are evidence in all possible worlds. But in a world of brains in vats in which the brains were informed or knew of their own situation, their experiences would not be evidence of objects as we take them to be. This would be true of any sceptical world believed to be such by its victims. Such worlds are possible. The sceptic's descriptions of them do not involve logical contradictions. We could even grant that we could not *all* be brains in vats, but that would leave open the possibility that any one of us is and could possibly be informed of this by our programmers. Thus, there is no necessary connection between experiential evidence and the real properties of objects.

Do we have *any* a priori knowledge of reality, as opposed to that which reflects only definitions of terms, including logical connectives and operators? Is there any a priori insight into the necessary structure of reality, knowledge of what is real but not contingent, that needs no inductive confirmation? Well-worn examples that seem to express such knowledge include the claims that nothing can be red and green all over and that, if A is taller than B, and B is taller than C, then A is taller than C (Bonjour 1998, pp. 100–3). It turns out, however, that such examples express lack of experience or imagination, instead of a priori insight into necessary truth. When I was much younger and clothing styles were much different, I owned an iridescent raincoat that looked red and green (as well as tan) all over. Whether the effect was achieved with discrete red and green threads is irrelevant here, since on most accounts of colour, whatever looks a certain colour to normal observers in normal conditions is that colour.

I must admit that I have never seen three people or objects that disconfirm the example regarding tallness, but imagine the following possible world. Imagine a world with only three rods and one observer (you). When rod A is visually compared to B, A appears to be taller; when B is visually paired with C, B appears taller; but when C is compared to A, C appears taller. When you attempt to place all three in the same visual field, you cannot take them all in by sight. One always disappears

out of the visual field. This world is possible: in fact, it is not so unlike the world of quantum mechanics, with its indeterminacy and measurement problem. And the world described seems not to be a world in which the supposed necessary truth about tallness holds true. Hence it is not a necessary truth (unless we simply stipulate a definition of 'tallness' in which only a transitive relation will bear that term).

Even if we are not so sceptical of necessary truths about reality, the relation between appearing and being a real property of objects seems not to be necessary. If we are to answer sceptical challenges to our beliefs about that relation, we therefore require an inductive argument, broadly construed. We have looked at the implications for the sceptical argument of the counterfactual analysis of knowledge. It is time to use the analysis we endorsed, that knowledge is belief best explained (in significant part) by truth. Here the question becomes whether the ways we are appeared to in experience are better explained by objects and their properties as we take them to be than by sceptical hypotheses such as the programming of brains in vats.

It will be objected to this approach right off that we cannot construe our beliefs about physical objects and their properties as a theory that best explains the patterns within appearances because we cannot even formulate such patterns without appeal to physical space and objects. We learn the language of appearing only after learning to hedge our claims about real properties, which we immediately and naturally ascribe without inference. More important, it is claimed, we know what experiences we have had and will have only by reference to locations in physical space among physical objects. Thus, to even formulate the supposed data we must presuppose the 'theory' that is supposed to explain the data, which 'theory' therefore does not require this explanatory justification.

The former point about learning, however, is again irrelevant to the question of justification, or demonstrating knowledge. Even if we must conceive of physical properties first, that does not show that there are such properties or that they are as we take them to be. Our evidence for our beliefs about them still consists in the ways they appear, although we don't initially conceive it that way. The latter point about formulating the data is first of all debatable. While we cannot translate statements about physical properties into statements about appearings, it is not so clear that we could not learn to weaken all claims about objects to claims about appearances, although doing so would be very cumbersome and awkward. Instead of talking of seeing an unsupported object fall, we could talk of a visual experience as of an unsupported object followed by a visual experience as of an object falling. But even if such reductions would not be universally possible for a skillful language user, this still only indicates a conceptual necessity, the way we must express ourselves. It still does not imply that the ways we conceive objects are the ways they are independent of our conceptions, and it still does not preclude the attempt to show that real properties as we conceive them best explain the ways they appear to us.

That physical objects and their properties do provide the best explanation we have for the ways we are appeared to seems easily established. Appeal to real physical objects first of all explains more deeply than do explanations for particular appearances in terms of regularities within experience itself. The former can explain the regularities themselves that are otherwise taken to be ultimate. And explanations in terms of physical objects and their properties are also superior to sceptical explanatory hypotheses such as programmes for brains in vats or deceptive Cartesian demons. The agreement of the different sense modalities on the qualities and dimensions of objects, for example, strongly suggests a realist explanation, and of course we also use the realist model successfully to predict future experiences. The physical realist picture explains both how physical objects interact with each other and, at least in part, how they cause experiences (by reflecting light, emitting sound waves, and so on). No such predictions or explanations are forthcoming from the brain in vat or evil demon hypotheses, which are therefore ad hoc and useless additions. The only way they have any explanatory power is by being parasitic on physical object explanations: the demon or programmer of brains must make it seem as if we are surrounded by interacting physical objects (Vogel 1990). Not only does this add nothing to the explanations available without these additions; it also raises natural but unanswerable questions, such as why the programmers or demon would deceive us in this way.

If the commonsense and scientific explanations for our experience are superior to the sceptical hypotheses, does this show that we know that we are not brains in vats? When a hypothesis is put forth only to explain certain data, and when a superior explanation is later offered, we have some reason to disbelieve the former hypothesis and the entities it posits solely for explanatory purposes. When the demon theory of disease was replaced by the germ theory, rational people ceased to believe in disease-causing demons. In that case, however, there was additional evidence against the existence of demons – the fact that no one has ever seen one (except perhaps in a highly irrational frame of mind). The best explanation for the latter fact is that there are none. In the brain in the vats case, we would not see ourselves as such if we were, and so we lack that additional evidence. Unlike the usual case of knowledge of negative existential propositions (propositions that certain things do not exist), we have no evidence against the existence of the brains (while also lacking evidence for their existence). It is this, we can claim, that explains the plausibility of the sceptical argument, while leaving it open to us to defend our knowledge claims against that argument via inference to the best explanation.

Two problems remain for this inference as an answer to the sceptic. The more tractable one is that an inference to the best available explanation at a given time is not necessarily an inference to the best explanation *tout court*, so that even if we can accept the latter as producing the true explanation, we cannot so easily

accept the former as doing so. We need additional reason to think that the available explanations exhaust the field of plausible ones. Here that demand seems easily met, however. There does not seem to be the remotest possibility of an explanation for our experience being developed that rivals that which appeals to physical objects and the scientific theories of physical reality.

The far more difficult problem is that an inference to the best explanation defeats the sceptic only if we can defend the principle that underlies such inference against sceptical challenge. We must be able to show that what appears to us to be a best explanation is likely to be true. In general, the most difficult part of any anti-sceptical epistemology will be to defend the fundamental principles of reasoning or the basic sources through which we seek knowledge. We have seen that we can defend perception as a source of knowledge through an inference to the best explanation. Similar arguments will be available for memory, knowledge of other minds, and ordinary induction (for example, the best explanation for many coherent memory impressions will appeal to earlier veridical perception or testimony; the best explanation for an observed ratio of instances in a class may well be a deeper or universal regularity, and so on). But this leaves the formidable task of defending the principle of inference on the basis of which these other sources of knowledge can be defended.

In indicating the nature of this defence, we must note first that we cannot know a priori that such fundamental cognitive principles as inference to the best explanation lead to truth. This is not a necessary truth, since sometimes such inference fails; and whether it is generally reliable depends on the type of brains we have and on our relation to our environment. Nor can it be necessary that inference to the best explanation is likely to be reliable. It is hard to imagine how it could be necessary that anything is likely (contrast Bonjour 1998, p. 214). In terms of the possible worlds model of probability, this would mean that in all galaxies or groups of possible worlds, the number of worlds in which the proposition in question is true exceeds the number in which it is false. But how could that be necessary and how could we know that it was: what is to keep the worlds in which the proposition is false from clustering?

In the absence of an a priori defence of inference to the best explanation and other fundamental cognitive principles and practices, we would require an inductive or empirical argument. One promising approach might be to argue that for creatures with such limited physical capacities and instincts as humans, basic cognitive capacities would have been naturally selected for their capacity to provide information or truth necessary to survival. It is plausible that humans would have had to infer correctly the proximity of predators from tracks or predator noises, for example, in order to survive, since they could not outrun them or otherwise protect themselves. There are three major obstacles that a generalised version of such an

argument would have to overcome, however. First, there is the question whether, in the environment in which our brains evolved or were selected for their cognitive capacities, truth was in general the key determinant of fitness or utility. Second, a natural next question is whether an affirmative answer to the first one suggests that inferences to explanations far removed from the environment in which natural selection took place continue to be truth preserving. And third, there is the problem of circularity in the argument.

We can only very briefly indicate answers to the first two questions here, since we will pay more attention to the third, which will introduce the topic of the structure of knowledge. In regard to the first question, it must be admitted that systematic distortion can be utile and even fitness enhancing, as when the exaggeration of colour contrasts enables us to see object boundaries more easily (not to mention that the perception of colour in itself may be a systematic distortion of objective reality). But such systematic distortions seem to occur precisely to enable us to obtain more vital veridical information about the environment, for example about the locations of various objects. It remains hard to see how creatures like us could have survived if our basic cognitive capacities did not generate true information about our environments. We must know the means necessary to our ends, including survival, as well as the consequences of our actions, and all this is a matter of inferences as well as perception and memory.

As for inferences far removed from contexts in which survival is at stake, we can ask generally whether there is reason to suspect that a cognitive capacity that produces truth in one context will cease doing so in another. More specifically, we can separate everyday inferences that can be later verified directly in perception from those which produce the more remote products of scientific theory. The former inferences are demonstrated to be true if perception is accepted as a source of knowledge, but the charge of circularity will be raised again, to be addressed shortly. Scientific inferences not only take explanations to deeper levels, but seek to correct for distorting subjective inputs into earlier, commonsense explanations. This gives us more instead of less reason to believe in the truth of such explanations, although there might be less reason here to believe that the best of available explanations at a given time is the best overall, hence the true explanation.

The problem of circularity can be pressed at every stage of this suggested defence of inference to the best explanation as a basic cognitive principle. First, it was suggested that brains with such cognitive capacities to provide truths necessary to survival were probably products of natural selection. But theory of natural selection, indeed appeal to the physical environment, is itself legitimated via inference to the best explanation. The attempt to legitimate a principle by appeal to the products of its own use is circular. Second, we noted that many such inferences can be confirmed by later perceptions, giving direct evidence of their truth-preserving

nature. But once more perception itself is certified as a source of knowledge only via inference to the best explanation for the ways things perceptually appear, again using the very cognitive principle that perception is supposed to help legitimate.

Addressing this problem of circularity takes us to the question of the structure of knowledge, to which we now turn.

THE STRUCTURE OF KNOWLEDGE

In responding to the charge of circularity, we should note first that the argument defending inference to the best explanation does not beg the question against the sceptic in the way that it would if one of its premises simply stated that the principle of inference is sound or truth preserving. Any principle, indeed any conclusion, could be 'defended' in this way, since any premise implies itself. Here the principle of inference was used, not mentioned, in its own defence. Such use does not trivially result in self-support, as the sceptic might suggest. We do not normally dream, for example, that dreaming is a reliable source of knowledge or hallucinate that hallucination is reliable. In the case of inference to the best explanation, it is not a foregone conclusion that it will be self-supporting, and the argument from evolution is controversial even if sound. But the principle, when applied with critical care, does iterate, in that it is more explanatorily coherent to believe that it leads to truth. It is doubtful whether there are any obviously unsound principles that iterate it this way or cohere with other fundamental sources of knowledge.

This coherence with other principles or sources is also significant even if circular. Crystal balls might predict their own reliability, but their predictions fail to be regularly confirmed by perceptions. This kind of mutual support adds significantly to the self-support of the principle, since it strongly suggests a common cause in truth. Why should percepts confirm earlier inferences if they do not reflect the same facts believed? The more seemingly independent facts that an explanation unifies, the more reason there is for accepting the explanation. Of course, a non-sceptical answer to the question just raised simply applies an inference to the best explanation again. But the near ubiquitous nature of the principle does not more firmly support a sceptical attitude toward it, although it does make more clear that any defence of it will have to be circular. While we have not relied heavily on necessary truths, it is worth pointing out here that it is a necessary truth that basic or fundamental epistemic principles can be supported only by themselves or their products and other fundamental principles (otherwise they would not be fundamental). We cannot escape the totality of our cognitive resources to verify that they lead to truth. Such confirmation can come only from within the circle of these resources. In the case of inference to the best explanation, which is involved in the legitimisation of

perception, memory, and simple induction, and in turn supported by them, the circle is certainly not narrow. Can we demand more of the epistemologist than the demonstration of such broad coherence in the set of our epistemic beliefs?

Circular reasoning is certainly not always vicious. We assess the skill of a tennis player, for example, by noting his execution of many skilful shots. But we judge a shot to be skilful and not merely lucky because it is made by a skilful player. The same shot made by a beginner would be just lucky. The reasoning is circular but sound and informative nonetheless (compare Sosa and Van Cleve 2001). That inference to the best explanation is supported by its own products makes for a circle, but not a vicious or self-defeating one. It is true that if the principle of inference did not lead to truth, then its being more explanatorily coherent to believe that it does would not seem to lend it much useful support. But in the absence of knowledge of the antecedent, the coherence of the principle with its products and with other basic epistemic sources can be seen to give us reason to accept it, just as inconsistency or incoherence with other principles would give us reason to reject a principle.

Sceptics, however, will have a strong rejoinder to this response to the charge of circularity. They will point out that perfect coherence in a set of beliefs, including epistemic beliefs, even together with truth, does not suffice for knowledge. Brains in vats may have coherent sets of beliefs, and when some of those beliefs about objects outside the vats also happen to be true, this does not give them knowledge. A coherent set of beliefs describes some possible world, but not necessarily the actual one. If we took the entire set of a person's beliefs and transferred them to another person in a different set of circumstances, the set would be equally coherent, but would contain much less knowledge (Sosa 1991, p. 203). It is clear once more from these examples that, in order to constitute knowledge, a coherent set of beliefs must be anchored to the actual world or surrounding environment in the right way.

Perceptual input in itself is insufficient to provide the required anchor. Such input could be provided by a deceiving demon or by the programmers of the brains in vats, and a complete madman could weirdly process perceptual input so as to make it cohere with his other mad beliefs. Such input and such processing would not help in the acquisition of knowledge even if it occasionally resulted in true belief. The perceptual experience must be of the right kind, and it must result in the right kind of belief to be of use in acquiring or demonstrating knowledge. But we cannot simply check this experience itself or the causal chains by which it produces beliefs in order to confirm that it and they are of the right kinds. The causal chains are largely inaccessible, and the experience must be conceptualised correctly in order to be of epistemic use. Experience itself need not be conceptualised (or it could not be the source of concepts), but until it is, it can play only an inaccessible causal role. To be of use in demonstrating knowledge, experiential input must be conceptualised in such a way that the best explanations for the beliefs it produces appeal to the

truth of those beliefs. It will be clear that the madman's beliefs fail this test from the fact that they fail to cohere with the beliefs of others, despite their internal coherence.

In our own case, in order to demonstrate knowledge we must defend our beliefs against sceptical challenge. Normally we defend a belief by citing evidence for it that is part of an explanatory chain leading to the fact believed. We do so for simple perceptual beliefs as well. For example, I defend my belief that there is a red object before me by noting that I am appeared to redly (but not seemingly by red beams of light). The claim that there is indeed a red appearance that explains my belief in the red object can be defended not only by inferring it as the explanation for that belief, but also directly by appeal to the belief that I am appeared to in that way (which is not to imply that I normally infer the belief about the object from the belief about the appearance). But here I arrive at a point in the demonstration where I no longer defend my belief by appeal to evidence or an explanatory chain. The best explanation for my belief that I am appeared to redly is simply that I am so appeared to, that 'red' is the correct concept or term in my vocabulary to apply to this experience, that I apply this term consistently to this type of experience. The latter is all that is required for the truth of my belief, and the only evidence for its being the case is my having formed the belief itself.

These beliefs about certain appearances form the foundations for the demonstration of knowledge in two related senses. First, they are shown to constitute knowledge without appeal to evidence or coherence with other beliefs. Second, they make up a set of most certain beliefs which with others must cohere, picking out one set of coherent beliefs as true of the actual world of the subject. Once this anchor is in place, it is doubtful that there are equally coherent but incompatible sets of beliefs with equal explanatory coherence, and the main objection to demonstrating knowledge by showing coherence in a set of beliefs falls away. In clarifying the sense in which these beliefs are foundations, several caveats must also be mentioned.

First, the relevant beliefs about appearances are not infallible, in contrast to their characterization by traditional foundationalism. In order to pick out patterns in experience from which objective properties can be inferred as causes in the demonstration of knowledge, they must refer to properties that are instantiated on different occasions and in different objects. While 'appears red' (as opposed to 'appears to be red') refers to a phenomenal property, a property picked out by what it is like to experience it, it is not defined ostensively on each occasion as referring only to whatever property is present in the visual field. A belief employing the latter ostensive concept might be infallible, but it would not be of any epistemic use. Beliefs in reinstantiated or reinstantiable properties are always fallible, since the concepts they employ can be misapplied or applied inconsistently. But for certain properties

that are naturally salient in experience, the best explanation for beliefs about how they appear will appeal only to the truth of these beliefs, as indicated earlier. The second caveat is that this will be true *only* for beliefs about how these naturally salient properties appear, for beliefs about red but not about C-sharp. The best explanation for my belief that I am appeared to C-sharply appeals to my musical training, and not simply to the fact that I am so appeared to. Third, it is not necessary in order to serve the function of foundations that the defence of every other belief trace a line back to one or more of these beliefs about appearances. Knowledge can have a web-like instead of linear structure as long as it is anchored to the world at key points.

CONCLUSION

At the end of the day, sceptics and non-sceptics can both stand their ground. Non-sceptics will hold that a set of coherent beliefs anchored to the world via suitable foundations, generated by a set of epistemic principles that are supported by their products and by each other, such that the best explanations for the beliefs being held appeal to their truth, constitutes the best indication we can have of knowledge of the real world. Non-sceptics will point out that nothing said earlier precludes the possibility that we are brains in vats or systematically deceived about the nature of reality. It remains possible that our fundamental epistemic principles and sources provide only a distorted view of the world independent of them. And the only argument to the contrary remains circular. The only surefire way to avoid these sceptical possibilities is to give up this notion of an independent world or the idea that we can intelligibly aim at knowledge of its properties. But if explanatory coherence among beliefs includes depth of explanations, and if appeal to real properties explains more deeply than explanations limited to the domain of experience, then this is not a viable epistemological option.

Absent this guarantee, we certainly have reason from within our world view for thinking that our epistemic principles and practices are sound and provide us with knowledge of the real world. But such reason must remain relative to those principles that determine what counts as a reason and so cannot convince the sceptic who will reject any principle that lacks independent support. Such fundamental and seemingly irreconcilable disagreement may be disappointing, but it keeps philosophy alive.