

OXFORD

THE
EXISTENCE
OF GOD

SECOND EDITION

RICHARD
SWINBURNE

THE EXISTENCE OF GOD

This page intentionally left blank

The Existence of God

RICHARD SWINBURNE

Second Edition

CLARENDON PRESS · OXFORD

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford ox2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Bangkok Buenos Aires Cape Town Chennai
Dar es Salaam Delhi Hong Kong Istanbul Karachi Kolkata
Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi
São Paulo Shanghai Taipei Tokyo Toronto

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

This edition © Richard Swinburne 2004

The moral rights of the author have been asserted

Database right Oxford University Press (maker)

First published 1979

Reissued with appendices 1991

This edition published 2004

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

ISBN 0-19-927167-4

ISBN 0-19-927168-2 (pbk.)

1 3 5 7 9 10 8 6 4 2

Typeset in 10.5pt/12pt Minion by Kolam Information Services Pvt. Ltd., Pondicherry, India

Printed in Great Britain on acid-free paper by

Biddles Ltd., King's Lynn, Norfolk

Preface to the Second Edition

The Existence of God is the central book of all that I have written on the philosophy of religion. It was originally published in 1979. A 'revised edition' was published in 1991, but the revision consisted merely in the addition of two appendices; the main text remained intact. The present revision is a much more substantial one. I have rephrased my accounts in Chapters 6, 7, and 8 of the cosmological and teleological arguments, incorporating here the material of the 1991 appendices, developing the argument from laws of nature by a discussion of the nature of laws of nature (depending on a rewritten Chapter 2) and improving my account of the argument from fine-tuning. I have altered Chapter 9 in the light of my subsequent work on consciousness; and Chapters 10 and 11 in the light of my subsequent work on the problem of evil. I have added three additional notes—one to show how arguments to the existence of one God are compatible with the Christian doctrine of the Trinity (God as 'three persons of one substance'), and two discussing recent influential variants of an argument from design. I have largely rearranged the material of Chapter 12 in order to make the argument more perspicuous. There are also smaller alterations at various other places in the book. In the course of these various alterations, I have connected what I have to say with recent important new books and articles. Although my views on many minor matters involved in the argument of the first edition of *The Existence of God* have changed, I remain convinced of the correctness of its general approach to the topic, and of its resulting conclusion. A diligent student of the earlier editions will, however, detect marginally more sympathy for the argument from evil against the existence of God, balanced by marginally greater confidence in the force of the argument from moral awareness for the existence of God (and also considerable confidence in the force of an argument from the miracle of the Resurrection of Jesus, to which for reasons of space I merely allude in this book, but

for which I have argued in detail in my book *The Resurrection of God Incarnate* (Clarendon Press, 2003)).

The first edition was based on two series of Wilde Lectures given in the University of Oxford in Hilary Term 1976 and in Hilary Term 1977; and on two Forwood Lectures given in the University of Liverpool in February 1977. I am grateful to those who originally elected me to these lectureships; and to everyone who has helped me subsequently in my understanding of the issues in oral discussion and in published criticism. My critics are many, and they have provided much help.

I am grateful to the editors and publishers of the journals concerned for permission to reuse material that was incorporated in the earlier editions, from these articles: 'Whole and Part in Cosmological Arguments', *Philosophy*, 44 (1969) 339–40; 'The Argument from Design', *Philosophy*, 43 (1968) 199–212; 'The Argument from Design—A Defence', *Religious Studies*, 8 (1972) 193–205; 'The Problem of Evil', in S. C. Brown (ed.), *Reason and Religion* (Cornell University Press, 1977); 'Natural Evil', *American Philosophical Quarterly*, 15 (1978), 295–301; 'Mackie, Induction, and God', *Religious Studies*, 19 (1983), 385–91; 'The Argument from the Fine-Tuning of the Universe', in J. Leslie (ed.), *Physical Cosmology and Philosophy* (Collier MacMillan, 1990). Thanks to editors and publishers for permission to use more recent material from the following articles: 'The Argument from Laws of Nature Reassessed', in M. Stone (ed.), *Reason, Faith and History: Essays in Honour of Paul Helm* (Ashgate, 2004), 'The Argument to God from Fine-Tuning Reassessed', in N. A. Manson (ed.), *God and Design: The Teleological Argument and Modern Science* (Routledge, 2003); 'What is so Good about Having a Body?', in T. W. Bartel (ed.), *Comparative Theology* (SPCK, 2003); and 'Prior Probabilities in the Argument from Fine-Tuning', forthcoming in a supplement to *Faith and Philosophy*. Thanks to the Oxford University Press for permission to reuse verbatim in Chapter 9 a large section of my shorter book *Is There a God?* (Oxford University Press, 1996); and in Chapter 11 passages from my book *Providence and the Problem of Evil* (Clarendon Press, 1998). And finally very many thanks to Sarah Barker for her patient typing and retyping of many versions of this new edition.

Contents

Introduction	1
1. Inductive Arguments	4
2. The Nature of Explanation	23
3. The Justification of Explanation	52
4. Complete Explanation	73
5. The Intrinsic Probability of Theism	93
6. The Explanatory Power of Theism: General Considerations	110
7. The Cosmological Argument	133
8. Teleological Arguments	153
9. Arguments from Consciousness and Morality	192
10. The Argument from Providence	219
11. The Problem of Evil	236
12. Arguments from History and Miracles	273
13. The Argument from Religious Experience	293
14. The Balance of Probability	328
Additional Note 1: The Trinity	343
Additional Note 2: Recent Arguments to Design from Biology	346
Additional Note 3: Plantinga's Argument against Evolutionary Naturalism	350
<i>Concordance</i>	355
<i>Index</i>	357

This page intentionally left blank

Introduction

The Existence of God is a sequel to *The Coherence of Theism*, originally published in 1977. *The Coherence of Theism* was concerned with what it means to say that there is a God and whether the claim that there is a God is internally coherent. *The Existence of God* is concerned with whether the claim is true; it is concerned to assess the weight of arguments from experience for and against this claim, and to reach a conclusion about whether on balance the arguments indicate that there is a God or that there is not. The present book assumes that the claim that there is a God is not demonstrably incoherent (i.e. logically impossible), and hence that it is proper to look around us for evidence of its truth or falsity. For argument in justification of this assumption I must refer to the earlier work. However, it is in no way necessary for a reader to have read the earlier work in order to understand this one; nor, with the exception just described, does this work in any way presuppose the results of the earlier one. The issues discussed in *The Existence of God* are ones of more general concern than those discussed in *The Coherence of Theism*. Most people have usually supposed that they understood in some very vague way what it meant to say that there was a God; and, so long as they supposed that human words were only a rough guide to what was claimed, that the claim was not demonstrably incoherent. Intense concern about the exact meaning of the claim and whether it is coherent has been primarily the concern of professional theologians and philosophers. But what has worried ordinary people down the centuries is whether the evidence of human experience shows that the claim is true or that it is false. That issue is the topic of this book. The book aims to discuss the topic in depth and with rigour.

The book is written in deep conviction of the possibility of reaching a fairly well-justified conclusion by rational argument on

this issue, perhaps the most important of all deep issues that stir the human mind. It is a conviction that was explicitly acknowledged by the vast majority of Christian (and non-Christian) philosophers from the thirteenth to the eighteenth centuries; and, I believe, shared, although discussed only briefly, by most Christian (and non-Christian) philosophers from the first to the twelfth centuries. By the nineteenth century, however, philosophical theology began to feel the powerful sceptical influence of Hume and Kant. These philosophers produced principles designed to show that reason could never reach justified conclusions about matters much beyond the range of immediate experience, and above all that reason could never reach a justified conclusion about the existence of God. In recent years many others have argued in the same spirit, so that, both among professional philosophers and outside their narrow circle, there is today deep scepticism about the power of reason to reach a justified conclusion about the existence of God.

As I construct my positive arguments, I shall briefly give my grounds for thinking that the principles of Hume and Kant are mistaken and that reason can reach justified conclusions outside the narrow boundaries drawn by those philosophers. Those who believe in the ability of modern science to reach justified (and exciting) conclusions about things far beyond immediate experience, such as subatomic particles and nuclear forces, the 'Big Bang' and cosmic evolution, ought to be highly sympathetic to my enterprise; Hume and Kant would not, on their own principles, have had a very sympathetic attitude to the claims of modern physical science.

I shall, however, argue that, although reason can reach a fairly well-justified conclusion about the existence of God, it can reach only a probable conclusion, not an indubitable one. For this reason, there is abundant room for faith in the practice of religion, and my trilogy on the philosophy of theism ends with a volume on *Faith and Reason*.

Recent developments in philosophy which I shall describe, especially developments in inductive logic, often called confirmation theory, provide tools of great value for the investigation of my topic. Confirmation theory involves some occasional use of symbols. I introduce these symbols in the text and explain their meaning with the aid of examples. There is no need for any reader unfamiliar with such symbols to take fright at them. My use of confirmation theory enables me to express my arguments with the rigour appropriate to

any detailed presentation of the evidence for and against a large-scale theory of the universe; and also enables me to bring out the close similarities that exist between religious theories and large-scale scientific theories. I do, however, owe an apology, as well as an explanation, to those who find it difficult to cope with symbols. The symbols are not very frequent, and I have been careful to express the main argument of the passages in which symbols occur in words as well.

Inductive Arguments

An argument starts from one or more premisses, which are propositions taken for granted for the purpose of the argument, and argues to a conclusion. An argument is a valid deductive argument if it is incoherent to suppose that its premisses are true but its conclusion false. For example, the following argument is a valid deductive argument:

(Premiss 1) No material bodies travel faster than light.

(Premiss 2) My car is a material body.

(Conclusion) My car does not travel faster than light.

In a valid deductive argument the premisses make the conclusion certain. There are arguments that are not deductively valid, but in which the premisses in some sense ‘support’ or ‘confirm’ or ‘give strength to’ the conclusion, and some or all arguments of this general kind are often characterized as ‘good’ or ‘correct’ or ‘strong’ inductive arguments. However, we need here to distinguish carefully between two different kinds of argument. There are arguments in which the premisses make the conclusion probable, that is, more probable than not—for example:

P_1 : 70% inhabitants of the Bogside are Catholic.

P_2 : Doherty is an inhabitant of the Bogside.

C: Doherty is Catholic.

The conjunction of the premisses makes the conclusion probable. However, many arguments that are called ‘correct’ inductive arguments are hardly to be regarded as of this type. Take the following argument:

P: All of 100 ravens observed in different parts of the world are black.

C: All ravens are black.

The normal way to construe this conclusion, in the context of a discussion of inductive arguments, is to suppose that it is about all ravens at all moments of time and points of space—and, even if you suppose that nothing on a distant planet would count as a raven, that means all ravens at all times in the earth's history and at all places on its surface. But, when the conclusion is interpreted this way, it becomes implausible to suppose that P makes C more probable than not. For it is not improbable to suppose that the blackness of observed ravens arises from a particular feature of modern ravens, a particular feature of their make-up not present in older ravens. To suppose that all ravens are always black seems to go a long way beyond the evidence recorded in P. C may, however, be true; and, most of us suppose, P increases the probability that it is true, but P does not make C probable.

Most of the arguments of scientists from their observational evidence to conclusions about what are the true laws of nature or to predictions about the results of future experiments or observations are not deductively valid, but are, it would be generally agreed, inductive arguments of one of the above two kinds. (I do not mean that they have the simple pattern of the easy examples given above, but only that they are arguments that have the defining characteristics of one of the two kinds.) The various astronomical observations made by Tycho Brahe, Kepler, Galileo, and other men of the seventeenth century were observations that favoured Newton's theory of motion, in the sense that they made it more likely to be true, more probable, than it would have been otherwise. The various botanical, geological, and breeding data described by Charles Darwin in *The Origin of Species* added to the probability of his theory of the evolution of animal species by natural selection of variations. It is an interesting question, to which I shall need to allude at a later stage, whether, in a typical scientific argument from various data of observation and experiment to a conclusion about what are the fundamental laws of physics or chemistry, the premisses make the conclusion probable or merely add to its probability. Laws of nature are normally supposed to be generalizations that not merely hold at all times and places, but would continue to hold under unrealized or unrealizable circumstances (for example, however humans interfere with the universe). Newton's theory of motion consists of his three laws of motion and his law of gravitational attraction. Did the various observations of the seventeenth century make it more probable than not that his theory

was true? I pass no judgement on this matter at this stage. However, on our normal way of looking at these matters, clearly observational evidence often makes more probable than not a particular prediction about the future. All the observational evidence about the past behaviour of sun, moon, planets, etc. makes it more probable than not that the earth will continue to spin on its axis for the next twenty-four hours and so that the sun will rise over the earth again tomorrow.

Let us call an argument in which the premisses make the conclusion probable a correct P-inductive argument. Let us call an argument in which the premisses add to the probability of the conclusion (that is, make the conclusion more likely or more probable than it would otherwise be) a correct C-inductive argument. In this latter case let us say that the premisses 'confirm' the conclusion. Among correct C-inductive arguments, some will obviously be stronger than others, in the sense that in some the premisses will raise the probability of the conclusion more than the premisses do in other arguments.

The point of arguments is to get people, in so far as they are rational, to accept conclusions. For this purpose it is not sufficient that their premisses should in some sense necessitate or probabilify their conclusion. It is also necessary that the premisses should be known to be true by those who dispute about the conclusion. There are plenty of valid arguments to the existence of God that are quite useless, because, although their premisses may be true, they are not known to be true by those who argue about religion—for example:

P₁: If life is meaningful, God exists.

P₂: Life is meaningful.

C: God exists.

This argument is certainly valid. If the premisses are true, the conclusion must be true. The premisses may be true; but atheists would deny either the first premiss or the second one. Since the premisses are not common items of knowledge to those who argue about religion, they do not form a suitable jumping-off ground for such argument. What are clearly of interest to people in an age of religious scepticism are arguments to the existence (or non-existence) of God in which the premisses are known to be true by people of all theistic or atheistic persuasions. I therefore define arguments from premisses known to be true by those who dispute about the

conclusion which are valid deductive, correct P-inductive, or correct C-inductive arguments, respectively good deductive, good P-inductive, and good C-inductive arguments. In investigating arguments for or against the existence of God, we need to investigate whether any of them is a good deductive, good P-inductive, or good C-inductive argument.

I take the proposition 'God exists' (and the equivalent proposition 'There is a God') to be logically equivalent to 'there exists necessarily a person¹ without a body (i.e. a spirit) who necessarily is eternal, perfectly free, omnipotent, omniscient, perfectly good, and the creator of all things'. I use 'God' as the name of the person picked out by this description. I understand by God's being eternal that he always has existed and always will exist. There is an alternative understanding of 'eternal' in the Christian tradition as 'timeless' or 'outside time'. This understanding did not, however, arrive in the Christian tradition under the fourth century AD; it is very difficult to make any sense of it, and, for reasons that I have given elsewhere,² it seems quite unnecessary for the theist to burden himself with this understanding of eternity. By God's being perfectly free I understand that no object or event or state (including past states of himself) in any way causally influences him to do the actions that he does—his own choice at the moment of action alone determines what he does. By God's being omnipotent I understand that he is able to do whatever it is logically possible (i.e. coherent to suppose) that he can do. By God's being omniscient I understand that he knows whatever it is logically possible that he know. By God's being perfectly good I understand that he always does a morally best action (when there is one), and does no morally bad action. By his being the creator of all things I understand that everything that exists at each moment of time (apart from himself) exists because, at that moment of time, he makes it exist, or permits it to exist. The meaning of this claim that there is a God will be developed in somewhat greater detail at points in later chapters, especially in Chapter 5.³ The claim that there is a God is called theism. Theism is, of course, the core belief of the creeds of Christianity, Judaism, and Islam.

¹ In understanding God as a person, while being fair to the Judaic and Islamic view of God, I am oversimplifying the Christian view. See my Additional Note 1.

² See *The Coherence of Theism* (Clarendon Press, 1993), ch. 12.

³ For more thorough analysis I must refer the reader to *The Coherence of Theism* and my book *The Christian God* (Clarendon Press, 1994).

In the course of human history many people have taken for granted the existence of God, and many others no doubt have taken for granted his non-existence. They have not had consciously formulated reasons for their beliefs. They have just believed. However, others who have believed have had reasons for their beliefs. As with most people's reasons for most of their beliefs, these reasons have often been very vague and inchoate. Sometimes, however, people have formulated some of their reasons for belief in a sharp and explicit form. Then we have something clearly recognizable as an argument for or against the existence of God. Those arguments that have been frequently discussed have been given names—and thus we have 'the cosmological argument', or 'the argument from religious experience'. Other arguments exist that have not been discussed frequently enough to gain a name. And people have had other reasons for belief or disbelief that have never been formulated explicitly enough to constitute an argument.

In the course of this book I shall discuss various of the reasons that people have had for believing in the existence of God, or in the non-existence of God, some of which have received a sufficiently precise form already to be codified in named arguments and others of which will need to be knocked into a clear shape. I shall discuss only arguments in which the premisses report what are (in some very general sense) features of human experience—for example, evident general truths about the world or features of private human experience. Such arguments I shall term *a posteriori* arguments. They claim that something that humans experience is grounds for believing that there is a God or that there is no God. I shall not discuss *a priori* arguments—these are arguments in which the premisses are logically necessary truths—namely, propositions that would be true whether or not there was a world of physical or spiritual beings. Among logically necessary truths are the truths of mathematics or logic. Hence I shall not discuss the traditional ontological argument⁴ for the existence of God, or any variants thereof. Nor shall I discuss arguments against

⁴ The traditional version of the ontological argument was put forward by Descartes and probably originally by St Anselm. It runs roughly as follows: 'God is by definition a most perfect being. A being which exists is more perfect than one which does not. Therefore, God, being most perfect, exists.' For ancient and modern versions of the argument and criticisms of it, see (e.g.) the collection edited by A. Plantinga, *The Ontological Argument* (MacMillan, 1968). For a very careful analysis leading to a rejection of the argument, see J. Barnes, *The Ontological Argument* (MacMillan, 1972).

the existence of God that claim that there is something incoherent or self-contradictory in the claim that there is a God. I think that ontological arguments for the existence of God are very much mere philosophers' arguments and do not codify any of the reasons that ordinary people have for believing that there is a God. The greatest theistic philosophers of religion have on the whole rejected ontological arguments and relied on *a posteriori* ones.⁵ Arguments against the existence of God that claim that theism is incoherent do, however, I admit, have some basis in the thought of ordinary people. I shall not, however, of course be able to discuss all the *a posteriori* reasons that people have had for believing that there is or that there is not a God. But I shall consider those that, in my view, are the most plausible and have had the greatest appeal in human history. In reaching my final conclusion about how probable it is that there is a God, I assume that no *a priori* arguments of either species,⁶ and no *a posteriori* arguments other than those that I discuss, have any significant force.

Although my theme is arguments for and against the existence of God, it will seem that I concentrate on arguments for the existence of God. I do discuss in a separate chapter the main argument against the existence of God—the argument from evil, which claims that the existence of pain and suffering in the world shows that there is no perfectly good and all-powerful being. But, apart from that argument (and the associated argument from hiddenness, which I also discuss there), the main reason that atheists have for believing that there is no God has been their claim that there is insufficient evidence, that the theist's arguments do not make the existence of God probable to any significant degree. The atheist's arguments, apart from the argument from evil, have been largely in the form of criticisms of the theist's arguments. I therefore discuss such arguments in the course of discussing each of the main arguments for the existence of God. In discussing arguments for the existence of God, I shall consider forms of cosmological and teleological argument, the argument from the existence of consciousness, the moral argument, arguments from miracle and revelation, and the argument from religious experience. A cosmological argument argues that the fact that there is a universe

⁵ e.g. St Thomas Aquinas. See his *Summa Theologiae*, Ia2.1.

⁶ I attempt to prove this for arguments that purport to show that theism is incoherent in *The Coherence of Theism*.

needs explaining and that God's having made it and kept it in being explains its existence. An argument from design argues that the fact that there is design in the world needs explaining, and that God's action provides that explanation. There are various forms of argument from design, according to the kind of design to which it draws attention. I discuss two different genera of the argument under the headings 'teleological arguments' and 'the argument from providence', and different species of each genus. The argument from the existence of consciousness argues that the fact that there are conscious beings is mysterious and inexplicable but for the action of God. Arguments from miracle and revelation cite various public phenomena in the course of human history as evidence of God's existence and activity. The argument from religious experience claims that various private experiences are experiences of God and thus show his existence.

Some of the issues that I discuss are ones that I have treated at greater length elsewhere; but the discussion in this book is, I hope, adequate—given the constraints imposed by the length of the book—to support the conclusions drawn here. For example, I discussed the problem of evil at book length in my book *Providence and the Problem of Evil*;⁷ but I hope that the discussion of it in Chapters 10 and 11 of the present book suffice to make it plausible that the kind and amount of evil that we find on Earth do not count significantly against the existence of God. Yet there is one respect in which my discussion in this book is manifestly incomplete. When I discuss arguments from miracles, I have space only to discuss which strange public phenomena (for example, a dead man coming to life) if they occurred would be evidence for the existence of God, but I do not have space to discuss the historical evidence for and against the occurrence of particular public phenomena. So in effect I discuss here only the form of an argument that needs filling out with detailed historical material.⁸

Kant produced a threefold classification of arguments for the existence of God that has had a permanent and to my mind far from beneficial influence on the subsequent discussion of this topic. He wrote:

⁷ (Clarendon Press, 1998).

⁸ I provide that filling-out with respect to the miracle crucial for the Christian religion, the purported Resurrection of Jesus Christ, in my book *The Resurrection of God Incarnate* (Clarendon Press, 2003).

There are only three possible ways of proving the existence of God by means of speculative reason. All paths leading to this goal begin either from determinate experience and the specific constitution of the world of sense as thereby known, and ascend from it, in accordance with the laws of causality, to the supreme cause outside the world; or they start from experience which is purely indeterminate, that is from experience of existence in general; or finally they abstract from all experience, and argue completely a priori, from mere concepts, to the existence of a supreme cause. The first proof is the *physico-theological*, the second the *cosmological*, the third the *ontological*. There are, and there can be, no others.⁹

The distinction is made in terms of the nature of the premiss. Either you start from a conceptual truth—in which case you have the ontological argument; or from ‘existence in general’—in which case you have the cosmological argument; or from the details of what Kant calls ‘determinate experience’, how things are in the world—in which case you have the ‘physico-theological’ argument.

My reason for claiming that this doctrine of Kant has had a far from beneficial influence on discussion of this topic is that by his use of the word ‘the’ Kant tends to assume that there can be only one argument of each type—whereas in fact there can quite clearly be many different arguments under each heading that are so different from each other that it would be misleading to call them forms of the same argument at all. There is, for example, no reason to suppose that all arguments to the existence of God in which the premisses are in some sense conceptual truths need have the form of the traditional ontological argument. Above all, there is no reason to suppose that all arguments from how things are in the world need have the form of the argument that Kant calls ‘physico-theological’, and has elsewhere been called the argument from design. This latter argument may itself have many forms. It may argue, for example, from the regular behaviour of objects in the world codified in laws of nature, or from the ready availability in the world of the things that humans and animals need to survive. In both cases there is an argument from a very general order in nature. But there are arguments too, as we have noted, from particular miracles, from the development of human history, or from particular religious experiences. Not all of these may be particularly good arguments but they deserve to be considered on their merits—Kant’s classification obscures their existence.

⁹ I. Kant, *Critique of Pure Reason*, B618–19, trans. N. Kemp Smith (MacMillan, 1964).

So then we shall consider the worth of various *a posteriori* arguments, not merely two, as listed by Kant. When we have our arguments in clear form, we shall need to ask—are they good deductive arguments, or good P-inductive arguments, or good C-inductive arguments? Sometimes the proponents of such arguments have not been clear whether the arguments were intended to be deductive or inductive, let alone about the kind of inductive arguments that they were intended to be.

One unfortunate feature of recent philosophy of religion has been a tendency to treat arguments for the existence of God in isolation from each other. There can, of course, be no objection to considering each argument initially, for the sake of simplicity of exposition, in isolation from others. But clearly the arguments may back each other up or alternatively weaken each other, and we need to consider whether or not they do. Sometimes, however, philosophers consider the arguments for the existence of God in isolation from each other, reasoning as follows: the cosmological argument does not prove the conclusion, the teleological argument does not prove the conclusion, etc., etc., therefore the arguments do not prove the conclusion. But this ‘divide and rule’ technique with the arguments is admissible. Even if the only kind of good argument was a valid deductive argument from premisses known to be true, it would be inadmissible. An argument from *p* to *r* may be invalid: another argument from *q* to *r* may be invalid. But, if you run the arguments together, you could well get a valid deductive argument; the argument from *p* and *q* to *r* may be valid. The argument from ‘all students have long hair’ to ‘Smith has long hair’ is invalid, and so is the argument from ‘Smith is a student’ to ‘Smith has long hair’; but the argument from ‘all students have long hair and Smith is a student’ to ‘Smith has long hair’ is valid.

That arguments may support and weaken each other is even more evident, when we are dealing with inductive arguments. That Smith has blood on his hands hardly makes it probable that Smith murdered Mrs Jones, nor (by itself) does the fact that Smith stood to gain from Mrs Jones’s death, nor (by itself) does the fact that Smith was near the scene of the murder at the time of its being committed, but all these phenomena together (perhaps with other phenomena as well) may indeed make the conclusion probable.¹⁰

¹⁰ Among those who seem to have assumed that there are no good arguments other than deductive ones, and that arguments are not cumulative, are both (the early) Alistair

In order to consider the cumulative effect of arguments, I shall consider them one by one, starting with the cosmological argument and including the arguments from evil and from hiddenness against the existence of God, and ask how much the premisses of each argument add to or subtract from the force of the previous arguments. To give advance notice of some of my conclusions, I shall argue that (neither separately nor in conjunction) are any of the arguments that I consider for or against the existence of God good deductive arguments. There are, of course, as I have pointed out, valid deductive arguments to the existence of God, but they start from premisses that are far from generally accepted. On the other hand, I shall argue that most of the arguments (taken separately and together) for the existence of God are good C-inductive arguments—that is to say, their premisses make it more probable (or likely) that God exists than it would otherwise be. Some of these arguments of course confirm the existence of God much more strongly than do others. I shall allow that the argument against the existence of God from evil is a good C-inductive argument of very limited force. I shall claim that the argument from hiddenness to the non-existence of God is not a good C-inductive argument. The crucial issue, however, is whether all the arguments taken together make it probable that God exists, whether the balance of all the relevant evidence favours the claim of theism or not. For clearly, in so far as the probability of a hypothesis is relevant to whether or not we ought to act on it, we ought to act on a hypothesis in so far as it is rendered probable by the total evidence available to us—all we know about the world, not just some limited piece of knowledge. The religious person claims that his

MacIntyre and Antony Flew. Thus MacIntyre: ‘One occasionally hears teachers of theology aver that although the proofs do not provide conclusive grounds for belief in God, they are at least pointers, indicators. But a fallacious argument points nowhere (except to the lack of logical acumen on the part of those who accept it). And three fallacious arguments are no better than one’ (A. MacIntyre, *Difficulties in Christian Belief* (SCM Press, 1959), 63). This passage is quoted with approval by Flew in his *God and Philosophy* (Hutchinson, 1966), 167, who remarks himself in another very similar passage: ‘It is occasionally suggested that some candidate proof, although admittedly failing as a proof, may sometimes do useful service as a pointer. This is a false exercise of the generosity so characteristic of examiners. A failed proof cannot serve as a pointer to anything, save perhaps to the weaknesses of those who have accepted it. Nor, for the same reason can it be put to work along with other throwouts as part of an accumulation of evidences. If one leaky bucket will not hold water that is no reason to think that ten can’ (ibid. 62–3). But, of course, arguments that are not deductively valid are often inductively strong; and, if you put three weak arguments together, you may often get a strong one, perhaps even a deductively valid one.

religious viewpoint makes sense of the whole of his experience; and his atheistic rival is liable to make a similar claim. In the final chapter I shall reach a conclusion on whether or not the balance of all the relevant evidence favours theism. I shall be fairly brief in dismissing the suggestions that any of the arguments separately or all the arguments taken together constitute a good deductive argument. I shall be fairly brief because many other philosophers have devoted their technical skills to this task, and relatively few philosophers today would accept that there are good deductive arguments to be had here. I shall devote most of my time to assessing the inductive strength of such arguments. I shall consider of each argument whether it is a good C-inductive argument, but only when we have all the arguments shall I ask whether, taken together, they make a good P-inductive argument. I proceed in this way because, as will appear, it is a lot easier to see when we have a good C-inductive argument than when we have a good P-inductive argument.

It will be useful to introduce at this stage the symbols of confirmation theory that I shall use from time to time in subsequent chapters. I represent by lower-case letters such as e , h , p , and q propositions. $P(p|q)$ represents the probability of p given q . Thus p might represent the proposition: 'The next toss of this coin will land heads', and q might represent the proposition: '505 of the last 1,000 tosses of this coin have landed heads'. Then $P(p|q)$ represents the probability that the next toss of the coin will land heads, given that 505 of the last 1,000 tosses have landed heads. (The value of $P(p|q)$ would then generally be supposed to be 0.505.) However, the relation between p and q may be of a much more complex kind; and clearly we normally assess the probability of claims on evidence other than or additional to that of relative frequencies. p may be some scientific hypothesis—say, Einstein's General Theory of Relativity—and q may be the conjunction of all the reports of the evidence of observation and experiment that scientists have collected relevant to the theory. Then $P(p|q)$ represents the inductive probability of Einstein's General Theory given all the reports of relevant observations and experiments. Inductive probability is thus to be distinguished from statistical probability, which is a property of classes of things (for example, inhabitants of a certain town, say Tunbridge Wells) and is a measure of the proportion of things in the class that have some other property (for example, voting Conservative in the 2001

Election). The probability of an inhabitant of Tunbridge Wells voting Conservative in 2001 is just the proportion of inhabitants of Tunbridge Wells who voted Conservative in 2001. (In English, the indefinite article—for example, ‘the probability of *an* inhabitant . . .’—often indicates that the probability is statistical.) The classes may be of actual things (for example, inhabitants of Tunbridge Wells), or of hypothetical things, things that would be generated by a certain process (for example, tosses of this coin, if we were to toss it for a very long time).

Inductive probability is also to be distinguished from physical probability. The physical (or natural) probability of an event (and so of the proposition that records it) is a matter of the extent to which at some earlier time the event is predetermined by its causes. An event that is made inevitable by the preceding state of the world has a physical probability of 1—its occurrence is physically necessary; an event whose non-occurrence is made inevitable by the preceding state of the world has a physical probability of 0—its occurrence is physically impossible. An event has a physical probability between 1 and 0 if it is not predetermined that it will happen or that it will not happen, but the preceding state of the world is biased in favour of its happening to the degree measured by the value of the probability: larger values of the probability indicate a greater bias in favour of its happening.¹¹ Physical and statistical probabilities may themselves constitute evidence that makes some hypothesis inductively probable; or other evidence may make it inductively probable that they have a certain value.

My concern with inductive probability is a concern with how probable q makes p , quite apart from who is doing the calculation, how clever he is, and his degree of confidence in the evidential force of q . Clearly in science and history and all other empirical inquiries we think that there are correct ways to assess whether and (within rough limits) how much some evidence supports some hypothesis. I shall set out these criteria in Chapter 3. In order to emphasize the objective character of the value $P(p|q)$ with which I am concerned and to distinguish it from measures of evidential support that

¹¹ I call probability of this kind ‘physical probability’ because this term has a certain currency, but I do not wish to imply that it applies only to physical objects or states. There can be in the sense defined some physical probability that some mental event occurs.

measure subjects' degree of confidence or are in part functions of subjects' abilities to work out the true measure of evidential support,¹² I shall in future call $P(p|q)$, the logical probability of p on q . This is clearly an *a priori* matter. If q represents all the relevant evidence, the value of $P(p|q)$ cannot depend on further evidence—it measures what the evidence you have already got shows. It is an *a posteriori* matter whether, in 1,000 tosses, 505 have landed heads; but an *a priori* matter whether that evidence gives a probability of 0.505 to the next toss landing heads.

A hypothesis up for investigation is often represented by h . Then $P(h|e \& k)$ represents the probability of a hypothesis h given evidence ($e \& k$).¹³ It is often useful to divide the evidence available to an observer into two parts—new evidence and background evidence; if this is done, the former is often represented by e and the latter by k . Background evidence (or background knowledge, as it is sometimes called) is the knowledge that we take for granted before new evidence turns up. Thus, suppose that detectives are investigating a murder. h could represent the hypothesis that Jones did the murder; e could represent the proposition that reports all the new evidence that detectives discover—for example, that Jones's fingerprints were found on the weapon, that he was near the scene of the murder at the time it was committed, etc., etc. k could represent the proposition reporting the detectives' general knowledge about how the world works—for example, that each person has a unique set of fingerprints, that people who touch metal and wood with bare hands usually leave their fingerprints on them, etc., etc. Then $P(h|e \& k)$ represents the probability that Jones did the murder, given detectives' total evidence.

For all propositions p and q $P(p|q) = 1$ if (and only if) q makes p certain—for example, if q entails p (that is, there is a deductively valid argument from q to p); and $P(p|q) = 0$ if (and only if) q makes $\sim p$ certain—for example, if q entails $\sim p$.¹⁴ $P(p|q) + P(\sim p|q) = 1$. So if $P(p|q) > 1/2$, then $P(p|q) > P(\sim p|q)$ and it is on q more

¹² For elucidation of the distinction between logical probability and other kinds of inductive probability, which I call 'epistemic probability' and 'subjective probability', see my *Epistemic Justification* (Clarendon Press, 2001), ch. 3.

¹³ ' $e \& k$ ' is the conjunction of e and k , the proposition 'both e and k '.

¹⁴ ' $\sim p$ ' is the negation of p , the proposition 'it is not the case that p '. '>' means 'is greater than'. '<' means 'is less than'. I shall also subsequently use '≥' to mean 'is greater than or equal to', and '≤' to mean 'is less than or equal to'.

probable that p than that $\sim p$. So (for background knowledge k) an argument from e to h will be a correct C-inductive argument if (and only if) $P(h|e \& k) > P(h|k)$, and a correct P-inductive argument if (and only if) $P(h|e \& k) > 1/2$. The division between new evidence and background evidence can be made where you like—often it is convenient to include all evidence derived from experience in e and to regard k as being what is called in confirmation theory mere ‘tautological evidence’, that is, in effect all our other irrelevant knowledge.

My strategy will be as follows. Let h be our hypothesis—‘God exists’. Let e_1, e_2, e_3 , and so on be the various propositions that people bring forward as evidence for or against his existence, the conjunction of which form e . Let e_1 be ‘there is a physical universe’. Then we have the argument from e_1 to h —a cosmological argument. In considering this argument I shall assume that we have no other relevant evidence, and so k will be mere tautological evidence. Then $P(h|e_1 \& k)$ represents the probability that God exists given that there is a physical universe—and also given mere tautological evidence, which latter can be ignored. If $P(h|e_1 \& k) > 1/2$, then the argument from e_1 to h is a good P-inductive argument. If $P(h|e_1 \& k) > P(h|k)$, then the argument is a good C-inductive argument. But, when considering the second argument, from e_2 (which will be the conformity of the universe to temporal order), I shall use k to represent the premiss of the first argument e_1 ; and so $P(h|e_2 \& k)$ will represent the probability that God exists, given that there is a physical universe and that it is subject to temporal order. And, when considering the third argument, from e_3 , k will represent the premiss of the second argument ($e_1 \& e_2$). And so on. In this way all relevant evidence will eventually be fed into our assessment. I shall consider some eleven arguments. I shall claim that for most of these e_n , where $n = 1, \dots, 11$, $P(h|e_n \& k) > P(h|k)$ —that is the argument is a good C-inductive argument for the existence of God, that two of the arguments (one for and one against) have no force ($P(h|e_n \& k) = P(h|k)$ in these cases) and that one argument against has force ($P(h|e_n \& k) < P(h|k)$) where e_n is the occurrence of evil. The crucial issue to which we will eventually come is whether $P(h|e_{11} \& k) > 1/2$.

In using the symbols of confirmation theory I do not assume that an expression of the form $P(p|q)$ always has an exact numerical value. It may merely have relations of greater or less value to other probabilities, including ones with a numerical value, without itself

having a numerical value— $P(h|e_1 \& k)$, for example, might be greater than $P(h|e_2 \& k)$ and less than $P(h|k)$ and less than $1/2$ without there being some number to which it is equal. Clearly, for example, we may judge one scientific theory to be more probable than another on the same evidence while denying that its probability has an exact numerical value; or we may judge a prediction to be more probable than not and so to have a probability of greater than $1/2$, while again denying that that probability has an exact numerical value.

Now it is sometimes said that the different arguments for the existence of God show different things. The cosmological argument shows at most the existence of some sort of necessary being; the argument from design shows at most some sort of arch-architect;¹⁵ the argument from miracles shows at most some sort of poltergeist—so what have they in common? This objection gets things back to front. There is no *one* thing that premisses show. In a deductive argument there are many different conclusions that can be drawn from a set of premisses. And in inductive arguments the premisses support different conclusions with different degrees of force.

What does ‘there is a print in the shape of a human foot on the sand’ show? It shows with different degrees of force many things—that sand is shapeable, that some creature has been on the sand, that a man has walked on the sand. The evidence makes probable the different propositions to different degrees. Our concern is with the effect of various pieces of evidence on the proposition in which we are interested—‘God exists’. Does each confirm it (that is, increase its probability)? Does it make it probable? Our concern is for various pieces of evidence e_n (including any k) and for $h =$ ‘God exists’ with the value of $P(h|e_n)$. This may well be for some e_n less than the value for some other interesting proposition h_1 , say, ‘there exists an impersonal cause of the universe’, of $P(h_1|e_n)$. That is, e_n may make h_1 more probable than it makes h . However, even though, say, $P(h_1|e_1) > P(h|e_1)$, it certainly does not follow that $P(h_1|e_1 \dots e_7) > P(h|e_1 \dots e_7)$. That is, ‘God exists’ may gain only a small amount of probability from e_1 , a small amount from e_2 , a small amount from e_3 , and so on. For each of e_1 , e_2 , e_3 , there may be some other proposition h_1 , h_2 , h_3 , which is in some sense a rival to ‘God exists’ for

¹⁵ See e.g. Kant’s treatment of the argument in *The Critique of Pure Reason*, B648–58. He writes (B655): ‘The utmost, therefore, that the argument can prove is an *architect* of the world who is always very much hampered by the adaptability of the material in which he works, not a *creator* of the world to whose idea everything is subject.’

which $P(h_n|e_n) > P(h|e_n)$; but, nevertheless, on the total evidence h may be more probable than each of the rivals.

A similar situation normally arises with any far-reaching scientific or historical theory. Each separate piece of evidence does not make the theory very probable, and indeed taken on its own makes some narrower theory much more probable. But the cumulative force of the evidence taken together gives great probability to the wide theory. Thus each of the various pieces of evidence that are cited as evidence in favour of the General Theory of Relativity do not by themselves make it very probable, but together they do give it quite a degree of probability. Each by itself (given the general background knowledge available in the early twentieth century) was evidence in favour of some rival but far less wide-ranging hypothesis than General Relativity. Thus the movement of Mercury's perihelion taken by itself would suggest only that there was a hitherto unknown planet lying between Mercury and the sun or that the sun was of an odd shape, rather than that General Relativity was true. Taken by itself it would not have given much probability to General Relativity; but taken with other pieces of evidence it did its bit in supporting the latter. It is along these lines that the theist may wish to answer the accusation that an argument such as the cosmological argument does not show the existence of the God of Abraham, Isaac, and Jacob. Not by itself, he may reply, but it does its small bit together with some very diverse arguments that do their small bit, to get to this conclusion.

Note that it is no objection to a P-inductive or C-inductive argument from e to h that some contrary hypothesis h^* is also compatible with e , as some writers on the philosophy of religion seem to think. They seem to think that if, for example, the order in the universe is compatible with 'God does not exist', then there is no good argument from it to 'God exists'. But one has only to think about the matter to realize that this is not so. In any non-deductive argument from e to h , not- h will be compatible with e ; and yet some non-deductive arguments are good arguments.

Note also a further interesting feature of good C-inductive arguments. In such an argument from e to h , $P(h|e \& k) > P(h|k)$. It may be the case that also for some contrary hypothesis h^* there is a good C-inductive argument from e —that is, also $P(h^*|e \& k) > P(h^*|k)$. The fact that certain evidence confirms a hypothesis does not mean that it does not also confirm a rival hypothesis. Once again, this should be immediately clear if one thinks about it. Suppose that a

detective has background information k , that either Smith, Brown, or Robinson did the crime, and that only one of them did. Then evidence (e) turns up that Robinson was somewhere else at the time the crime was committed. e adds to the probability that Brown did the crime, and it also adds to the probability that Smith did the crime. Despite this, one sometimes reads writers on the philosophy of religion dismissing some consideration that is adduced as evidence for the existence of God on the grounds that it supports a rival hypothesis equally well.

So then our task will be to assess the worth of different arguments to the conclusion 'God exists'. How are we to do this? In the case of deductive arguments, philosophers have a moderately clear idea of what makes a valid argument, and so are in a position to look at various arguments and see if they are valid. But our main concern will be with inductive arguments. How are we to set about assessing the probability of 'God exists' on different pieces of evidence? To do this we need to know for what fillings of p and q $P(p|q)$ becomes high or low. There is, however, fortunately no need to undertake any very general examination of this question. This is because all important *a posteriori* arguments for the existence of God have a common characteristic. They all purport to be arguments to a (causal) explanation of the phenomena described in the premisses in terms of the action of an agent who intentionally brought about those phenomena. A cosmological argument argues from the existence of the world to a person, God, who intentionally brought it about. An argument from design argues from the design of the world to a person, God, who intentionally made it thus. All the other arguments are arguments from particular features of the world to a God who intentionally made the world with those features.

Not all inductive arguments are arguments to an explanation. When we argue from the sun having risen at intervals of approximately twenty-four hours over the last many thousand years to the claim that it will rise tomorrow, we are not arguing to an explanation. Its rising tomorrow does not explain its previous rising. Yet when the geologist argues from various deformations to the occurrence of an earthquake millions of years ago, he is arguing to an explanation; he is arguing from phenomena to an event that brought those phenomena about. However, not all arguments to an explanation are arguments to the intentional action of an agent. An intentional action is an action that some agent does, meaning to do it. It is one, therefore,

that the agent has some reason or purpose for doing—either the minimal purpose of doing it for its own sake or some further purpose that is forwarded by doing the action. Since he acts for reasons or purposes on which he chooses to act, we may term such an agent a rational agent. Persons are rational agents;¹⁶ but they are not the only ones—animals too often perform intentional actions. By contrast, however, inanimate objects and events do not have purposes on which they choose to act and which they seek to fulfil, but rather they bring about their effects unthinkingly. The geologist's argument from deformations to the occurrence of an earthquake is an argument to an explanation of the deformations, but not an argument to an explanation in terms of the intentional action of a rational agent. However, when a detective argues from various bloodstains on the woodwork, fingerprints on the metal, Smith's corpse on the floor, money missing from the safe, Jones's having much extra money, to Jones's having intentionally killed Smith and stolen his money, he is arguing to an explanation of the various phenomena in terms of the intentional action of a rational agent. Since persons are paradigm cases of rational agents, I will term explanation in terms of the intentional action of a rational agent personal explanation. In Chapter 2 I shall analyse the nature of personal explanation more fully and I shall contrast it with the other accepted pattern of explaining mundane phenomena, which I shall call scientific explanation.¹⁷ In Chapter 3 I shall go on to consider when it is right to invoke personal explanation and when it is right to invoke scientific explanation. A crucial issue that arises there is when is it reasonable to suppose that phenomena do have a (causal) explanation, and when, by contrast, is it reasonable to suppose that phenomena are just brute facts, things that explain other things, but do not themselves have an explanation. This issue of what is the proper terminus for explanation

¹⁶ I understand by a person a rational agent who has at least the complexity of sensations, desires, beliefs, etc., typical of human beings. See *The Coherence of Theism*, 102–3.

¹⁷ This term is a little misleading, because by contrasting 'scientific explanation' with 'personal explanation' I seem to imply that there cannot be a 'science' of persons, in the sense of an integrated theory of how persons behave. I do not wish to imply this. But I give to the kind of explanation that I analyse in Chapter 2 this name 'scientific explanation' because it is explanation of the kind used in most sciences. Elsewhere (in my book *Is There a God?* (Oxford University Press, 1996)) I have used instead the term 'inanimate explanation', but this too has a misleading implication that I wish to avoid, that there cannot be scientific laws connecting physical and mental events.

will be discussed in Chapter 4. It is one that is crucial for theism. For the theist claims that the various phenomena that constitute his evidence—for example, the existence of the world and its conformity to order—need explanation; and that this is provided by the action of God, whose existence and action need no explanation. So Chapter 2 will bring out the nature of the theist's explanations, and Chapters 3 and 4 will provide essential tools for answering the question of when it is right to invoke them. With these tools we shall then be in a position to look in detail at the theist's arguments.

The Nature of Explanation

General Considerations

When the theist argues from phenomena such as the existence of the universe or some feature of the universe to the existence of God, he is arguing, we have seen, to a causal explanation of the phenomena in terms of the intentional action of a person. Explanation in terms of the intentional action of a person is the normal case of what I termed personal explanation. We give a personal explanation of my being in London by my having gone there in order to give a lecture; or of the letter's being on the table by my wife's having put it there in order to remind me to post it. However, as we have seen, not all explanations are personal explanations. Other explanations of the occurrence of phenomena seem to have a distinct common structure and these I will call scientific explanations. This chapter will be concerned with analysing the structure of explanations of the two kinds; and the next chapter will consider when each is to be invoked.

When someone is said to have provided an explanation of the occurrence of some phenomenon (that is, an event or state of affairs), this is ambiguous. What is meant may be that he has provided a true explanation of the phenomenon, or it may be merely that he has suggested a possible explanation of the phenomenon. Our interest in explanations is interest in true explanations. What is it to provide a true explanation of the occurrence of a phenomenon *E*? It is to state truly what (object or event) brought *E* about (or caused *E*), and why it was efficacious. To explain the occurrence of the high tide is to state what brought about the tide—the moon, water, and the rest of the earth being in such-and-such locations at such-and-such times, and why the moon etc. had that effect—because of the inverse square law of attraction acting between all bodies. We can thus detect two

components of an explanation of a phenomenon E —the ‘what’ that made E happen and the ‘why’ that made E happen. The ‘what’ will be what I may term some other independent actual factors—other events, processes, states, objects, and their properties at certain times. By these factors’ being independent I mean that the ‘what’ is not the same event or process as E nor part of it; nor is it an object that is a participant in E at the time of E , nor is it a state or property of E or the objects that participate in E at the time of E ’s occurrence. Only something different from E can make E happen. By the factors’ being actual I mean only that any events, processes, and states cited occurred; that any object cited existed and had the properties cited.

To say that certain factors $A . . . D$ brought about E entails at least that each, in the conditions of its occurrence, made it more physically probable that E would occur; it influenced E ’s occurrence. Normally, perhaps, each of the factors is necessary, given the others and the world being in other ways the same, for the occurrence of E —that is, without any one of them, the world otherwise remaining the same, E would not have occurred. Normally, perhaps, too the set of factors together is sufficient for the occurrence of E —that is, given their occurrence, E must necessarily occur. We may call all the factors together that make up the ‘what’ the cause of E . Alternatively, or more usually, we distinguish one as the ‘cause’ of E (the effect), and call the others the conditions that were necessary for the cause to have its effect (or at least made it physically probable that it would have the effect); which we call the cause is sometimes a somewhat arbitrary matter. Normally it will be the most unexpected member of the set of factors, or the one, the occurrence of which involves the sharpest change from the previous state of the world. Thus, suppose someone lights a match close to petrol at a certain temperature and a certain pressure, and all of this produces an explosion. We may describe the ignition of the match and the petrol’s being at that temperature and pressure as jointly the cause of the explosion. But it would be more natural to describe the ignition of the match as the cause of the explosion, and the petrol’s being at that temperature and pressure conditions necessary for the cause to have its effect. My terminology will be as follows. I shall call a set of factors that together were sufficient for the occurrence of an event E a *full cause* of E . Any member of a set of factors that contributes towards bringing E about I shall call a *cause* of E .

To set out the 'why' of an explanation is to say why the cause, under the specified conditions, had the effect that it had. Thus it might be to cite a law of nature that all events of a certain kind exemplified by the cause bring about events of a certain other kind exemplified by the effect. To cite the 'why' is to cite what I shall call the *reason* why the cause under the conditions of its occurrence had the effect that it had. I am thus using the word 'reason' in a wide but natural sense—in a wider sense than the sense in which a reason for something is always someone's reason for bringing it about. In saying that something was *the* reason for some effect I do not necessarily imply that it *was someone's* reason for bringing about the effect.

Now, if there is a full cause *C* of *E* and a reason *R* that guarantees *C*'s efficacy, there will be what I shall call a *full explanation* of *E*. For, given *R* and *C*, there will be nothing remaining unexplained about the occurrence of *E*. In this case, the 'what' and 'why' together will deductively entail the occurrence of *E*. But, if there is no full cause of *E* (for example, there occur factors that facilitate the occurrence of *E*, but do not necessitate it) or no reason that ensured that the cause would have the effect that it did, there will be at most what I shall call a *partial explanation* of *E*. Any explanation involving factors or reasons that did not make the occurrence of *E* physically necessary but made it physically more probable than it would otherwise have been, I will term a *partial explanation*. *E* may be given a *partial explanation* because there is no full explanation of *E*. Alternatively, it may well be the case that, even if a full explanation exists, people are in no position to provide it, yet they can give some explanation—they can state some of the causes that make up the 'what' and some of the reasons for their efficacy. In that case they are providing an explanation, but only a partial one.

Also, of course, people may take for granted or not be interested in certain aspects of a full explanation and for that reason give only a partial explanation. A geologist interested in the history of geological formations may explain a present formation by telling the historical story of successive stages in its evolution. In telling this story, he may not bother to cite the physico-chemical laws that are responsible for one stage succeeding another, simply because he is not interested in these. For that reason his explanation is only partial. The context often determines which answers to our questions about 'the explanation' of some phenomena will satisfy us. But, while in other contexts

of discussion we may not need to give full explanations even if they are available, in the contexts of scientific and metaphysical discussion it is often of crucial importance to know whether there is a full explanation of some phenomenon and what its character is.

Scientific Explanation

Explanations are of different patterns according to the different kinds of cause and reason that feature in them. Explanation of the kind used in science I shall call scientific explanation. The classical account of the nature of scientific explanation is that set out carefully by C. G. Hempel and P. Oppenheim, and subsequently championed by Hempel.¹ On the Hempelian account the causes are a group of events (states of affairs or changes thereof) *C*, known as the 'initial conditions', one of which we may arbitrarily select as 'the' cause. The 'why' is a set of natural laws *L*. In the normal case these will be universal generalizations, having the form 'all *A*'s are so-and-so' or 'all *A*'s do so-and-so'—for example, 'all copper put in nitric acid dissolves under such-and-such conditions of temperature and pressure'. *C* and *L* then fully explain *E* if *E* follows deductively from them. We explain a particular explosion by the ignition of a particular volume of gunpowder in certain conditions of temperature, pressure, and humidity, and the generalization that under such circumstances ignited gunpowder explodes. We explain a particular piece of litmus paper's turning red by its having been immersed in acid and the generalization that litmus paper being immersed in acid always turns red. Sophisticated scientific explanations invoke many laws or generalizations and a complex description of previous events, of which it is a somewhat remote deductive consequence that the event or state to be explained occurs. It is a consequence of Newton's laws and arrangements of the sun and planets thousands of years ago that they are in the positions in which they are today, and the former explain their being in those positions.

This normal pattern of scientific explanation is called by Hempel deductive–nomological explanation, or D–N explanation—

¹ For a simple exposition, see C. G. Hempel, *Philosophy of Natural Science* (Prentice-Hall, 1966), ch. 5. The original article dealing only with deductive-nomological explanation is C. G. Hempel and P. Oppenheim, 'Studies in the Logic of Explanation', *Philosophy of Science*, 15 (1948), 135–75.

'deductive' because E is deduced from L and C , and 'nomological', from the Greek *nomos*, 'law', because laws are involved in the explanation. A D-N explanation of an event is a full explanation. However, sometimes the law involved may be a probabilistic law—that is, claim that ' n per cent A 's are B ', where n is intermediate between 100 and 0. It may be a law of genetics that '90 per cent offspring of such-and-such a mating have blue eyes' (or 'there is an 0.9 probability of an offspring of such-and-such a mating having blue eyes'. The probability in this case is a statistical probability.) In such cases, according to Hempel, a law L together with initial conditions C will explain E if L and C make it highly probable that E . (The high probability is in this case an inductive probability, a measure of how much evidence supports some hypothesis, in this case that E occurs.) Thus, if an individual a is an offspring of the stated mating, this together with the law suggested above makes it probable that a has blue eyes; then Hempel holds, the law and the initial conditions together explain a having blue eyes. However, the notion of the inductive probability being 'high' is very vague; and plausibly the law and initial conditions may provide some sort of explanation of an event even if the probability is not very high—so long as the law and initial conditions make the occurrence of the event more probable than it would otherwise be. So, following others,² I shall amend Hempel's account of statistical explanation as follows: a law L and initial conditions C explain an event E if they increase the probability of the occurrence of E . Clearly explanation that involves probabilistic laws is only partial explanation. There is still something unexplained in why the initial conditions were on this occasion efficacious.

Science does not explain only particular events, but it may also explain laws. If it is a consequence of L_1 that, perhaps under particular conditions C , L_2 operates, then L_1 (together with C) explains the operation of L_2 . (If the consequence is deductive, the explanation is a full one; if L_1 only makes the operation of L_2 probable, the explanation is only partial.) More fundamental laws explain the

² This amendment is a simplified version of the amendment proposed by Wesley Salmon. See his paper 'Statistical Explanation' in Wesley C. Salmon (ed.), *Statistical Explanation and Statistical Relevance* (University of Pittsburgh Press, 1971). The considerations that led Salmon to propose a somewhat more complicated theory are, I believe, taken into account by my subsequent requirement that the probability in the law must be a physical probability.

operation of less fundamental laws. Given a certain assumption about the constitution of gases, Newton's laws of motion explain the operation of the Van der Waals gas law. One set of laws is often said to explain another also when a slightly looser relation holds. L_1 (perhaps together with some C) may entail and render it probable that phenomena will be as predicted by L_2 —to a high degree of approximation. It then follows that the true laws of nature in the realm of L_2 are very slightly different from L_2 , but that L_2 is a very close approximation to them. Newton's laws of motion have the consequence that, given the distribution of sun and planets through space, Kepler's laws of planetary motion will hold to a high degree of approximation. I shall follow common usage and say that in such circumstances L_2 operates to a high degree of approximation, and that L_1 explains the operation of L_2 .

It is Hempel's claim that explanation that does not at first sight seem to fit into this scientific pattern can really quite easily be so fitted. Thus we use this scientific pattern of explanation not only when doing science of any degree of sophistication but in much everyday explanation of happenings. We explain the cheese's being mouldy by its having been left in a warm place for two weeks and by the generalization that almost always cheese turns mouldy within two weeks if it is in a warm place. Our explanation may often take the form of explaining some phenomenon as brought about, not by an event, but by an object. We may say that the breaking of the window was brought about by a brick, but what we are saying here, it is urged, is that the breaking was caused by some event involving the brick—for example, its fast motion; and this reduction to the scientific pattern seems initially plausible enough.

This account, however, needs amplification in order to distinguish between merely accidentally true universal or probabilistic generalizations and true laws of nature that intuitively involve some sort of physical necessity or probability. A universal generalization 'all ravens are black' and 'this is a raven' would not explain 'this is black' unless the generalization were a claim that there is some sort of causal connection between being a raven and being black—namely, that ravens must be black—of physical necessity. Similarly, we need to add that a statistical generalization ' n per cent of A 's are B ' (a statistical probability of $n/100$ of an A being B) does not explain a particular A 's being B unless it asserts some sort of causal connec-

tion between being *A* and being *B*. This will be so if it is claiming that each *A* has a $n/100$ physical probability of being *A*. By the physical probability of an event, it will be recalled, I mean a certain bias or tendency in nature. If nature is deterministic, the only physical probabilities in nature are probabilities of 1 (physical necessity) or 0 (physical impossibility). But, if there is a certain amount of indeterminism in nature, then there are physical probabilities between 1 and 0. When probabilistic generalizations are concerned with these, then we may call them probabilistic laws—most interpreters of Quantum Theory, for example, claim that the basic formulae of Quantum Theory are fundamental laws of this kind. In the latter case ‘*n* per cent of *A*’s are *B*’ together with ‘this is an *A*’ would (partially) explain ‘this is a *B*’ if its being an *A* made it physically probable to degree $n/100$ that that thing would be a *B*. Only so would there be some sort of causal connection between being *A* and being *B*, which we need if ‘*n* per cent of *A*’s are *B*’ is to explain an *A*’s being *B*. By contrast, John’s voting Conservative is not to be explained by the fact that his name appears on page 591 of the telephone directory and 70 per cent of those on that page vote Conservative. For the latter generalization just states how things happen to be; it is not to be understood as stating that being on that page pushes people in the direction of voting Conservative.

I shall in future call Hempel’s account amended in respect of explanation by probabilistic laws in the way described and amplified somehow or other so as to make a distinction between true generalizations and laws that involve physical necessity or probability the amended Hempelian account. But what this model amounts to depends on how we spell out the notion of a law of nature and so of the physical necessity or probability involved in a law. One view, originating from Hume, is the regularity view. On this view, ‘laws of nature’ are simply the ways things behave—have behaved, are behaving, and will behave. ‘All copper put in nitric acid dissolves under such-and-such conditions of temperature and pressure’ is a true universal law of nature if and only if all bits of copper when put in nitric acid under those conditions always have dissolved, now dissolve, and will dissolve. ‘50 per cent of atoms of C_{14} decay within 5,600 years’ is a true statistical law if and only if, taking the whole history of the universe, half the atoms of C_{14} have decayed within 5,600 years. We do, however, need the distinction between laws of nature, and accidental generalizations that are true

merely by accident.³ 'All spheres of gold are less than one mile in diameter' may be a true universal generalization, but it holds only in virtue of the accident of no civilization anywhere in the universe having put or going to put enough effort into constructing such a sphere. Regularity theory has reached a developed form that tries to take account of this distinction, in the work of David Lewis.

For Lewis, 'regularities earn their lawhood not by themselves, but by the joint efforts of a system in which they figure either as axioms or as theorems'.⁴ The best system is the system of regularities, which has (relative to rivals) the best combination of strength and simplicity. Strength is a matter of how much it successfully predicts (that is, whether it makes many actual events, past, present or future—whether observed or not—probable; and very few actual events improbable); simplicity is a matter of the regularities fitting together, and no doubt, each having internal simplicity in a way that Lewis does not, but no doubt could, spell out. The true laws are the regularities of the best system. Accidental generalizations are the regularities that do not fit into such a system. They float loosely without being derivable from more fundamental regularities. So 'all spheres of gold are less than one mile in diameter', even if true, is probably not a law, because it does not follow from the best system—as is evidenced by the fact that it certainly does not follow from our current best approximation to the ultimate best system—a conjunction of Relativity Theory and Quantum Theory. Similarly with probabilistic laws—if and only if '90 per cent of *A*'s are *B*' were a consequence of the best system of regularities would it be a law of nature. If (and only if) it follows from such a best system that a particular *A* will be followed by a particular *B* (and certain other complicated conditions hold), then that *A* causes that *B*. Lewis's account of laws of nature is part of his campaign on behalf of 'Humean supervenience', that everything there is supervenes (logically) on 'a vast mosaic of local matters of particular fact', which he interprets as a spatio-temporal arrangement of intrinsic properties,

³ Hempel makes the distinction in terms of 'laws' being those true generalizations that are grounded in accepted theory. But this is unsatisfactory, since some laws may never be discovered and so never become part of accepted theory.

⁴ David Lewis, 'A Subjectivist's Guide to Objective Chance', Postscript, in *Philosophical Papers* (Oxford University Press, 1986), ii. 122.

or 'qualities'.⁵ Laws of nature and causation are for Lewis among the things thus supervenient.

There seem, however, to be overwhelming objections to any Humean account, including Lewis's, if laws of nature are supposed to explain anything—and in particular to explain whether and why one thing causes another, as Humeans suppose that they do. For, since whether some regularity constitutes a law depends, on this account, nor merely on what has happened but on what will happen in the whole future history of the universe, it follows that whether *A* causes *B* now depends on that future history. Yet, how can what is yet to happen (in maybe two billion years' time) make it the case that *A* now causes *B*, and thus explain why *B* happens? Whether *A* causes *B* is surely a matter of what happens now, and whether the world ends in two billion years' time cannot make any difference to whether *A* now causes *B*? What is yet to happen can make no difference to what is the true explanation of why *B* occurs (namely, that *A* occurred and caused *B*)—though, of course, it might make a difference to what we justifiably believe to be the true explanation. (Put another way, that some proposed explanation is the simplest explanation of the data, past and future is evidence that it is the true explanation; but it does not constitute it being the true explanation.) Further, it is because of their role in causation that laws of nature are said to generate counterfactuals. Suppose that 'all copper expands when heated' is a law of nature, but that I do not heat a certain piece of copper; it is all the same fairly evidently the case that, 'if that copper had been heated, it would have expanded'. But, if a law simply states what does (or did or will) happen, how can it provide any ground for asserting the counterfactual? It would do that only if there were some deeper kind of necessity built into it than that provided by fitting into a system. Fitting into a system could be evidence only of that deeper kind of necessity.

So, dismissing Humean accounts of laws for good reason, let us consider alternative accounts of the physical necessity (and physical probability) involved in laws of nature that do not analyse it away in terms of patterns of actual events. Physical necessity may be thought of either as separate from the objects that are governed by it, or as a constitutive aspect of those objects. The former approach leads to a

⁵ *Ibid.*, pp. ix–x.

picture of the world as consisting of events (constituted by substances having, gaining or losing properties), on the one hand, and laws of nature (involving physical necessity or probability), on the other hand; and it can be developed so as to allow for the possibility of there being universes in which there are no events, but merely laws of nature.⁶ Laws of nature are thus ontologically concrete entities.

The version of this account that has been much discussed recently is the version that claims that laws of nature are relations between universals.⁷ (Universals are properties that can be fully instantiated in many different objects. Thus 'brown' is a universal, because innumerable different things can be brown.) It being a fundamental law of nature that 'all photons travel at 300,000 km/sec. relative to every inertial reference frame' consists in there being such a connection between the universal 'being a photon' and the universal 'travelling at 300,000 km/sec. relative to every inertial reference frame'. These universals are tied together, but the tie is not a logically necessary one—that is, it is not on this view part of what it is to be a photon that it travels at that speed. But it is physically necessary, and the physical necessity is a matter of the two universals being tied together. One can perhaps begin to make sense of this suggestion if one thinks of the causing of states of affairs (for example, the bringing into existence of a photon) as making properties, which are universals, to be instantiated; and this involving the bringing of them down to Earth from an eternal Heaven, together with whatever is involved with those universals—namely, other universals (for example, travelling at 300,000 km/sec.) connected thereto. But why should we believe that there is such a Platonist heaven in which universals are tied together? And how can universals act on the world? This is a very

⁶ Thus 'I hold . . . that many empty [possible] universes exist. As I see it, there is a world devoid of all material objects and events in which the general principles of Newtonian mechanics are laws; there is another empty world in which the general principles of Aristotelian physics are laws' (John W. Carroll, *Laws of Nature* (Cambridge University Press, 1994), 64 n. 4).

⁷ See e.g. D. M. Armstrong, *A World of States of Affairs* (Cambridge University Press, 1997), Michael Tooley, 'The Nature of Laws', *Canadian Journal of Philosophy*, 7 (1977), 667–98, and F. I. Dretske, 'Laws of Nature', *Philosophy of Science*, 44 (1977), 248–68. In expounding this account, I do so in terms of Tooley's view that the ties between universals exist in a Platonist heaven before and independently of any instantiation in mundane objects, rather than Armstrong's view that the universals exist only in so far as they are instantiated. On the latter view it seems inexplicable why any particular tie was instantiated, rather than any other one.

mysterious causal relation between the non-spatio-temporal world and our world for which we have no analogue.

The alternative to thinking of the physical necessity involved in laws of nature as separate from the objects governed by it is to think of it as a constitutive aspect of those objects. The way in which this is normally developed is what we may call the substances-powers-and-liabilities (S-P-L) account of laws of nature. The 'objects' (the 'what') that cause are individual substances—this planet, those molecules of water. They cause effects in virtue of their powers to do so and their liabilities (deterministic or probabilistic) to exercise those powers under certain conditions, often when caused to do so by other substances. Powers and liabilities (the 'why') are thus among the properties of substances. Laws of nature are then just regularities—not of mere spatio-temporal succession (as with Hume), but regularities in the causal powers (manifested and unmanifested) of substances of various kinds. That heated copper expands is a law is just a matter of every piece of copper having the causal power to expand, and the liability to exercise that power when heated. As a matter of contingent fact, substances fall into kinds, such that all objects of the same kind have the same powers and liabilities. The powers and liabilities of large-scale things (lumps of copper) derive from the powers and liabilities of the small-scale things that compose them (atoms; and ultimately quarks, electrons, etc.). And, given a satisfactory theory integrating all science, all ultimate particulars will have exactly the same powers and liabilities (for example, the power to cause an effect proportional in a certain way to their mass, charge, spin, etc., and the liability to exercise that under conditions varying with the mass, charge, spin, etc., of other objects).

This account of the ultimate determinants of what happens as merely substances and their causal powers and liabilities provides explanation of what happens in familiar terms. As I shall consider more fully shortly, we ourselves have causal powers that we, unlike inanimate objects, can choose to exercise. The S-P-L way of explaining things was the way familiar to the ancient and medieval world, before talk of 'laws of nature' became common in the sixteenth century. It was revived by Rom Harré and E. H. Madden in *Causal Powers*.⁸ When talk of 'laws of nature' became common in the sixteenth century, they were supposed to be God's laws for nature,

⁸ (Blackwell, 1975).

and so such talk has its natural place in a theistic world view. But if there is a God and he makes things in the world behave as they do, he surely operates not directly, but by sustaining the laws of nature—which means, on this account, by determining which powers and liabilities substances have, and conserving those powers and liabilities in substances. The basic structure of explanation in terms of substances, powers and liabilities does not presuppose that there is a God who operates in this way.

With the S–P–L account, unlike with the Humean account and the universals account, we have moved away from the Hempelian structure of scientific explanation in a crucial respect. For ‘laws of nature’ no longer play any causal role in explaining particular phenomena. What causes the expansion of a particular piece of copper is that piece of copper, its power to expand and its liability to exercise that power when heated. The regularity involved in other pieces of copper having similar powers and liabilities is no part of the explanation. While causation is essentially involved in laws, laws are not essentially involved in causation. The S–P–L account of laws of nature and of the explanation of particular events seems to me more satisfactory than the other accounts. The regularities in the causal powers and liabilities of particular substances, and so in their behaviour, which constitute the ‘laws of nature’, do entail that particular substances will have particular powers and liabilities; and so, any evidence that makes it probable that such and such (e.g. ‘all *A*’s do so-and-so in circumstances *C*’) is a law of nature is evidence that makes it (inductively) probable that a particular instance of it holds (for example, that this *A* has the power to do such-and-such and the liability to exercise it in circumstances *C*). But the law does not explain why these substances have those powers and liabilities. And so the S–P–L account raises the question of why so many substances have similar powers and liabilities to each other (why does each substance in the universe have the power to attract each other substance in the way stated in, for example, Newton’s ‘laws’), and we will return to that question in Chapter 8. But, as we shall see, a question the same in essentials arises also on the other accounts of laws of nature; and the argument of this book does not depend on my preferred account of laws of nature and so of scientific explanation. So I shall normally operate simply with the amended Hempelian model without presupposing how it is to be spelled out, or corrected. However, at crucial points I shall draw the reader’s attention to

alternative accounts of laws of nature and so of scientific explanation, and especially to the S–P–L account. I now move on to contrast scientific explanation with personal explanation.

Personal Explanation

The other pattern of explanation that we use all the time in explaining mundane phenomena is what I shall call personal explanation. In personal explanation the occurrence of a phenomenon *E* is explained as brought about by a rational agent *P* doing some action intentionally. The central case of this, with which we shall be primarily concerned, is where *P* brings about *E* intentionally—that is, brings *E* about, meaning so to do. The other case is where *P* brings about *E* unintentionally in consequence of doing something else intentionally—we shall come to this case briefly later. In the central case *E* occurred because *P* meant *E* to occur through what he was doing. What an agent meant to occur through his agency may be called the intention (or purpose—I shall use these terms interchangeably) *J* in the agent's action, for example, that *E* occur. *E* is then explained by *P* having intention *J*. *E* may be the motion of my hand, *P* be myself, and *J* my intention that *E* occur. *E* is then what I shall call the result of an intentional action *A* of bringing *E* about.⁹ In the example cited, *A* is my moving my hand. However, *E* is only partially explained by *P*'s having intention *J*. For a person may have the intention to bring about some effect and yet fail to do so. I may mean my hand to move through my agency, and yet the hand may fail to move because someone is holding it down; in consequence the only action that I perform is that of trying to move my hand. If *E* does result from *P* and *J*, a full explanation will tell us why, how it was that *P*'s intention was efficacious.

This leads us to the well-known distinction¹⁰ among intentional actions between basic actions and mediated actions. Roughly

⁹ The occurrence of the result of an action is thus entailed by the performance of the action. The result of an action is to be distinguished from a consequence of the action. A consequence of an action is something brought about but not entailed by performance of the action. This distinction is due to von Wright. See G. H. von Wright, *Norm and Action* (Routledge & Kegan Paul, 1963), 39 ff.

¹⁰ First set out by A. C. Danto in his 'Basic Actions', *American Philosophical Quarterly*, 2 (1965), 141–8.

speaking, a basic action is something that an agent just does, does not do by doing anything else. A mediated action is an action that is not a basic action, one that an agent does by doing something else. I signal by moving my hand. I break the door down by giving it a kick. The former is a mediated action; the latter a basic action. Now, if bringing about *E* is a basic action, the answer to the question how it was that *P*'s intention was efficacious will simply be that bringing about *E* was among the basic powers or capacities *X* that *P* had at that time—that is, was among the basic actions that *P* could do at will (that is, would succeed in doing if he formed the intention to do them). Bringing about the motion of our arms or legs, lips, eyes or eyebrows, etc., is for most of us most of the time among our basic powers. *E* is fully explained when we have cited the agent *P*, his intention *J* that *E* occur, and his basic powers *X*, which include the power to bring about *E*; for, given all three, *E* cannot but occur. Of course, often in such cases it is so obvious why *E* occurred that we do not bother to give the explanation, but the explanation is true nevertheless. We may not bother to comment, when someone is walking along, that his legs moved because he moved them (that is, that he brought about their motion, meaning so to do), but it is true nevertheless. Sometimes, however, this sort of explanation is not at all obvious—it may on occasion be the explanation of a person's ears wiggling or her heart stopping beating that she brought about these things intentionally.

If bringing about *E* is a mediated action, the answer to the question how it was that *P*'s intention was efficacious will be more complicated. It will be that *E* was the intended consequence of some basic action of *P*'s, *A*—that is, a consequence that *P* meant to occur through his performing *A*, which consists in bringing about some state of affairs *S*. *P* has the intention *J* that *E* occur as a consequence of the occurrence of *A* (and so *J* contained within it the intention that *S* occur). For *P* to have this latter intention, he must believe that his doing *A* will (no less probably than his doing any other basic action) have his bringing about *E* as a consequence (normally perhaps by *S* causing *E*). The explanation of how *P*'s intention was efficacious is that bringing about *S* is among *P*'s basic powers *X*, and that the bringing-about of *S* had as a consequence the occurrence of *E*. There will often be a scientific explanation of the latter. *S* may cause *E* in accord with natural laws *L*, because it is a consequence of *L* that in circumstances *D* (which in fact hold) *S* is followed by *E*. So, in this case, *P*, *J*, *X*, *D*, and *L* fully explain the

occurrence of *E*. *E* is brought about by *P* having a certain intention *J*, which in consequence of his basic powers *X* brings about some state of affairs *S* in circumstances *D*, which laws of nature *L* then ensure will bring about *E*. Thus a full explanation of the door being flat on the ground is that I, exercising my basic powers, brought it about that my foot moved quickly into contact with the door, meaning this to occur and it to cause the door to be flat; the door hinges, the mass of the door, the mass and velocity of my foot were in fact such that it was a consequence of the laws of mechanics that the impact of my foot with the door was followed by the flattening of the door. In the above analyses I use the word 'consequence' in a wide sense. The connection between *A* and *E* may be either causal or logical. This may be as in the above example, because *S*, the result of *A*, causes *E*. It may also be because, given current circumstances *D*, the performance of *A* constitutes bringing about *E*. Thus, given conventions in banking and motoring, my writing my name in a certain place has as a consequence that a cheque bears my signature, and my putting my arm out of the car window has as a consequence that a signal indicating a turn to the right is made.

So, to summarize, in the central case of personal explanation we explain a phenomenon *E* as brought about intentionally by a rational agent *P*. If the bringing-about of *E* is a basic action *A*, we need to cite further an intention *J* of *P* that *E* occur and to state that bringing about *E* is among the things that *P* is able to do at will—namely, among *P*'s basic powers *X*. *P*, *J*, and *X* provide a full explanation of *E*. Of course, we can often go further and explain how it is that *P* has intention *J* (for example, by stating that he formed this intention in order to forward some wider intention, as when we explain that he formed the intention to sign a cheque because he had the intention to pay you money). Or we can explain how it is that *P* has those powers (for example, by stating which nerves and muscles need to be operative for *P* to have these powers). But *P*, *J*, and *X* suffice to explain *E*—whether or not we can explain how it is that *J* and *X* hold. If the bringing-about of *E* is a mediated action, things are more complicated. We cite *P* and his intention *J* to bring about *E* as a consequence of a basic action *A*; we explain that the performance of *A* was among *P*'s basic powers, and we explain how the performance of *A* had *E* as a consequence. Again, the occurrence and operation of the factors cited here may themselves be explained further; but they do not need to be for us to have a full explanation. When there is only a

basic action involved, the agent *P* is the cause of the effect; his intention and powers provide the reason for the efficacy of the cause. Where the action is a mediated action, further factors are added. The two figures at the top of the diagram on p. 39 summarize these results for basic actions, and for mediated actions in cases where a natural law *L* brings it about that *S* has *E* as a consequence. Causes and the conditions for their operation (the 'what') are shown to the left of the arrows: reasons (the 'why') are shown above the arrows; effects are shown to the right of the arrows.

There is, I claimed earlier, a second kind of personal explanation. Here we explain the occurrence of *E* as brought about unintentionally by a rational agent *P* bringing about something else intentionally; *E* is an unintended consequence of an intentional action. For example, in standing up I may unintentionally knock over a cup. Here the knocking-over of the cup is caused by my occupying a certain standing position, which was a state of affairs brought about intentionally by me. I did not mean the cup to be knocked over, but, given the circumstances (the original position of the cup, etc.), my occupying the standing position causes the knocking-over of the cup in virtue of the laws of mechanics *L*. My concern henceforward will be only with the central case of intentional action where the effect is brought about intentionally.

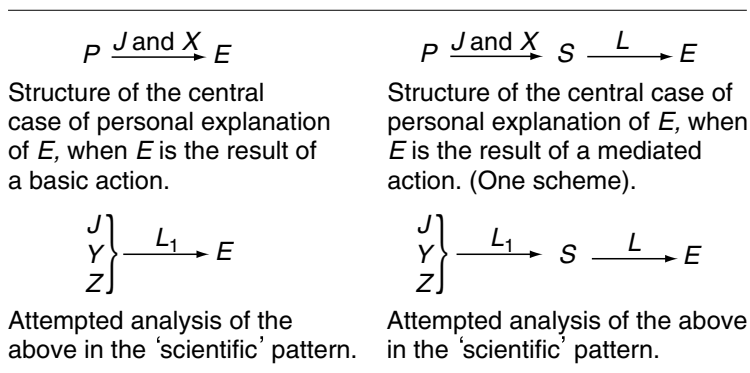
Personal Explanation Unanalysable in Terms of Scientific Explanation

Personal explanation looks very different from scientific explanation. In scientific explanation in the amended Hempelian model we explain an event *E* by past events or states *C* and natural laws *L*. In personal explanation we explain *E* as brought about by an agent *P* (not by an event or state) in order to realize intentions for the future. Despite the apparent difference, it has, however, been argued by some philosophers, seminally by Donald Davidson¹¹ and by many others at greater length, that really personal explanation conforms to the scientific pattern. In my terminology and using the Hempelian

¹¹ D. Davidson, 'Actions, Reasons, and Causes', *Journal of Philosophy*, 60 (1963), 685–700.

model of scientific explanation, a Davidson-like suggestion amounts to the following.

Suppose, first, that *E* is the result of a basic action. Then, to say that *P* brought *E* about intentionally is just to say that an event involving *P*—that is, *P*'s intention that *E* occur—*J*, brought it about. To say that *P* had the power to bring about *E* is just to say that *P*'s bodily condition *Y* (brain states, muscle states, etc.) and environmental conditions *Z* (no one having bound *P*'s arm, etc.) and laws *L*₁ are such that an intention¹² such as *J* is followed by the event intended, *E*. We then have a scientific explanation as set out in the diagram.



Suppose, next, that *E* is the result of a mediated action. Then to say that *P* brought it about is to say that an event involving *P*—that is, *P*'s intention *J*—under the current bodily and environmental conditions *Y* and *Z* brought about (in accordance with laws *L*₁) the result of the basic action *S*, which had as a consequence *E*. We have seen that there are different ways in which *S* may have *E* as a consequence. One way is that *S* may bring about *E* in accord with the normal scientific pattern of causation—that is, in virtue of some law of nature *L*. This is the scheme depicted in the diagram. The other way in which *S* may

¹² Davidson holds that mental events such as 'intentions' are identical with brain events, and that the laws involved are laws connecting these brain events (under a physical description, and not a mental description such as 'intentions') with other physical events. I shall shortly give reason for adopting a system of categories that rule out the postulated identity, from which it follows that if there are laws involved here they are psycho-physical laws.

have *E* as a consequence can also, it is suggested, easily be fitted into the scientific pattern of explanation. So, on this reductionist view, personal explanation is in essence really scientific explanation. There are not explanations of events of two kinds—only explanations of one kind. Events brought about by actions are just those that include intentions among their causes.

In order to show what is wrong with this, I wish to make two points—first, that the intention in an action that an agent is performing is not the same as any brain event that might be connected with it; and, secondly, that having an intention (in the sense with which we are concerned¹³) is not a passive state of an agent, but just is the agent exercising causal influence (which will cause the effect intended if and only if the agent has the requisite power).

I understand by a substance a thing (other than a property) that has properties; tables, planets, atoms, humans, and other persons are all substances. (Sometimes, when there is no danger of misunderstanding, I use ‘object’ as a synonym for ‘substance’.) Substances have properties—that is, characteristics that can characterize them and other substances as well. In this sense all properties are universals; brown is a property, and different things can be brown. Properties include both monadic properties, which characterize individual substances, and relational properties, which link two or more substances. Being square, weighing 10 kilos, or being-taller-than are properties, the former two being monadic properties, the latter being a relational property that relates two substances (one thing is taller than another thing). On these definitions there is no more to the history of the world than substances coming into existence, gaining and losing properties (including relations to other substances), and then ceasing to exist. It is useful to have a word for these things such that there is no more to the history of the world than all these things; and a natural word to choose for that category of thing is the word ‘event’. I propose to use it in this sense: that an event consists in the instantiation of a property in a substance (or substances, or in events) at a time or the coming into being,

¹³ It is important to distinguish the intention in an agent’s action or the intention with which he acts, which is what concerns us here, from an intention to do something at some future date. The latter is not something manifested in action and may be a state of some sort. The former exists only in so far as an agent performs some intentional action—even if only the minimal action of trying to do something. The account in the text is meant to apply only to intentions in actions.

or ceasing to exist, of a substance. Events include the table being square now, or John being taller than James on 30 March 2001 at 10.00 a.m., or me coming into existence on 26 December 1934. In order to fulfil the purpose of the definition of 'event', we need so to individuate properties that, if you knew which properties had been instantiated in what when, you would know (or could deduce) everything that had happened. This will involve, for example, counting being red and reflecting light of such-and-such a wavelength as different properties—for you could (just by looking at it in normal light) know that something was red without knowing (or being able to deduce) that it reflected light of such-and-such a wavelength, and conversely.

It follows immediately that having an intention cannot be the same event as having any brain event, for you could know that someone was intending to do such-and-such in his action without knowing that he was in a particular brain state or any brain state, and conversely. These are two different events connected with a subject, even if perhaps of physical necessity they always go together. It is true that other criteria for two events being the same event might yield a different result—that the two events were the same; but then, to tell the whole history of the world on those other criteria, it would not be enough to know that some event (for example, some brain state) had taken place; you would need to know that it had two different somethings, say 'characteristics'—a brain characteristic, and an intention characteristic—associated with it. Some sort of dualism is unavoidable here, and I suggest that my proposed use of the word 'event' provides a neat system of categories by using which we can describe the world fully, a system of categories not too distant from ordinary usage.

So intentions are not brain events, even if closely connected with brain events. In the sense to be defined in Chapter 9, they are mental events. The next issue is what sort of mental events are they. Is having an intention a passive state, some state in which the agent finds himself—like having a sensation or a belief? Davidson thinks of intentions as 'desires' and let us read him (despite what he writes elsewhere) as supposing that these are mental events distinct from brain events. These desires may need some event such as a perception or a neural event, to make them cause other events. Then, he claims, actions are events that have passive mental states, desires for their occurrence, among their causes. Personal

explanation is analysable in terms of the production of effects by such desires.

Despite the fact that it is the most plausible form of reductionist theory, like all the others, Davidson's theory is open to a fatal objection. The basic idea of all such theories is that an agent's bringing about an effect intentionally—that is, meaning so to do—which is how we defined the agent's bringing about an effect having an intention so to do—is to be analysed as the causing of that effect by some passive state of the agent or some event involving him. But all such analyses fail because, if an intention (or wish or desire) of *P* to bring about *E* is some passive state or event, it could bring about *E* without *P*'s having intentionally brought about *E*. Causation by an intention (so understood) does not guarantee intentional action.

The classic objection to the reductionist theory was formulated as follows by Richard Taylor. Here the causal factor is termed a 'desire', but it could equally well be termed a 'want' or an 'intention'.

Suppose . . . that a member of an audience keenly desires to attract the speaker's attention but, being shy, only fidgets uncomfortably in his seat and blushes. We may suppose, further, that he does attract the speaker's attention by his very fidgeting; but he did not fidget *in order* to catch the speaker's attention, even though he desired that result and might well have realized that such behaviour was going to produce it.¹⁴

Here we have a case of a desire for *E* causing *E*, and yet there is no action. The basic point is that desires, wants, etc. may occur and yet the agent for some reason may not act to fulfil his desire or want. Nevertheless, in such a case, possibly without the agent's knowledge, the desire may bring about the intended effect—without the agent's bringing about the effect intentionally. An agent's bringing something about intentionally is not analysable as his intention bringing that thing about, if an intention is supposed to be a passive mental event or state. The same applies if we substitute, for 'intention', 'desire', 'want', or any similar term.¹⁵ So a Davidson-type analysis

¹⁴ R. Taylor, *Action and Purpose* (Prentice-Hall, 1966), 248–9.

¹⁵ A response by Alvin Goldman (*A Theory of Human Action*, (Prentice-Hall, 1970)) is to admit that intentional acts have to be caused by 'action plans' or desires 'in a certain characteristic way', and to claim that in Taylor's example we do not have a case of that way. And what is this characteristic way? Goldman writes: 'To this question, I confess, I do not have a fully detailed answer. But neither do I think that it is incumbent on me, *qua*

seems to fail. To say that *P* brought something about intentionally is not to say that some passive state of *P* or event involving *P*, such as an intention, brought that thing about. There seems to be no other plausible way in which personal explanations can be analysed into the scientific pattern, and so it would appear that personal explanation is of a distinct type from scientific explanation (on the amended Hempelian model of the latter). (Note that in future I shall understand by a 'desire' to do some action a causally influential inclination to do the action, which may or may not coincide with a judgement by the agent that it is overall good to do the action. If it does not coincide, the agent has to choose whether to resist his desire or yield to it.)

If intentions are not states or events that happen to an agent, they must themselves be actions. Having an intention is not something that happens to an agent, but something she does. For me to have the intention in acting of moving my hand is to do what (if I were to fail or find it difficult to move the hand) would be called 'trying' to move my hand. In the past, having such an intention has been given the technical name of making a 'volition' to move my hand. The basic mistake that reductionist analyses make is (in the terminology introduced at the beginning of the chapter) to treat intentions as belonging to the 'what' rather than to the 'why' of explanation. When one explains an occurrence as brought about by an agent having some intention, one is not by the word 'intention' describing some occurrent state or event that caused the occurrence, but one is stating that the agent brought about that occurrence and did so because he meant to do so. To act intentionally is to exercise causal agency in a certain direction, which will succeed in producing the intended effect if the agent has the requisite power. An intention—to avoid a puddle, say—explains why at a certain time a man with normal basic powers (and that involves, physically, a normal brain and the operation of normal psycho-physical laws) behaved as he did, made such

philosopher, to give an answer to this question. A complete explanation of how wants and beliefs lead to intentional acts would require intensive neurophysiological information, and I do not think it fair to demand of a *philosophical* analysis, that it provide this information' (ibid. 62). But this really will not do. For hundreds and hundreds of years people have been able to distinguish, among cases where wants cause the events wanted, those cases where an action was performed. We have distinct concepts here that we know how to apply. It is indeed incumbent upon a philosopher to analyse the difference—although it is not up to him to say which neurophysiological goings-on are physically necessary to produce a case of an action being performed.

movements as in fact led to his feet bypassing the puddle. That this account is correct is brought out by the linguistic fact that explanations in terms of intentions can easily be paraphrased in terms of explanations in which there occur no nouns that could conceivably be regarded as denoting occurrent states or events. To say that a man's intention in making certain movements was to avoid the puddle is to say that he made them in order to avoid the puddle, or so as to avoid the puddle. But no such paraphrase is possible for the initial conditions which are cited in normal scientific explanations.

Although intentions, like laws of nature, belong to the 'why', the reasons, of explanation, there are, of course, vast differences between laws of nature and intentions. Intentions are such that necessarily the agent whose they are 'goes along' with them, is aware of them, and has privileged access to them in the sense that he is in a better position than outsiders to know about them. Laws of nature are not necessarily known to anyone, nor necessarily does any person 'go along with them' or have privileged access to them. But that the 'why' is here known and adopted by an agent is one of the differences between personal and scientific explanation. The other main difference is that, in personal explanation, talk about a substance which explains, namely a person, is not reducible to talk about occurrent states of or events involving that person. The contrast between scientific and personal explanation remains even on the S-P-L account of the former, although the two patterns are much closer to each other on this account. That is a reason for preferring the S-P-L model; it brings out that both personal and scientific explanation are species of the same genus, causal explanation. In both, the cause is a substance or substances. In both, the substance has certain powers, and produces the effect in virtue of its powers. The difference is that, in scientific explanation, the substance has liabilities to exercise its powers under certain circumstances; it is either physically necessary or probable that it will, and it has no intention or purpose in doing what it does; whereas in personal explanation the substance (the person) acts intentionally, doing the action that—given its beliefs—will most probably fulfil its intentions. There is no parallel for that in the scientific case. In consequence, even if scientific explanation can be expressed in terms of an event (the substance being in certain circumstances) rather than the substance itself causing the effect, personal explanation cannot be expressed in this way. A person causing an effect is not analysable

as a passive state of that person or an event involving that person causing the effect.

Can there be Two Explanations of a Phenomenon?

So far in this chapter I have been concerned to characterize the structures of the two types of explanation that we use in explaining the occurrence of phenomena, and to show how they differ from each other. I now turn to the question of whether there can be only one true explanation of some phenomenon. For, if so, then, if there is a personal explanation of some phenomenon, there cannot be a scientific one, and conversely. I suggest that there can be two true distinct explanations of some phenomenon *E*, if one or other of three conditions is satisfied, but that otherwise there cannot be.

Clearly there can be two true distinct explanations of *E*, when one or other or both are partial explanations of *E*. For the one may combine with the other to make a fuller explanation. Thus a man's death from cancer may be explained by (1) his smoking and a law about the proportion of smokers who die from cancer, and by (2) his parents' having died of cancer and a law about the proportion of those whose parents die of cancer who themselves die of cancer. Since (1) and (2) only make probable but do not necessitate the man's death from cancer, they are only partial explanations. Clearly they can be combined into a fuller explanation in terms of the man's smoking and his parents having died of cancer and the proportion of those who smoke and whose parents have died of cancer who die of cancer.

But can there be two different full explanations of a phenomenon? The answer is still yes—if the occurrence of the causes (the 'what') and the operation of the reasons (the 'why') cited in one explanation are to be explained at least in part by the occurrence of the causes and the operation of the reasons cited in the other explanation. For example, the present position of Mars is explained by its position in the last few days and the laws of planetary motion, formulated more or less correctly by Kepler. Where it has been recently and the laws stating how planets move determine where Mars will be today. Yet the present position of Mars is also explained by its position and velocity last year and those of all other heavenly bodies, and Newton's laws of motion. Newton's laws state how material bodies change their

velocities under the influence of other bodies. Both are full explanations, and yet they are clearly compatible. This is because Newton's laws and the positions and velocities of the planets explain their (approximate) conformity to Kepler's laws. Kepler's laws operate because Newton's laws operate and the sun and the planets have the initial positions and velocities that they have, and are far distant from other massive bodies. It is for this reason that the motion of a human hand is often explicable by both personal and scientific explanation. The motion of my hand may be fully explained by goings-on in the nerves and muscles of my arm, and physiological laws. It may also be fully explained by me bringing it about, having the intention and power so to do. Yet in this case the causes and reasons cited in each explanation provide a partial explanation of the occurrence and operation of the causes and reasons cited in the other. The goings-on in my nerves and muscles are brought about unintentionally by my bringing about the motion of my hand intentionally. Also, the operation of physiological laws provides part of the explanation of my having the power to move my hand—only because nervous discharges are propagated as they are, am I able to move my hand. So there is here a twofold reason why two full explanations can each fully explain the motion of my hand.

But can there be two distinct full explanations of some phenomenon *E*, when neither in any way explains the occurrence or operation of the causes and reasons involved in the other? Yes, again, so long as there is overdetermination. In overdetermination each of the full explanations gives causes and reasons sufficient for the occurrence of the effect, but neither pair on its own is necessary since the other pair would have produced the effect on its own. If someone dies as a result of being poisoned by *A* at the same time as he is shot by *B*, we have such overdetermination. But such coordination will be a coincidence, barring a common cause of the actions of *A* and *B* (for example, *C* who employed both *A* and *B* to murder the same victim in order to ensure that he really died). It cannot be necessary for the production of the effect to have two distinct full explanations, when neither in any way explains the occurrence or operation of the causes and reasons involved in the other; unless the occurrence and operation of the causes and reasons involved in both are explicable, at least in part, by the causes and reasons of a third full explanation (a common cause). It follows, given that scientific and personal explan-

ation are the only two possible kinds of explanation,¹⁶ and barring accidental overdetermination, that there can be a full personal explanation and a full scientific explanation of some phenomenon only if one in part explains the occurrence and operation of the components of the other—either the scientific explanation at least in part explains the causes and reasons in the personal explanation, or conversely; or there is a further full explanation (either personal or scientific) that explains the causes and reasons operative in both the other explanations.

Explanation by the Action of God

In this chapter so far I have been concerned to analyse the structure of personal explanation, and to show its relation to scientific explanation. I have done this because, when the theist claims that the action

¹⁶ Three recent writers have suggested that there is a third possible kind of causal explanation of phenomena (axiarchic explanation) and that it is in terms of an explanation of this kind that the phenomena that I shall discuss in subsequent chapters are to be explained. This is that phenomena come into existence because it is good that they should exist. See John Leslie, *Value and Existence* (Blackwell, 1979); Derek Parfit, 'The Puzzle of Reality: Why does the universe exist?', *Times Literary Supplement*, 3 July 1992, repr. in P. Van Inwagen and D. W. Zimmerman (eds.), *Metaphysics: The Big Questions* (Blackwell, 1998); and Hugh Rice, *God and Goodness* (Oxford University Press, 2000). In this book I am proposing a personal explanation of the existence of the universe with its various characteristics in terms of a person, God, who brought them about because he believed them good. But the suggestion of Leslie and others is not this, but that there is an impersonal principle at work bringing into being good things because they are good. This is not a personal explanation, and neither is it a scientific explanation—for (in the Hempelian model) laws of nature operate on already existing states of affairs, and (in the S–P–L model) if substances are to cause states of affairs, those substances have already to exist. The suggestion of Leslie and others is that the axiarchic principle operates to bring forth good things out of nothing. The trouble with this suggestion is that, while there are innumerable instances of mundane phenomena rightly explained by a personal or scientific explanation (that is in terms of the action of ordinary persons or inanimate substances), these are no mundane examples of anything coming into existence because it is good that it should. Food never appears on the tables of the hungry because it is good that it should, but only because some person puts it there because he believes it to be good that it should be there. So we have no criteria that we can extrapolate from mundane situations for judging when an explanation of this kind is probably true and when it is not. We have criteria for judging when purported scientific or personal explanations are or are not probably true, which I set out in Chapter 3. But in the absence of criteria for judging the worth of an axiarchic explanation of the existence of the universe, we can have no grounds for supposing that such an explanation is probably true.

of God explains various phenomena, such as the existence and orderliness of the world, he is proposing a personal explanation of these phenomena. However, personal explanations of phenomena by the action of God differ from most mundane personal explanations in two important respects, on which I must now comment in conclusion of this chapter.

The first is that a personal explanation of the occurrence of a phenomenon *E* in terms of God's bringing it about, meaning so to do, cannot be even in part explicable scientifically. We have seen that a personal explanation may often, at any rate in part, be explained by a scientific explanation—and conversely. Thus a person having the powers that she has may be explained in part by her having nerves and muscles and by the operation of various physiological laws. Her having the intentions that she has may also be given a scientific explanation, and perhaps a human's existence may also be explained in this way. The fact that personal explanation cannot be analysed in terms of scientific explanation does not mean that its operation on a particular occasion cannot be given a scientific explanation. However, it seems coherent to suppose that there should be a personal explanation of the occurrence of some event *E* by the agency of an agent *P* having the intention *J* to bring about *E* and the power so to do, without all this being in any way susceptible of a scientific explanation. To start with, an agent might have the power to perform certain basic actions without his having that power being dependent on any physical states or natural laws. His capacity to perform these actions might be an ultimate brute fact (or only explicable by another personal explanation). Likewise, an agent's having the intentions in actions that he does, his choice of intentional actions, may not be susceptible of scientific explanation.

To see the above, note that there is at present with respect to some of the intentions that we form no plausible scientific explanation of why we form these intentions, rather than any other ones. And yet our explanations of other things in terms of these intentions would still be explanations even if there was no explanation of why we formed these intentions. Then we have basic powers to bring about mental images of different geometrical shapes. There might be a partial scientific explanation of my having this power in terms of my brain being in a certain state giving me this power. Yet there is no contradiction in supposing that powers of visualiza-

tion are not dependent on the brain, or on anything else. Maybe we just have such powers. But that would not affect the fact that my having a certain mental image could be explained by my basic power to produce such images. And, although we normally suppose (correctly) that there is a scientific explanation of the existence of this body that is mine, there is no scientific explanation of how it comes about that this body is mine (rather than someone else's) and so no scientific explanation of my existence at all. For this world could have been the same in all its physical aspects, and yet a different person could have operated through this body. (I develop this point more fully in Chapter 9.) And yet the fact that science cannot explain my existence does not mean that there is no true explanation of things in terms of me bringing them about. Personal explanation may explain without there being a scientific explanation of the occurrence and operation of the factors involved in it.

When the theist claims that the existence of the world and its various features is to be explained by the action of God bringing these things about meaning so to do, he will claim that God's action cannot be explained scientifically, even in part. God is supposed to be perfectly free. God's existence and powers do not depend on the states of the physical world or the laws of its operation—rather, vice versa. Nor are God's intentions scientifically explicable. But all this, as we have now seen, does not in any way weaken the explanatory value of the personal explanation. God's bringing about some event may be explicable by a wider personal explanation. He may bring about *E* in order thereby to bring about *F*; *F* may be an event that takes a considerable period of time, and *E* may be the first stage of *F*. But the theist claims that this kind of explanation is the only kind of explanation of God's actions that can be provided. God's own intentions alone explain his doing what he does. God's basic actions are supposed to include creating the universe *e nihilo* (that is, not out of existing matter), keeping it in existence, making things behave in accord with natural laws, and occasionally intervening in the universe (sometimes by setting those laws aside). Creating matter *e nihilo* is not something that humans are able to do, but it is easy enough to conceive of their doing it. It is logically possible that I could just find myself able as easily to make appear before me an inkwell or to make a sixth finger grow, as I am at present able to move my hand. Various tests (for example, sealing off the room and keeping its contents carefully weighed) could show that the inkwell

or finger were not made of existing matter. Creating *e nihilo* is a perfectly conceivable basic act.

The other important respect in which personal explanations of phenomena by the action of God differ from most mundane personal explanations is that God is supposed to be a person without a body—that is, a spirit. It is important to make clear at this stage what it is for a person not to have a body. We can best do so by asking a different question—what is it that I am saying when I say that this body, the body behind the desk, is *my* body? First, that I can move, as basic actions of mine, many parts of this body, whereas I can make a difference to anything else only by moving parts of this body. To move the arm over there (your arm), I have to grasp it with this arm, but I can move this arm straight off. Secondly, my having a mental life of thought and feeling and intention depends causally on the operation of this body, and in considerable part which mental events I have (in particular my sensations, feelings, and perceptual beliefs) are caused by events in this body. In so far as these events are caused, it is events in this body that cause them; and other events (for example, occurrences in the room) cause my mental events only by causing events in this body that cause the mental events. In consequence, thirdly, while I am aware of goings-on in this body without causal influences from outside the body impinging on it (I know the position of these limbs and feel the emptiness of this stomach), I can come to know about things outside the body only through their effects on this body. I see the desk and so know where it is only because light rays from the desk impinge on these eyes. I learn what you tell me only because by talking you set up air vibrations that impinge on these ears. And, fourthly, I look out on the world from where this body is. It is things around this body that I see well, things further away that I see less well.

Clearly a person has a body if there is a physical object (that is, a substance) to which he is related in all of the above four ways. And clearly a person does not have a body if there is no physical object to which he is related in any of the above ways. If a person is related to different physical objects in each of these ways, we shall have to say that his body is of a different kind or more widely extended than ours. And if he is related to a physical object only (or only to some

degree) in some of these ways, we shall have to say that he is embodied only to some degree.¹⁷

Now, on the traditional account of God, God is supposed not to be embodied in any of these ways. There is no physical object, not even the whole universe, through which he has to act in order to make differences to other things. He could abolish this physical universe at a stroke and create another one, and he can exert causal influence on non-embodied creatures without needing to operate through anything physical in order to do so. Nor is God dependent on anything physical or anything else for his life of thought. And he knows about everything without being dependant on any physical process for the acquisition of his knowledge. And he does not have any particular perspective on the world. He knows how things are without being dependent for his knowledge on a particular pattern of sensations arising from a particular viewpoint. So God is in no way embodied. He can, of course, move any part of the physical universe as a basic action, and knows without inference about the state of every part of the universe; but that does not make the physical universe his body, because he is not dependent on the universe for this ability and knowledge.

So then in the arguments to the existence to God the theist argues from the existence and order of the world and various features of it to a person, God, who brought these things about, meaning so to do. In this chapter I have been concerned to analyse what it is to explain an event as brought about by some person meaning so to do; and in conclusion I have drawn attention to two special features of personal explanations in terms of the action of God.

Having investigated in this chapter the structure of personal explanation, in the next chapter I will consider the evidence that justifies us in putting it forward, the evidence that makes it probable that an explanation of the personal type rather than one of the scientific type is the true explanation of some phenomenon. We will then be in a position to see whether the evidence recorded in the premisses of arguments to the existence of God constitutes such evidence.

¹⁷ The different ways in which persons can be embodied will be analysed more fully in Chapter 6.

The Justification of Explanation

What are the grounds for believing that some proposed explanation of a phenomenon *E* is a true explanation? (I write ‘a’ true explanation rather than ‘the’ true explanation, for as we have seen there may be many true explanations of the same phenomenon.)

The Justification of Scientific Explanation

To start with, what are the grounds for supposing that a proposed scientific explanation is a true one? In answering this, I assume, to start with, the amended Hempelian account of scientific explanation (outlined on pp. 26–9 ff.). My answer will be very brief, since my concern is more with personal explanation, but I think that it is sufficiently general to be acceptable to most philosophers of science. On the amended Hempelian account the occurrence of a phenomenon *E* is explained if laws of nature *L* and other particular phenomena *C* called the initial conditions physically necessitate (or make more probable) the occurrence of *E*. A proposed explanation will be a true one if the purported law *L* that it cites is in fact a law of nature and the cited initial conditions in fact occurred (and *L* and *C* do entail that it is physically necessary or physically more probable than it would otherwise be that *E* occur). A proposed explanation will probably be true in so far as (given the entailment, just mentioned) it is probable that *L* is a law of nature and probable that *C* occurred. It is probable that a universal statement such as ‘all material bodies attract each other with forces proportional to the product of their masses and inversely proportional to the square of their

distance apart' is a law of nature in so far as it belongs to a scientific theory that has high prior probability and great explanatory power.¹

The prior probability of a theory is its probability before we consider the detailed evidence of observation cited in its support. The prior probability of a theory depends on the degree of its fit with background knowledge (an *a posteriori* matter), and on its simplicity and scope (features internal to the theory and so an *a priori* matter). A theory fits with our general background knowledge of how the world works in so far as the kinds of entities and laws that it postulates are similar to those that probably (on our evidence) exist and operate in other fields. Thus a theory about the behaviour of argon at low temperatures would fit well with background knowledge, in so far as it postulated similar behaviour for argon to that postulated by other theories rendered probable on the same criteria for similar substances—for example, another inert gas, neon, at low temperatures.

Its degree of simplicity and its scope determine the intrinsic probability of a theory, its probability independent of its relation to any evidence. The simpler a theory, the more probable it is. The simplicity of a theory, in my view, is a matter of it postulating few (logically independent) entities, few properties of entities, few kinds of entities, few kinds of properties, properties more readily observable, few separate laws with few terms relating few variables, the simplest formulation of each law being mathematically simple. A theory of fundamental particles, for example, would be simple to the extent to which it postulates only a few kinds of particle with such properties (for example, mass and electric charge) of which we can observe other instances on the larger scale, whose behaviour is governed by simple mathematical formulae. A theory is simpler and so has greater prior probability to the extent to which these criteria are satisfied. But, of course, it is often the case that only a theory that is less than perfectly simple can satisfy the other criteria (for example, explanatory power) for probable truth. The best theory may be less than perfectly simple; but, other things being equal, the simpler, the more probably true.

¹ For full discussion of these criteria, which I describe in the next few pages, see my *Epistemic Justification* (Clarendon Press, 2001), ch. 4.

I count a property *P* as more readily observable than a property *Q* if one can discover whether or not an object is *P* without discovering whether or not it is *Q* but not vice versa (I understand ‘observability’ to include ‘experienceability’) The well-known philosophical example of ‘grue’ will illustrate this criterion. We may define an object as being ‘grue’ at a time *t* if and only if either it is green and *t* is earlier than AD 2050 or it is blue and *t* is AD 2050 or later. Then all objects observed so far (before AD 2050) that are green are also grue, and conversely. But our discovery that large numbers of emeralds are green and so also grue would not make it much more probable that there is a law that all emeralds are grue but would make it much more probable that there is a law that all emeralds are green, and the cited criterion explains why that is so. An object can be observed to be green or not without knowing what the date is, but to discover whether an object is grue we need to observe its colour (in the ordinary sense) and also discover the date. In this sense ‘grue’ is more remote from observation than ‘green’. (If it be suggested that a tribe might learn the meaning of ‘grue’, not by means of this definition but by examples of things said to be ‘grue’, then ‘grue’ will mean the same as ‘green’—since both would be defined by the same examples, and so there would be no conflict between ‘all emeralds are green’ and ‘all emeralds are grue’.²) Of course it is probable, as physics has taught us, that the fundamental laws of nature concern properties that are far from readily observable (for example, hypercharge and isospin); but that is because postulated laws of the latter kind have proved to have greater explanatory power than postulated laws in terms of observable properties. Other things being equal (which so often is not the case), laws in terms of observable properties are as such more likely to be true.

One formulation of a law is mathematically simpler than another in so far as the latter uses terms defined by terms used in the former but not vice versa. Mathematical operations can then be ordered in terms of simplicity—addition is simpler than multiplication, multiplication than powers; scalars than vectors, vectors than tensors, and so on. This requirement also has the consequence that simpler theories use small integers rather than large integers, and integers rather than integers followed by a complicated fraction. Thus for phenomena made equally probable (to the degree to which we can make

² For the sources and further discussion of the ‘grue’ problem, see *Epistemic Justification*. 88–9.

measurements), we should prefer the hypothesis of an attractive force between objects inversely proportional to r^2 (the square of their distance apart), rather than one inversely proportional to $r^{20\dots(100\text{ zeros})\dots 01}$. Interestingly, however, hypotheses attributing infinite values of properties to objects are simpler than ones attributing large finite values. For we can understand, for example, the notion of an infinite velocity (the velocity being greater than any number of finite units of velocity) without needing to know what the googplex is ($10^{10^{10}}$). And scientific practice shows this preference for infinite values over large finite values of a property. It preferred to postulate that light had an infinite velocity rather than a particular large finite velocity—for example, 301,000 km/sec.—until data were found that were very improbable on the former hypothesis. But note that the preference for the infinite over the large finite applies only to degrees of properties and not to numbers of independent entities. This difference arises, I suppose, because degrees of properties merge so as not to act independently—you cannot divide a velocity of 4 ft./sec. into two individuable velocities of 2 ft./sec. A velocity is a whole in the way that, say, a number of separably individuable planets are not. So, for example, we must not postulate an infinite number of planets in order to explain the motion of observable stars if we can explain that motion equally well by means of a large finite number of planets.

In assessing the simplicity of a scientific theory, in terms of the mathematical simplicity of its equations, scientific practice shows that we must use the simplest formulation of that theory. A theory telling us what entities there are, what properties they have and how they interact may be formulated in many different ways—that is, by means of many different but logically equivalent equations. ' $x = y$ ' is equivalent to ' $x = y + \frac{dy^3}{dy} - 3y^2$ ' and more generally to its conjunction with the most complicated mathematical theorem. But it is by its simplest formulation (for example, the former in the example) that we judge the simplicity of a theory. It lays bare the forces at work.

Yet, a theory's intrinsic probability is diminished in so far as its scope is great. What I mean by this is that, in so far as it purports to apply to more and more objects and to tell you more and more about them, it is less probable. Clearly the more you assert, the more likely you are to make a mistake. The force of this criterion is to render theories less probable in so far as they are about all material bodies

rather than (for example) just all bodies near the earth, or about all metals rather than just about copper. But typically, if a theory loses scope, it loses simplicity too, because any restriction of scope is often arbitrary and complicating. Why an arbitrary restriction to all bodies near the earth? A claim about the behaviour of all material bodies seems simpler. For this reason I do not think that the criterion of small scope is of great importance in determining prior probability; and so I shall concentrate largely on the other two criteria of prior probability, referring to this one only at crucial points. A theory has explanatory power in so far as it entails or makes probable the occurrence of many diverse phenomena that are all observed to occur, and the occurrence of which is not otherwise to be expected.

Thus, Newton's theory of motion, as put forward in his *Principia* in 1689, consisting of his three laws of motion and his law of gravitational attraction, satisfied these criteria very well; which made it probable that each of the proposed laws was indeed a law of nature. The theory was simple because there were only four very general laws of very great mathematical simplicity stating the mechanical relations that hold between all material bodies (that is, bodies having mass, a property that we feel on the human scale). Thus the law of gravitation stated that all material bodies attract each other in pairs with forces proportional to the product of the masses of each, m and m' , and inversely proportional to the square of their distances apart (r), $F = mm'/r^2$. The relations are simple because the distance is not raised to a complicated power (for example, we do not have $r^{2.0003}$ or $r^{\log 2}$), there is only one term (for example, we do not have $mm'/r^2 + mm'/r^4 + mm'/r^6$), and so on. Since the theory purported to cover all the mechanical behaviour of earthly and heavenly bodies, there was not in 1689 much other scientific knowledge with which it could fit. Its scope is very great in so far as it tells us about all material bodies, but it is concerned only with their mechanical interactions, not, for example, with their electrical interactions. The theory also had enormous explanatory power in that it rendered very probable the observed behaviour of bodies of very different kinds in very different circumstances—the motions of planets, the rise and fall of tides, the interactions of colliding bodies, the movements of pendula, etc. This aspect of a theory's explanatory power, I will call in future its predictive power. However, for a theory to have great explanatory power, the phenomena that it

predicts must be such that but for it they would not otherwise be expected. This will hold in so far as any other theories with significant prior probability do not predict them nearly as well as the theory in question; and so the evidence will have low prior probability, since only one theory (our theory) makes it at all to be expected. If another equally simple theory had predicted all that Newton's theory predicted in all its detail, the evidence would not support Newton's theory nearly as strongly. But no other simple theory could predict that. A theory has high explanatory power in so far as it has high predictive power and the evidence has low prior probability. So, except for the point about its enormous scope, Newton's theory satisfied the stated criteria very well. The fact that overall it was judged enormously probable illustrates my point that the criterion of scope is of far less importance than the other criteria.

Our grounds for believing that initial conditions *C* occurred are either that they were observed to occur, or, less directly, that the supposition that *C* occurred has itself great prior probability and explanatory power. It is for a reason of the latter kind that we suppose unobserved entities such as distant planets to exist. We observe a distant star moving in a certain way, and we can explain this if we suppose that there is close to it a massive planet that, in accordance with Newton's laws, is exerting on it an attraction so as to make it move in that way. If we suppose that Newton's laws operate (for which there is the vast amount of evidence that I have just outlined), we can account simply for the behaviour of the star by postulating at least one unobserved body that is exerting a gravitational force on the star. Such behaviour would otherwise be very improbable.³ It is clearly simpler to suppose that there is only one such body, and so this is the supposition with maximum prior probability and explanatory power.

It is also for a reason of this kind that we suppose unobservable entities such as atoms, molecules, photons, and protons to exist, to interact, and to have effects. We can explain the occurrence of certain clicks of Geiger counters and spots on photographic plates by the

³ If we were, for example, to postulate that some force of attraction other than the gravitational force were at work, we should be postulating the operation of a force determining star motion other than the force that determined all other star motion, and this would lead to a more complicated world picture than the supposition in the text.

supposition that certain such particles have produced them.⁴ So then, to summarize, our grounds for judging a proposed scientific explanation h of a phenomenon E to be probably true, where e is our observational knowledge, which includes E , are the prior probability of h and its explanatory power with respect to e .

I stress the enormous importance of the criterion of simplicity, an importance that is not always appreciated. Sometimes people ignore it and say that what makes a theory probable is just its explanatory power, or, worse still, just the fact that we can deduce from it statements reporting the phenomena that have been observed, our data or evidence. The trouble with this claim is that, for any finite collection of phenomena, there will always be an infinite number of different theories of equal scope such that from each (together with statements of initial conditions) can be deduced statements reporting the phenomena observed with perfect accuracy (and it may be that but for some one of these theories these phenomena are not to be expected). The theories agree in leading us to expect what has been observed so far, but disagree in their subsequent predictions. We may wait for new observations of phenomena to enable us to choose between theories. But, however many theories we eliminate by finding them incompatible with observations, we will always be left with an infinite number of theories between which to choose, on grounds other than explanatory power. If there are no theories of neighbouring fields with which some theories may fit better than others, the crucial criterion is that of simplicity. (And when our theories are very large scale, there will be little in the way of theories of neighbouring fields.)

This point may be illustrated by what is known as the 'curve-fitting' problem. Consider Kepler studying the motion of Mars. Suppose that he has as data a large finite number of past positions

⁴ These and similar examples, which will be discussed in Ch. 4, point to the obvious fact that science is often able to locate the cause of phenomena in some unobservable entity or process. Both Hume and Kant wrote when science had not had the success that it has had today in discovering the unobservable causes of observable events; and their philosophy of religion is often vitiated by the implicit or explicit principle that we could be justified in postulating a cause of some observable event only if that cause was also something observable. Thus Kant: 'If the empirically valid law of causality is to lead to the original being, the latter must belong to the chain of objects of experience' (*Critique of Pure Reason*, trans. N. Kemp-Smith (Macmillan, 1964), A636). It is sufficient to reflect on the evident success of chemistry and physics, in providing good grounds to believe in the existence of atoms, electrons, photons, etc., to realize that that principle is quite mistaken.

of Mars.⁵ He wishes to know the path along which Mars is moving, knowledge that will enable him to predict its future positions. He can mark on a map of the sky the past positions; but through those positions he can draw an infinite number of different curves, which diverge from each other in the future. One theory is, of course, that Mars moves in an ellipse. Another is that Mars moves in a spiral that diverges hardly at all from an ellipse during the period studied so far, but will diverge significantly hereafter. Another is that Mars moves along a path that describes increasingly large ellipses and eventually becomes parabolic. And so on. Of course very few of these theories would have been set out and seriously considered by Kepler or anyone else investigating the field. But my point is that, if the sole criterion for judging between theories was their ability to predict, all these theories would be equally likely to be true, for all of them would have been so far equally successful in predicting. The fact that many of the theories were not seriously considered is grounds for supposing that some other criterion was at work, and clearly it was the criterion of simplicity. Most theories that predict the data are theories that describe Mars as moving in a very contorted curve that can only be described by a very complicated equation. The theory that Mars moved in an ellipse was very simple one.

There must be a criterion to choose between the infinite number of theories that are equally successful in predicting the observations already made, if we are ever to be able to make any justified predictions for the future. The history of science reveals that, in the absence of background knowledge, that criterion is basically the criterion of simplicity. Without using this criterion we could make no progress at all in rational inquiry. *Simplex sigillum veri* ('The simple is the sign of the true') is a dominant theme of this book, as will become apparent in due course. All that I have been concerned to show here is the crucial influence of the criterion of simplicity within science. If we are to adopt in our investigations into religion the criteria of rational inquiry that are used in science and ordinary life, we must use this criterion there.

I stress that, as we deal with theories of larger and larger scope, there will be less and less background knowledge with which these

⁵ This supposition, of course, gives to Kepler far more knowledge than the historical Kepler had. The historical Kepler knew only the bearing of Mars from the earth at various times, not its distance from the earth as well. But I make my supposition to make exposition easier.

theories have to fit. More and more of the observational evidence falls into the category of data that the theory needs to explain, rather than data that it takes for granted in explaining other things. Newton sought a general theory of mechanics at a time when there was little by way of relevant data about non-mechanical phenomena. Since his time, as more data have been acquired and theories developed about electricity, magnetism, radiation, and so on, scientists have tried to develop a more fundamental theory of larger and larger scope to explain all these lower-level theories. In assessing candidate fundamental theories, 'fit with background knowledge' has become of less and less importance. A 'Theory of Everything' will have no contingent background evidence by which to determine prior probability. Prior probability must then be determined by purely *a priori* considerations.

But note that this does not make theories of narrow scope more 'empirical', less dependent on the *a priori* criterion of simplicity than theories of wide scope. For whether a narrow theory 'fits with' background knowledge is a matter of the simplicity of the conjunction of the postulated theory with that background knowledge. The theory about the behaviour of argon in a certain respect at low temperatures fits with a theory about the behaviour of neon and other inert gases at low temperatures if the conjunction of these theories is simple—for example, claiming that all inert gases conform in their behaviour to a certain simple equation at low temperatures—simpler than the conjunction of some other theory about argon with the theory about neon and other inert gases. Simplicity is an all-important and unavoidable criterion in assessing the probability of any scientific theory.

One consequence of the declining relevance of background knowledge as we deal with theories of larger and larger scope is that there becomes less reason to postulate entities and properties similar to those that play a role in theories of neighbouring fields. A theory about the behaviour of argon at low temperatures must postulate that argon consists of molecules that have mass and are subject to the laws of mechanics and gravitational attraction—because that is what we suppose about other gases. But, when we move to big theories of large scope that purport to explain so much more, we can (in so far as they satisfy the criteria of simplicity) postulate new kinds of entities and properties unlike those that occur in the lower-level theories that the big theory of large scope purports to explain. You

cannot suppose that argon is composed of quarks, while other gases are composed of molecules that are not made of quarks. But you can put forward a theory of the constitution of all protons and neutrons, that they are made of quarks, a quite new kind of entity with strange properties of a kind not hitherto observed.

Justification of Personal Explanation

We shall find that the same criteria of prior probability (determined by simplicity, scope, and fit with background knowledge, if any) and explanatory power are at work in assessing the probability of a hypothesis of personal explanation, a hypothesis that a certain agent produced some effect in virtue of certain beliefs, intentions, and powers.

We attribute effects to the actions of other human beings in so far as possible by attributing to them the same powers to perform basic actions, intentions, and ways of acquiring beliefs as similar as possible to our own (principle of charity) and as simple as possible (principle of simplicity), a picture that leads us to expect the public behaviour that in fact we find. Thus we suppose, other things being equal, that other humans who receive the same visual sensations as ourselves will come to hold the same beliefs: for example, that someone who has the sensations that we have when we see an aeroplane land will come to hold the belief that an aeroplane has landed (although, of course, it is possible that he may not, if he has never seen or heard of an aeroplane before). We suppose, other things being equal, that other humans have the same powers to move arms, legs, mouth, eyes, lips, etc., as we do ourselves. We suppose, other things being equal, that other people have similar intentions to ourselves—for example, to convey true information by what they say, when it is not too inconvenient. It is by making this assumption⁶ that we learn to interpret strange languages. All this involves applying the principle of charity. And we suppose that people do not change their intentions and beliefs suddenly and at random—that their intentions remain constant over a period, and that their beliefs change under the reception of sensory stimuli in regular ways. This is an application of the principle of simplicity. But

⁶ I develop this point in more detail on pp. 126–7.

the picture that we build up of people must be such as to lead us to expect the behaviour that we find. If we suppose that a man has the sole intention of posting a letter and believes that the postbox is up the road to the left, our supposition must be withdrawn if he walks down the road to the right. The principle of charity is really an application of the principle of simplicity as I have defined it—for, in making the assumption that the intentions of others and their ways of acquiring beliefs are similar to our own, we make a simpler assumption than the assumption that they differ.

We extend this general account to include animals. To animals with appendages (for example, mouth, legs, etc.) similar to our own, in so far as we can we attribute similar powers to move those appendages. Likewise, in so far as we can, we attribute to animals intentions and ways of acquiring beliefs similar to our own—for example, an intention to get food when the animal has not eaten for some time; and, if visual stimuli keep landing on the animal's eye from food in a cupboard, a belief that there is food in that cupboard. Again, however, our picture must lead us to expect the behaviour that we find, and must be modified or extended in order so to do. We attribute to animals powers other than our own—for example, the power to move a tail as a basic action; and we deny to them intentions similar to ours—for example, the intention to utter a complicated statement—in order to have a simple picture of their powers, intentions, and beliefs that leads us to expect the animal movements that we observe.

In the cases of humans and animals we assume that we can recognize human and animal bodies, the vehicles for the basic actions of rational agents. But we are also, I suggest, prepared to explain other phenomena as due to the actions of rational agents with powers, beliefs, and intentions in so far as such explanations satisfy the criteria of prior probability and explanatory power that we have seen at work in assessing scientific explanations. To start with, we are prepared to believe some physical object of a quite new and strange kind to be the body of a rational agent in so far as we can give a simple account of its behaviour in personal terms. We might discover on another planet some physical object *P* that undergoes various changes of bodily state and makes movements, such as a movement of some appendage *E*, but it is open to question whether *P* is a rational agent. If we are to suppose *P* to be a rational agent and *E* to be a result of an action, we are clearly making a supposition that

does not fit in very well with our background knowledge of the world. We are going to suppose the existence of a rational agent very different in its history, appearance, and physiology from the ones with which we are familiar. We may need also to suppose that *P* has basic powers, intentions, and ways of acquiring beliefs (for example, not via stimuli impinging on sense organs similar to our own) very different from those of humans. The suppositions about *P* may also need to be complex ones—for example, we may need to suppose that *P* has inconsistent and quickly changing intentions. Now clearly the greater the differences that we suppose between *P* and known rational agents, and the less simple our picture of *P*, the less probable it is that *P* is a rational agent. Yet, however initially improbable our suggestion, its explanatory power could be so great as to render it probable nevertheless. If a certain supposition about *P*, how he acquires beliefs, and what are his powers and intentions, made very probable certain movements in or of his body that would otherwise be very improbable, then we would have good grounds for believing it true. Thus among the suppositions about *P*'s intentions, beliefs, and basic powers may be detailed suppositions about *P*'s language—namely, that he utters certain syllables with the intention thereby of conveying certain information and that he believes that the uttering of those syllables will convey that information, and so on for various other syllables. If this supposition makes probable his utterance of various syllables rather than various other syllables under certain circumstances, and he does utter the former syllables, then that raises the probability of the supposition about *P*. Much evidence of this type would raise that probability greatly.

It is possible that we might find certain otherwise inexplicable phenomena that could be explained by the action of a non-embodied agent, such as a ghost or a poltergeist. The phenomena to be explained may be that books, chairs, inkwells, etc. start flying about my room. We postulate a poltergeist *P* with certain intentions, beliefs, and powers to be responsible. Clearly we have to suppose *P* to be very unlike other rational agents known to us both in his powers and in his ways of acquiring beliefs. (This cannot be, for example, via sense organs.) But we can suppose *P* to have beliefs influenced as are ours by how things are, and to have intentions of the kind that we have—for example, intentions of a kind typical of human beings with certain characters and histories. For example, we can suppose *P* to have previously been a certain embodied person who had been

greatly injured by *X* and who had greatly loved *Y*, *X* and *Y* both being still alive. Then we suppose *P* to be like many of us if we suppose him with such a history to have the intention to harm *X* and to save *Y* from harm, and to communicate with *Y*. If we suppose *P* to be in these ways similar to us, the supposition that *P* exists fits in to some extent with our background knowledge, although, in postulating a non-embodied person, clearly to a large extent it does not. The supposition is more probable in so far as it is simple—that is, postulates few constant intentions, simple ways of acquiring beliefs, and unchanging powers. Above all, the supposition will be rendered probable if it has high explanatory power. It would have this, for example, if the books, chairs, inkwells, etc. hit *X*, or form themselves into words that warn *Y* of impending danger; and so on. We would expect this kind of thing to happen if *P* is as we have supposed, far more than we would ordinarily expect it to happen. It is crucial that there does not exist a more probable scientific explanation of the goings-on (apart from any that explains or is explained by the intentions of poltergeists). It must not be more probable that the phenomena of books, chairs, and inkwells flying around are due to a hurricane or to my having suddenly acquired great mass and so great power of gravitational attraction, where these latter are explicable by normal scientific explanation. It is also crucial that there should not exist a probable personal explanation in terms of the action of an embodied agent (other than one that is explicable by or explains the action of a poltergeist). If an embodied agent moved the chairs by telekinesis, then a poltergeist did not.

The examples that I have taken so far of cases where we may claim some phenomenon to be the result of an action of a rational agent, other than a human being using normal basic powers, are cases where we claim the phenomenon to be the result of a basic action. If we are to claim a phenomenon to be the result of a mediated action, we must find a basic action of which it is an intended consequence. So, if we are to discover the results of mediated actions of agents other than humans, we must first discover the results of their basic actions. Once we have found basic actions other than normal ones, we can use criteria of the kind to which I referred earlier to determine whether their consequences were intended by the agent, and so whether they are to be considered the results of mediated actions.

In all this we see the investigator using the criteria of prior probability and explanatory power to judge the worth of proposed hypotheses of personal explanation, just as with proposed theories of scientific explanation. The prior probability of a hypothesis is, we saw, a matter of its fit with background knowledge, its simplicity, and its lack of scope. Fit with background knowledge in the case of the hypotheses of personal explanation considered so far is a matter of postulating persons similar to known persons in their history and physiology, their basic powers, their intentions, and their ways of acquiring beliefs. We saw that the less similar to known persons (namely, humans) were the postulated persons, the less probable it was that they exist. Simplicity in the case of hypotheses of personal explanation is a matter, first, of postulating few persons. You don't postulate that twenty persons have caused various phenomena—for example, twenty footprints on a beach—if you can explain the phenomena by postulating that only one person has caused the phenomena—that the footprints were made by one person walking. Then it involves postulating few properties, few constant intentions, and continuing basic powers; and simple laws—constant predictable ways in which persons acquire beliefs from their surroundings. A hypothesis will have smaller scope and so be more likely to be true in so far as it tells you about the causes of fewer phenomena; or provides less detail about a person's intentions, powers, etc.; but, even if it is a detailed hypothesis, we could have enough evidence to make it likely to be true. The explanatory power of a hypothesis of personal explanation is, first, a matter of its ability to predict the phenomena that we in fact observe. Thus a theory that *P* has the power of bending spoons at a distance and the intention of doing what people ask him to do leads us to predict that, when we ask him to bend the spoons, the spoons will bend. If our predictions come off, that is indeed evidence for the theory. But, as I noted with scientific explanation, the explanatory power of a hypothesis depends also on its evidence being such 'as was not otherwise to be expected'. That is, the evidence must not be probable on other hypotheses with relatively high prior probability and predictive power. It was for this latter reason that it was of great importance, if we were to accept theories about poltergeists or men with the ability to bend spoons at a distance, that there should be no other probable explanation of these phenomena.

The hypothesis of theism that seeks to explain the existence of the universe and its various features is, as we have seen, a hypothesis of

personal explanation; and so it is to be assessed by these criteria. But note that it is a hypothesis of enormous scope. A physical ‘Theory of Everything’ purports to explain everything physical; theism purports to explain everything logically contingent (apart from itself). In consequence there will be no background knowledge with which it has to fit. It will not, therefore, be a disadvantage to it if it postulates a person in many ways rather unlike the embodied human persons so familiar to us. In considering the arguments for the existence of God, we shall begin with a situation of tautological background knowledge, and so the dissimilarities between human persons and the postulated God will not as such affect the prior probability of theism. The fact, for example, that humans normally execute their intentions through a chain of neural events culminating in bodily events does not form part of our background knowledge when judging the probability of the existence of a God of whom that is not true. But, of course, the argument for theism must take account of this fact about humans, if not as background knowledge, then as evidence to be explained by the hypothesis. The proponent of theism needs to explain why a bodiless God should create embodied humans, and I seek to do that in Chapter 6. The fact that all the material bodies that we observe have diameters of more than 1 mm is not normally construed as background knowledge, making it improbable that there exist fundamental particles of far smaller diameters, but rather as something requiring explanation in part by the latter hypothesis.

Bayes’s Theorem

We can now put our points about the probability of a hypothesis h on evidence e depending directly on the prior probability of h and the predictive power of h , and inversely on the prior probability of e , into symbolic form. Where k is our general background knowledge, of what there is in the world and how it works, e is our phenomena to be explained and other relevant observational evidence, and h is our hypothesis, $P(h|e \& k)$ is a function of the prior probability of h , $P(h|k)$; and of its explanatory power with respect to e . This latter is a factor that increases with the predictive power of h , $P(e|h \& k)$; and decreases with the prior probability of e , $P(e|k)$. $P(e|h \& k)$ is a measure of how likely the observed phenomena e are to occur if the hypothesis h is true (given our background knowledge k). So, the

more h makes e probable, to be expected, the greater is $\frac{P(e|h \& k)}{P(e|k)}$. $P(e|k)$ measures the prior probability of e , how likely e is to occur anyway, whether or not h , given only k . Clearly, the more evidence we have, the more diverse and otherwise inexplicable is our evidence, the lower (relative to $P(e|h \& k)$) is $P(e|k)$, and so again the greater is $\frac{P(e|h \& k)}{P(e|k)}$.

These points are made explicit by a basic theorem of confirmation theory, Bayes's theorem,⁷ which runs as follows:

$$P(h|e \& k) = \frac{P(e|h \& k) P(h|k)}{P(e|k)}.$$

This theorem follows directly from the axioms of the mathematical calculus of probability, for the truth of which are good independent grounds.⁸ But in putting it forward I do not appeal much to these latter grounds, but mainly to the ones given so far in this chapter (although the particular way in which $P(h|e \& k)$ increases with $P(h|k)$ and $P(e|h \& k)$, but decreases with $P(e|k)$ does not depend on anything that I have said so far, but must depend on this derivation).

$P(h|k)$, the prior probability of h , depends as we have seen, in the normal case both on the internal simplicity of h (and its narrowness of scope) and also on how well h fits in with our general background knowledge of the world contained in k . However, as we saw in Chapter 1, any division of evidence between e and k will be a somewhat arbitrary one. Normally it is convenient to call the latest piece of observational evidence e and the rest k ; but sometimes it is convenient to let e be all observational evidence and let k be mere 'tautological evidence'. In the latter case the prior probability $P(h|k)$ is what I shall call the intrinsic probability of h , and will depend mainly on the simplicity of h (as well as to a lesser extent on its narrowness of scope). But, if k contains logically contingent evidence of what there is in the world and how it works, $P(h|k)$ will depend also on how well h fits in with that evidence. Where k is mere

⁷ $\frac{P(e|h \& k) P(h|k)}{P(e|k)}$ means, of course, $P(e|h \& k)$ multiplied by $P(h|k)$, and the result divided by $P(e|k)$.

⁸ In my view the best arguments for these axioms are that they codify the judgements of what is evidence for what, which seem to us intuitively correct. On this, see my *Epistemic Justification*, chs. 3 and 4.

'tautological evidence', $P(e|k)$ will be what I shall call the intrinsic probability of e .

I have claimed that Bayes's theorem is true, but I had better make clear what I mean by saying this. I mean that, in so far as for various e , h , and k , the probabilities occurring in it can be given a numerical value, it correctly states the numerical relationships that hold between them. In so far as they cannot be given precise numerical values, my claim that Bayes's theorem is true is simply the claim that all statements of comparative probability that are entailed by the theorem are true. By statements of comparative probability I mean statements about one probability being greater than, or equal to, or less than another probability. (Such statements are sometimes all that we can justifiably assert about some probabilities—see pp. 17–18.) Thus it follows from Bayes's theorem that, if there are two hypotheses h_1 and h_2 such that $P(e|h_1 \& k) = P(e|h_2 \& k)$, then $P(h_1|e \& k) > P(h_2|e \& k)$ if and only if $P(h_1|k) > P(h_2|k)$. This says that, if h_1 and h_2 both make it equally probable that we will find evidence e , given background knowledge k , then one of them h_1 will be more probable than the other on the total evidence e and k , if and only if h_1 was more probable than h_2 on the background evidence alone. Put more technically: if h_1 and h_2 have equal predictive power, h_1 will have greater posterior probability (that is, probability on the total evidence e and k) than h_2 , if and only if it has greater prior probability. So, if there are two scientific theories equally successful in predicting certain observations, one of them will be more probable than the other if and only if it was more probable before the observations were made. Or again it follows from Bayes's theorem that, if $P(h_1|k) = P(h_2|k)$, then $P(h_1|e \& k) > P(h_2|e \& k)$ if and only if $P(e|h_1 \& k) > P(e|h_2 \& k)$. This says that, if two hypotheses are equally probable before certain evidence e is obtained, one of them will be more probable than the other on the total evidence, if and only if, given that hypothesis, it is more probable that e will be found than it is given the other hypothesis. (In an extreme case, h_1 may entail e —it may be a deductive consequence of h_1 that e will occur—and h_2 may entail $\sim e$, that e will not occur.)

Here is another example, slightly different from those used so far, to illustrate the working of Bayes's theorem. Let h be the hypothesis that Jones robbed Barclays Bank, e be the evidence that he was near the bank at the time of the crime, and k be the background knowledge that Jones robbed another bank, Lloyds Bank, on another

occasion. Then $P(h|e \& k)$ is determined by the explanatory power of h , $\frac{P(e|h \& k)}{P(e|k)}$, and the prior probability of h , $P(h|k)$. $P(e|h \& k)$ is the probability that e , given both h and k . In this case this is 1, since, if Jones robbed the bank, he must have been near the scene at the time. $P(e|k)$ is the probability that he would be near the scene at the time, given that he had robbed Lloyds bank. This will be greater than $P(h|k)$, the probability, given that he robbed Lloyds Bank, that he robbed Barclays Bank; since he might have had quite innocent reasons for being where he was. The probability that he robbed Barclays Bank is then the prior probability that he did multiplied by the extent to which the hypothesis that he did it makes e more to be expected than it would otherwise be.

It will be useful at this stage, before proceeding with the main argument, to make another important point about confirmation that can be illustrated by Bayes's theorem. It is sometimes said that we are justified in accepting a hypothesis only if we have tested it by finding that it predicts certain events and then waited to see whether or not those events happen; and only if they do are we justified in accepting the hypothesis. Now it seems to me that, although we often test hypotheses in this way, we do not have to do so if they are to be rendered probable by our evidence and so we are to be justified in accepting them. The suggestion that hypotheses must predict successfully (interpreted in the above literal sense) if they are to be rendered probable by evidence is certainly not implied by Bayes's theorem. It is a matter of indifference, as regards that theorem, whether e is observed before or after the formulation of h . All that matters is the relations of probability holding between e and h . And surely the theorem is correct in that respect. My calling $P(e|h \& k)$ the 'predictive power' of h was not meant to imply that e was discovered only subsequent to the postulation of h that predicted it.

Newton's theory of motion was judged to be highly probable on the evidence available in the late seventeenth century, even though it made no new immediately testable predictions, other than the predictions that were already made by laws that were already known and that it explained (for example, Kepler's laws of planetary motion and Galileo's law of fall). Its high probability arose solely from its being a very simple fundamental theory from which those diverse laws are deducible. More generally, whether e renders h probable surely cannot depend crucially on whether we had thought of h before we saw e . Probability would become a highly subjective matter instead of an

objective relationship between evidence and hypothesis if that were so. Bayes's theorem is, however, able to explain why often, indeed *normally*, we are interested in predictions that we can check subsequently to the formulation of the theory. This is because only when we have the theory (h) do we know just which evidence is such as to make $\frac{P(e|h \& k)}{P(e|k)}$ very great; only then do we know which evidence we need to have in order that the theory be rendered highly probable. We are not very likely to have that evidence in our hands already—we normally need to look for it. Nevertheless, we *may* have it in our hands already. Hence it is in itself no objection to the hypothesis that there is a God, that it does not yield predictions such that we can know only tomorrow, and not today, whether they succeed. The theist's evidence may render his hypothesis probable without this condition being satisfied.⁹

It follows immediately from Bayes's theorem that $P(h|e \& k) > P(h|k)$ if and only if $P(e|h \& k) > P(e|k)$. This important principle is one that Mackie called the 'relevance criterion'.¹⁰ It follows from it by a fairly short step of logic that $P(h|e \& k) > P(h|k)$ if and only if $P(e|h \& k) > P(e|\sim h \& k)$. This says that a hypothesis h is confirmed by evidence e if and only if that evidence is more likely to occur if the hypothesis is true than if it is false. The result is surely correct. It is implicit in many judgements that we make in ordinary life. His fingerprints on the safe confirm the supposition that Jones robbed the safe if and only if they are more likely to be on the safe if he did the crime than if he did not. If they are equally likely to be on the safe whether or not he robbed the safe (for example, because Jones is the manager of the shop in which the safe is situated and often opens it), they do not confirm the supposition that he robbed the safe. It follows that an argument from e to h is a correct C-inductive argument if (and only if) e is more likely to be found if h is true than if h is false.

While the simplest theory (among theories of equal scope and fit with background evidence) has the greatest intrinsic probability, how probable is it, relative to slightly less simple theories? The practice of scientists, historians, etc. shows that they judge a *very* simple theory to be very much more probable than less simple theories. If you can

⁹ For full discussion of the irrelevance of prediction in a literal sense, see *Epistemic Justification*, app., 'Predictivism'.

¹⁰ See J. L. Mackie, 'The Relevance Criterion of Confirmation', *British Journal for the Philosophy of Science*, 20 (1969), 27–40.

explain many clues by the hypothesis that one agent caused them, that is far more probable than a theory with the same explanatory power to explain the clues that postulates that two agents caused the clues. And a theory that postulates an inverse square law of attraction is intrinsically far more probable than one that postulates a law of attraction varying with the 2.01th power of distance apart. And a theory that relates only a few postulated variables and yet is able to explain as much as a theory that relates many variables is far preferable.

However, when we come to not quite so simple theories (which we may need to do when very simple theories prove to have low explanatory power), the intrinsic probability of the simplest remaining theory is not much greater than that of the next less simple theory. The theory that the force of attraction varies as the 2.01th power is barely more probable intrinsically than the theory that it varies as the 2.012th power. Still, the practice of scientists and other investigators suggests that, although there will always be an infinite number of very complex theories able to explain the data with any arbitrary explanatory power, they always judge it more probable on the data available at some time that the correct explanation lies within some group of simplest theories. Too complex theories are regarded as too improbable to be true, if there are even moderately simple theories with significant explanatory power.

The theist argues from the world, the fact of its existence, and its detailed characteristics to a God who brought it about. Since the structure of his argument is the same as that of an argument from a narrower range of phenomena to a non-embodied person, such as a poltergeist, who brought them about as results of his actions, we must use the same criteria embodied in Bayes's theorem—bearing in mind, as I have emphasized, the differences between the range of the phenomena. The theist argues from all the phenomena of experience, not from a small range of them. We shall have to let e represent in turn the different facets of the world that the theist brings as evidence for the existence of God and the facets that the atheist brings as evidence against the existence of God. h will be the hypothesis that God exists, and, to start with, k will be mere tautological evidence. In order to assess $P(h|e \& k)$ in each case we shall then need to assess $\frac{P(e|h \& k)}{P(e|k)}$ and $P(h|k)$. The probability, on the evidence, of God's existence will depend on how well the hypothesis of God's existence is able to explain the occurrence of phenomena that would otherwise

be highly unlikely; and on its prior probability, which (since there will be no background knowledge) means its intrinsic probability, dependent on its scope and its simplicity. Theism is a hypothesis of vast scope; but, of course, of the same vast scope as any other world view—for example, physicalism. And so, for comparing it with other world views, we can ignore the issue of its scope. And, as I suggested earlier, with the example of a scientific theory—Newton’s theory—scope is anyway a criterion of far less importance than simplicity in determining intrinsic probability. The crucial determinant of the prior probability of theism must be simplicity. I shall assess the simplicity and so the intrinsic probability of theism in Chapter 5. The explanatory power of theism will vary, as we have seen, with the different e . But, before discussing the evidential force of different e —that is, the different arguments—we shall need to consider general principles involved in determining the values of the two probabilities, $P(e|h \& k)$ and $P(e|k)$, which determine the explanatory power. $P(e|k)$ is a matter of how likely the various phenomena are to occur anyway—that is, whether or not God brings them about. It follows from the calculus that:

$$P(e|k) = P(e|h \& k) P(h|k) + P(e|\sim h \& k) (P \sim h|k).$$

The first conjunct on the right-hand side ($P(e|h \& k)P(h|k)$) simply repeats the top line of the right-hand side of Bayes’s theorem. So by the theorem $P(h|e \& k)$ will be close to 1 if and only if the second conjunct ($P(e|\sim h \& k)P(\sim h|k)$) is low (relative to the first conjunct). This second conjunct measures how likely it is that e will occur if there is no God. This will be low if it is not probable that any other cause would bring e about or that e would occur uncaused. I shall discuss in Chapter 4 the general principles involved in assessing the latter, and especially what are the grounds for claiming that some phenomena occur uncaused (that is, have no explanation). In Chapter 6 I shall consider in general terms how $P(e|h \& k)$ is to be assessed; that is, the kind of e that God is likely to bring about. Note that I take h simply as ‘there is a God’. By itself it provides merely a partial explanation of e . It needs to be conjoined with an intention to bring about e in order to provide a full explanation of e . The value of $P(e|h \& k)$ will, for the various e , depend on how probable it is that God will have that intention.

Complete Explanation

I argued in the last chapter that the grounds for postulating a personal explanation of phenomena lie in the prior probability and explanatory power of that explanation. The explanatory power of an explanation depends crucially on the prior probability of the phenomena, how likely it is that phenomena would occur irrespective of whether the proposed explanation of them is correct. So a crucial factor in determining the force of an argument from phenomena to God's existence is whether those phenomena would be at all likely to have occurred but for God's agency. They could exist apart from God's agency, if something else could have brought them into being—namely, if there is a scientific explanation (or other personal explanation) of their existence—or if they could exist uncaused, without there being any explanation of their existence. In this chapter I shall consider what kind of phenomena could not be given a scientific (or other personal) explanation and then go on to consider what kind of phenomena could exist uncaused. This latter is the issue of what is the proper terminus of explanation; when do we have grounds for supposing that we have reached the end of the explanatory road, and when not. This chapter is concerned to develop general principles, which can subsequently be applied to the arguments for the existence of God. The theist claims that the phenomena that he cites—for example, the existence of the universe—could not exist uncaused, but that God could exist uncaused. We shall need to investigate whether the existence of God is a more satisfactory terminus for explanation than the existence of the universe with its various characteristics. If there is a God, it follows, as I shall show briefly in Chapter 5, that explanation stops with him—a being would not be God if something other than God could explain his existence. But what needs to be shown is that the existence of God forms a more

natural stopping place for explanation than, say, the existence of the universe.

The Scientifically Inexplicable

What, to start with, are the grounds for supposing that phenomena do not have a scientific explanation? Phenomena of two kinds can be shown not to be explicable scientifically. First, there are phenomena that are too *odd* to be fitted into the established pattern of scientific explanation, and, secondly, there are phenomena that are too *big* to be fitted into any pattern of scientific explanation.

To show phenomena too odd to be explicable scientifically the theist needs to show that there is good evidence for a scientific system *h* covering a certain range of phenomena, but that it is not a consequence of *h* that certain phenomena (within the general range of *h*) occur; and that any attempt to amend or expand *h* to allow it to predict *e* would make *h* so complex that it would be very improbable that it is true. Theists have claimed various particular phenomena to be too odd to be scientifically explicable. Among these, if we assume that they occur, are violations of laws of nature such as levitations, or people getting better from polio in a minute, or blood suddenly liquefying, or people walking on water, events that theists claim to be miracles. I shall discuss the issue of miracles in Chapter 12.

But the oddness of events need not be confined to the particular; there may be events of certain *kinds* that cannot be explained by science. There may be much evidence for a certain scientific system, and yet it be a consequence of this system that a scientific explanation of certain kinds of event is ruled out. Theists have sometimes claimed that the occurrence of living organisms or of conscious human beings is scientifically inexplicable. I do not think that much of a case can be made out for the former, but I do think that a substantial case can be made out for the latter and I shall discuss that case in Chapter 9. I shall analyse in detail and illustrate with examples in these two later chapters, 9 and 12, the kind of argument used to show phenomena to be too odd to be scientifically explicable. Here I have only sketched it in schematic outline.

The other phenomena that cannot be explained scientifically are phenomena that are too *big* for science, and too big not merely for some particular well-established scientific system, but for any

scientific system. In considering, in Chapter 2, the nature of scientific explanation, we saw that science explains why some event or state of affairs occurs. It does this, on the Hempelian model, in terms of a prior state of affairs and some natural law. It also explains why certain natural laws operate, and it does this in terms of more fundamental laws of nature—for example, it explains the operation of Galileo's law in terms of the operation of Newton's laws. But what, as I shall show more precisely in Chapter 7, science could not explain is why there are any states of affairs at all; it can explain only why, given that there are such states, this state is followed by that state. Nor could it explain, as I shall show more precisely in Chapter 8, why the most fundamental natural laws of all hold. Either these are brute facts about the world, or they have an explanation of a different kind.

We shall see that the scientifically inexplicable, the odd and the big, form the normal starting points for arguments to the existence of God. Cosmological and most teleological arguments argue from phenomena allegedly too big for science to explain; whereas most other arguments argue from phenomena allegedly too odd for science to explain. The arguments need to show also that there is no personal explanation in terms of the action of an embodied agent. This done, what has been shown is that an explanation in terms of a very powerful non-embodied agent is the only possible explanation of the phenomena. It then follows that either theism or something like it is true, or that the phenomena are just brute inexplicable facts, the stopping point of explanation.

Full, Complete, Ultimate, and Absolute Explanation

The main issue of this chapter is, then, what are proper stopping points for explanation, when do we have grounds for supposing that certain phenomena constitute such stopping points, and when do we have grounds for supposing that they themselves have an explanation. Having established general principles, we can then ask whether it is right to suppose that the physical universe, or the regular operation of scientific laws, or such particular events as (assuming that it happened) the resurrection of Christ from the dead, are just brute facts, or whether they are phenomena that it is reasonable to suppose to have a further explanation.

Before I come to deal properly with the central issue. I need first to make certain distinctions. The question of what is a ‘terminus’ of explanation needs a more technical explication. I have already in Chapter 2 introduced the concepts of full explanation and partial explanation of some phenomenon. An explanation of E by F is a full one if F includes both a cause, C , and a reason, R , which together necessitated the occurrence of E . (Recall that on the Hempelian account of scientific explanation C are initial conditions and R natural laws; whereas in personal explanation C is a person, and R that person’s intentions, beliefs, and basic powers.) If C and R together provide a full explanation of E , nothing else logically contingent beside C and R needed to be so in order for the occurrence of E to be guaranteed, and so a proposition reporting C and R entails a proposition reporting E . Thus a scientific explanation of an eclipse of the moon E is a full one if it cites natural laws involved, L (Newton’s laws of motion and laws of light propagation), and precedent states of the world C necessary for the occurrence of E (the positions and masses of moon, sun, and earth, and the absence of other heavenly bodies in the region) and if L and C together entail E . An explanation of E is only a partial one if the explanation includes factors that contributed to bringing about the occurrence of E (made it physically more probable), but these factors did not necessitate the occurrence of E .

Now a full explanation really does by itself explain why something happened. It does so quite independently of whether or not there is an explanation of how any states it cites came to be (for example, why the sun was where it was) or why any reasons that it cites operate (for example, why certain natural laws hold). To suppose otherwise is to commit a fallacy that we may call ‘the completist fallacy’. Clearly it is a fallacy. For if it were really the case that F could not explain E unless there is an explanation of F , nothing in the universe could be explained, unless there were explanations of such things as the origin of our galaxy—which is absurd. It is, however, a common fallacy. Thus Hume objects in the *Dialogues* to postulating a God who planned the world as an explanation of its order, on the grounds that the postulated existence of a rational agent who produces the order of the world would itself need explaining. Picturing such an agent as a mind, and a mind as an arrangement of ideas, Hume phrases the objection as follows: ‘a mental world or universe of ideas requires a cause as much as does a material world or universe of

objects.¹ Hume himself provides the obvious answer to this—that it is no objection to explaining *E* by *F* that we cannot explain *F*. But then he suggests that the *F* in this case, the mind, is just as mysterious as the ordered universe. Men never ‘thought it satisfactory to explain a particular effect by a particular cause which was no more to be accounted for than the effect itself.’² But that is plainly false. We can give a perfectly good explanation of how it came about that Jones lost his fortune in terms of the way the Monte Carlo roulette wheel spun as it did, while judging that there was no explanation of how the roulette wheel spun, this being something utterly beyond accounting for.

Nevertheless, although a full explanation of *E* (in the sense of ‘full explanation’ that I have delineated) leaves no facet of *E* unexplained, further explaining can often be done—there may be explanations of why the factors cited in the *explanans* are operative and how they came to exist. Let us concentrate for the moment on explanations of the operation of the factors at the time³ *t* at which they bring about *E*. Let the factors be *C*, the cause that brings about *E* at the time, and *R*, the reason for *C*’s efficacy. Let *C* bring about *E* at time *t*. The existence of *C* at *t* may depend on some other factor *B*, which at time *t* makes *C* exist. Thus, suppose my arm by moving makes a stick move at that time, and the stick’s motion makes a stone move at that time. The motion of the stone (*E*) is brought about by the motion of the stick (*C*), which in turn is brought about by the motion of my arm (*B*). Then the present existence of a cause is dependent on the

¹ David Hume, *Dialogues concerning Natural Religion* (first published 1779; ed. H. D. Aiken (Hafner, 1948)), 33.

² *Ibid.* 36.

³ A cause *C* acting ‘at the time’ so as to bring about *E* needs to be interpreted carefully so as not to entail literally simultaneous causation (the period of *C*’s action being of exactly the same length as *E*), nor so as to rule it out. This is because some thinkers (including myself) believe that literally simultaneous causation is logically impossible, for the same reason as backward causation (a cause causing an earlier effect) is logically impossible. (For my reasons for this belief, see my *The Christian God* (Clarendon Press, 1994), 82.) The cause acting ‘at the time’ also needs to be interpreted in such a way as to take account of the fact that the actions of causes and the occurrence of effects are events that take time (last for a period of time, however short). So I suggest that we understand by *C* acting at *t* to bring about *E* at *t* as every segment (however small) of *C*’s action over a period of time including *t* causing a segment of *E*, where *C*’s action and *E* both end at the same instant. (Then, however small is any final segment of *C*’s action, there will be some—even smaller—segment of *E* that it causes; and no initial segment of *C*’s action, however small, will be without effect on some segment of *E*.)

present operation of its own cause. Likewise the operation of the reason R may depend on some higher-level reason S , which at the time of R 's operation makes R to operate. Thus Galileo's law of fall operates on earth because earth has such-and-such a mass and Newton's laws operate.

I now define a complete explanation of the occurrence of E as follows. A complete explanation of the occurrence of E is a full explanation of its occurrence in which all the factors cited are such that there is no explanation (either full or partial) of their existence or operation in terms of factors operative at the time of their existence or operation. Thus, suppose that a high tide is brought about by sun, moon, earth, water, etc. being in certain positions and by the operation of Newton's laws. Here is, let us suppose, a full explanation. Suppose too that Newton's laws operate here because this region of the universe is relatively empty of matter and Einstein's laws of General Relativity operate. These factors act contemporaneously to make Newton's laws operate. Suppose, too, that nothing at this time makes sun, moon, etc. be where they are (even though some past cause was responsible for their being where they are). Nor does anything at this time make Einstein's laws operate or this region of the universe be relatively empty. Then there is a complete explanation of the high tide in terms of the operation of Einstein's laws, the universe in this region being relatively empty of matter, and sun, moon, earth, water, etc. being where they are.⁴

Complete explanation is a special kind of full explanation. I now delineate as a special kind of complete explanation what I shall call ultimate explanation. To speak loosely to start with, we have an ultimate explanation of some phenomenon E if we can state not merely which factors C and R operated at the time to bring E about, and which contemporaneous factors made C and R exist and operate at that time, and so on until we reach factors for the contemporaneous existence and operation of which there is no

⁴ It may be that, while C and R provide a full explanation of E , either C or R or both have no full explanation in terms of factors operative at the time of their operation, but only a partial explanation of this kind. So long as there is an end to any regress of partial explanation, there will, nevertheless, on my definition be a complete explanation of E . For suppose that event B (in virtue of a law S) partially explains the contemporaneous occurrence of C , but that neither B nor S has an explanation (either full or partial) in terms of factors operative at the time of its occurrence, then a complete explanation of the occurrence of E is provided conjointly by B , C , R , and S .

explanation; but also state the factors that originally brought *C* and *R* about, and which factors originally brought those factors about, and so on until we reach factors for the existence and operation of which there is no explanation. Less loosely, I define an ultimate explanation of *E* as a complete explanation of *E*, in which the factors *C* and *R* cited are such that their existence and operation have no explanation either full or partial in terms of any other factors. Those factors are ultimate brute facts. Suppose that there is no God, that the universe began with a bang in a state *X* at a time *t*, that it is governed by deterministic laws *L* (whose operation is not further explicable); and that, in accord with *L*, *X* brought about a state *Y*, and *Y* brought about a state *Z*, and *Z* brought about *E*. Then (*X* and *L*); and (*Y* and *L*), and (*Z* and *L*) are each complete explanations of *E*; but only (*X* and *L*) is an ultimate explanation of *E*.

Finally, let us delineate as a special kind of ultimate explanation what I shall call absolute explanation. An absolute explanation of *E* is an ultimate explanation of *E* in which the existence and operation of each of the factors cited are either self-explanatory or logically necessary. Other explanations cite brute facts that form the starting points of explanations; there are no brute facts in absolute explanations—here everything really is explained.

I do not believe that there can be any absolute explanations of logically contingent phenomena. For surely never does anything explain itself. *P*'s existence at t_2 may be explained in part by *P*'s existence at t_1 . But *P*'s existence at t_1 could not explain *P*'s existence at t_1 . *P*'s existence at t_1 might be the ultimate brute fact about the universe, but it would not explain itself. Nor can anything logically necessary provide any explanation of anything logically contingent. For a full explanation is, we have seen, such that the *explanandum* (that is, the phenomenon requiring explanation) is deducible from it. But you cannot deduce anything logically contingent from anything logically necessary. And a partial explanation is in terms of something that in the context made the occurrence of the *explanandum* more probable, without which things would more probably have gone some other (logically possible) way. Yet a world in which some logically necessary truth did not hold is an incoherent supposition, not one in which things would probably have gone some other way. These are among many reasons why it must be held that God is a logically contingent being, although maybe one necessary in other ways.

So for these reasons let us leave aside consideration of absolute explanation, and return to explanations of other types. I suggest that the arguments to the existence of God with which we are concerned are arguments to a complete explanation of phenomena. They all claim that God's widest intention at some time brings about certain phenomena at that time; and that nothing else at that time explains either his existence or his forming that intention. His intention involved in a complete explanation has no causal explanation at all, since he is perfectly free. Whether his existence at that time has an explanation in terms of something earlier depends on just how temporal a being God is. If God is a being who by his intention at each moment of time keeps himself in existence at the following moment, then his existence at a given moment, which provides part of a complete explanation of some mundane phenomenon, would itself have a further explanation in terms of his earlier intentions. Only a complete explanation that explains phenomena in terms of God's intention at a time, his existence for some beginningless period of time and his intention at each moment of existence to continue to exist would provide an ultimate explanation of those phenomena. If, on the other hand, God is the kind of being who is necessarily eternal in the sense that, if he exists at any time, he exists at all times, his existence at any time would have no further explanation. In the next chapter I shall argue in favour of God's essence being a necessarily eternal one.⁵ In that case any complete explanation in terms of God's intention at a time will also be an ultimate explanation. So let us now consider the grounds for supposing that some explanation is a complete explanation.

Justifying a Claim to Complete Explanation

The basic considerations in judging a proposed explanation probable are, as we have seen, the prior probability and the explanatory power of the proposed explanatory hypothesis. In so far as the proposed hypothesis renders probable, or better necessitates, the phenomenon

⁵ I am understanding 'eternal' in the sense of 'everlasting'. A third possibility is to regard God as being outside time. But in that case it is hard to see how he can be regarded as acting. I give reasons for finding this third view an incoherent view, and I develop the other two views in a full way in the last four chapters of *The Coherence of Theism* (Clarendon Press, 1993).

to be explained and the occurrence of other phenomena too that are observed to occur, and in so far as the occurrence of the phenomena is otherwise very improbable, then it has great explanatory power. It has prior probability, basically, in so far as it is simple and fits in with our general knowledge of the world. It is simple in so far as it postulates few entities and reasons (that is, laws or intentions) of a simple kind. It fits in with our general background knowledge of the world in so far as the entities and reasons that it postulates are of a kind with entities and reasons that we have reason to believe to exist and to operate in other spheres. I would suggest that our grounds for believing that objects (events, states, etc.) or reasons do not have a further explanation in terms of factors acting at the time, and so that any explanation provided by those objects or reasons is a complete explanation, are any grounds for believing that the latter could be explained themselves only by postulating causes and reasons (acting at the time) having no more explanatory power or prior probability than the *explananda*—that is, the things to be explained—or having more of one only if they have correspondingly less of the other.

I will now illustrate this claim with examples of cases where people have believed with reason that, as far as scientific explanation is concerned, they have reached a complete explanation. On the Hempelian model we explain a phenomenon *E* by initial conditions *C* and scientific laws *L*. A complete scientific explanation of *E* will cite the most fundamental laws of nature, and initial conditions that have no explanation in terms of contemporaneous states or events. I illustrate my thesis about our grounds for believing that we have reached a terminus of explanation by considering the grounds for believing that we have reached the most fundamental laws of nature.

We often explain the operation of one law by the operation of another—for example, Galileo's law of fall is explained by the operation of Newton's laws. Our grounds for accepting such an explanation will be that there is a gain of either explanatory power or prior probability without a great loss of the other. In my example there is a gain of explanatory power in that Newton's laws successfully predict both the operation of Galileo's law (to a high degree of accuracy) and the occurrence of further phenomena. It is a deductive consequence of Newton's laws and certain true statements of initial conditions (for example, that the earth has such-and-such a mass and radius) that Galileo's law operates to a high degree of approximation. Newton's laws also successfully predict the motions

of the planets, the behaviour of the tides, the interactions of colliding bodies, etc. Further, Newton's laws have high prior probability. As we also saw earlier, they are extremely simple. The question of fitting in with background knowledge hardly arises, because Newton's theory purports to cover such a large field that there is not a great deal outside it with which to compare the entities and laws that it postulates. However, the later discovery of inverse square laws of electrostatic and magnetostatic attraction (namely, laws of the same mathematical type as Newton's laws of gravitational attraction) gave Newton's theory some prior probability for this reason too. So the attempt to explain Galileo's law by Newton's laws was a move that resulted in a gain of explanatory power in no way cancelled out by any loss of prior probability. Hence Galileo's law of fall is indeed explained by the operation of Newton's laws. Newton's laws in their turn (or rather their operation where bodies have relatively large mass and relatively small velocity) are explained by Einstein's field equations of General Relativity. In passing from Newton's laws to Einstein's, there is perhaps a loss of simplicity (though Einstein himself considered that his laws had a simple form). But there is some considerable gain in explanatory power. We can derive from Einstein's laws not merely Newton's laws of mechanics, but various accurate predictions about the behaviour of light and other electromagnetic radiation; and also some successful predictions about the mechanical behaviour of bodies that are different from those made by Newton's laws.

Now my thesis is that we would be justified in believing that some law or laws were the terminus of explanation, were not to be further explained, if we had grounds for believing that any gain of explanatory power would be outweighed by a corresponding loss of prior probability and that any gain in prior probability would be outweighed by a corresponding loss of explanatory power. We would have these grounds if we already had simple laws that fitted in well with our background knowledge, and if we had grounds for believing that any attempt to amend our laws or to derive them from more fundamental laws, in order to increase the explanatory power of science, would make them or other scientific laws very complicated, or not to fit in with our background knowledge, at the expense of little gain of explanatory power. I suggest that the scientists of the later eighteenth century who held that Newton's laws were the most fundamental laws of nature had just such grounds. For the phenom-

ena that Newton's laws did not explain were light, chemical and biological phenomena, and a few miscellaneous phenomena, such as electric and magnetic attraction. Newton had outlined a plausible account of how his laws could explain light. In chemical and biological phenomena and also in the phenomena of electric and magnetic attraction there was obviously some force or forces at work that had significant effects only over very short distances. It looked as if one could cope with these phenomena by adding a fifth law concerned with such forces to the other four, without amending the latter.⁶ The fifth law would lay down the equation governing the operation of such forces, which would be subject to the general propositions about forces contained in the first three laws. Because these forces were operative only over very short distances, and were in consequence different from the gravitational force that showed its strength over longer distances, there was no reason to suppose that the law of gravity would require any amendment. So it looked as if there could be no gain in explanatory power by amending Newton's laws or deriving them from more fundamental laws. Nor, in view of their incredible simplicity (and, in so far as it arose, good fit with background knowledge), did it look as if they could possibly be improved on that score. For these reasons, the scientists of the eighteenth century were, I suggest, justified in believing that science had reached a terminus in Newton's laws.⁷ However, many new phenomena first discovered by science in the twentieth century proved not at all amenable to the kind of account that would fit in with Newton's theory. A whole range of phenomena concerned with the behaviour of light and other electromagnetic radiation, as well as subatomic phenomena, proved to be not at all as Newton's theory led people to expect. In consequence, a theory such as Einstein's General Theory of Relativity, which led us to expect many of these phenomena, as well as the gravitational phenomena, would be a theory of much greater explanatory power than Newton's, and so some loss of simplicity would be tolerable in order to obtain an overall theory. Yet my point remains that the scientists of the eighteenth century had a

⁶ For Newton's speculations on these matters, see I. Newton, *Optics*, Queries, 29, 30, and 31.

⁷ And, of course, many of them believed just this. Recall Halley's Ode prefixed to the *Principia*; and Pope's famous couplet:

Nature and Nature's laws lay hid in night
God said, Let Newton be, and all was light.

reasonable belief when they believed that Newton's laws were not susceptible of further explanation.

So then we have seen at work the criteria of prior probability and explanatory power in giving us grounds, within scientific explanation, for believing that we have reached a terminus in regress of explanation by laws. The same considerations are, I suggest, at work in giving us grounds for believing that, within scientific explanation, we have located initial conditions *C* that are not dependent for their present existence on some further state of affairs *B*. Thus, to use again the example used in Chapter 3, we may explain some phenomenon by the motion of some star *S*. Yet *S* is moving in a way that is to be expected if Newton's laws (or laws similar thereto) are the true laws of motion only if there is an unobserved planet *P* that is exerting an attractive force on *S*. It would complicate science vastly if we supposed that the laws of motion were somewhat different from Newton's, simply in order to account for the motion of *S*. It is far simpler to postulate *P*. We suppose that we have reached initial conditions that are not dependent on further states of affairs if, as far as we can see, there would be no overall gain in explanatory power or prior probability in postulating further states. To postulate *P* increases the explanatory power of science (in enabling it to explain the motion of *S*) while keeping its laws simple, and, although it adds to the entities that it postulates, the new entity is of a type (planet) well known in science—the supposition of its existence fits well with our background knowledge. However, contemporary science never postulates entities whose action is responsible for the current existence (as opposed, for example, to the motion) of distinct material bodies. Nothing distinct from *S* keeps *S* in being by its current action. But contemporary science claims this only because there would be no gain in prior probability or explanatory power if it were to claim otherwise. (Of course, as this book suggests, there may be reasons from outside science for making such a claim.)

Similar considerations arise with the problem of determining the kinds of constituents of material objects. Science postulates that observable material objects are made of unobservable constituents—for example, of molecules of various kinds linked in various ways. It postulates entities that make up material objects, and whose interaction explains the behaviour of observable material objects. Observing thousands of macroscopic substances combining in different ratios to make other substances, chemists postulated that

these substances were made of atoms of only 100 different kinds and they postulated certain laws of their interaction sufficient to account for the behaviour of the macroscopic substances. The atomic theory was such that it led us to expect a whole host of chemical phenomena, some of which were already known and some of which were subsequently discovered, phenomena that there was no other reason for expecting. Also in postulating that macroscopic substances were made only of atoms of 100 different kinds, it explained the phenomena in terms of a simpler picture. Seeking for greater simplicity, scientists naturally sought to give an account of how the atoms of different kinds differed from each other in terms of their being made of different numbers of yet more elementary building blocks of, say, two or three distinct kinds. They hoped that such an account would explain the valency of the different atoms, why they entered into just the chemical combinations that they did. Much of the subsequent history of fundamental physics is the history of the failure of such attempts, the failure to find one or two distinct elementary building blocks, such that the observable world could plausibly be regarded as composed of these and its behaviour as constituted by their interaction. At first in the early twentieth century a nice, simple picture seemed to be emerging—atoms seemed to be made of electrons and protons alone. But, alas, new phenomena turned up such that science had to postulate neutrons, photons and positrons, neutrinos, π -mesons, K -mesons, muons, and so on, in order to explain the phenomena. The variety of fundamental particles became almost as great as the variety of atoms. And then physics sought a gain in simplicity by postulating that some particles (for example, protons and neutrons) were made up of yet smaller particles, quarks. But then there turned out to be several different kinds of quarks. Of course there was a gain in explanatory power involved in postulating the various fundamental particles—various newly discovered physical phenomena were now predictable. But there was no gain in simplicity, or fit with background knowledge—though, as we have seen, this latter is hardly a consideration that arises with very fundamental theories that purport to explain everything. Physics is still devoting its attention to finding an underlying pattern in the variety of fundamental particles.

Although physics has not yet reached a situation where it can rest content, I suggest that physicists would be able to recognize grounds for believing that there were no further entities responsible for some

observed or postulated behaviour. Suppose that physics had had the success that physicists of the early twentieth century hoped for. Suppose that all known chemical and physical behaviour were predictable by postulating one positive and one negative kind of particle, of equal mass but opposite charge, out of different numbers and arrangements of which all bodies were made and whose interactions constituted their behaviour. Postulating further entities could then result in no gain of explanatory power. Nor to all appearances could it result in any gain of simplicity, the crucial element determining prior probability. For I defy anyone to imagine a simpler kind of scientific explanation of data of the kind with which these physicists were concerned. This is not to say that there could not be one; only that it is unreasonable to suppose that there could be one. Here we would have a reasonable stopping point for explanation (within science). We know with kinds of entities, as with laws, what a reasonable stopping point would be like.

So far I have phrased my points in terms of the amended Hempelian account of scientific explanation. Let me now put it in terms of the substances-powers-and-liabilities account. To account for the existence and behaviour of observable objects science postulates that objects (substances) of observable kinds have certain powers and liabilities. It postulates that objects have various powers (for example) of attracting and repelling other objects, and various liabilities to be acted on by yet other objects so as to be made to exert their powers. It postulates that there are objects additional to ones that can be seen—for example, distant planets known only by their effects on observable stars—and that observable objects are composed of various unobservable objects with certain powers and liabilities. It does all this only in order to explain what is observed. What is postulated must have considerable power to explain what is observed, and there must be greater prior probability in supposing that the postulated objects, powers, and liabilities exist than in supposing that the things to be explained exist unexplained. The basic consideration here is simplicity. The explanation is to be accepted (as we saw in Chapter 3) in so far as it postulates *few* entities (as few unobserved planets as possible), entities of few *kinds* (few kinds of fundamental particles), few and simple kinds of powers and liabilities (for example, all material bodies, not just nitrogen atoms on earth, having certain powers and having powers describable by simple formulae—for

example, the power of attracting in accordance with Newton's mm'/r^2 formula). We accept an explanation with postulated entities and properties in so far as it leads us to expect the phenomena to be explained and other phenomena, which latter would not be otherwise expected, and in so far as it provides or belongs to a simpler world picture than the one in terms of the phenomena alone. We move beyond an explanation to a new one only if we can get greater explanatory power (new things are explained) or the total world picture becomes simpler. If there are no further phenomena to be explained and the entities and properties postulated have a simplicity that would be hard to better in the ways shown by our examples, we have good grounds for believing that we have reached a complete explanation.

Similar considerations apply to personal explanation. Here, it will be recalled, we explain an event E as brought about by a person P with certain basic powers, beliefs, and intentions. The powers, beliefs, and intentions belong to the 'why' of explanation. As with scientific laws, we seek to explain each of these factors in terms of simpler factors fitting in with background knowledge with greater explanatory power. Thus, in explaining a human being's behaviour, although we may begin by postulating a separate intention for each action, we seek to postulate a number of wider-ranging intentions of the kind that other humans have such that, given them, it would be predictable that that human—given his beliefs—would have the more specific intention that he has at some time. Thus we may explain a man's intention to open the door in terms of his belief that the door's being open is a necessary condition of his going out of it, and his belief that his going out of it is a necessary condition of his posting a letter, and his intention to post a letter. The latter intention will (together with certain beliefs) explain not merely the former intention, but many other intentions that the man has on the way to his posting the letter (for example, an intention to go down the stairs, an intention to cross the road, etc.). Similarly we explain beliefs by wider and simpler beliefs—a belief that this body that feels heavy will fall rapidly if I let go of it, in terms of a belief that all bodies that feel heavy will fall rapidly if I let go of them, and the latter belief in terms of a belief that all bodies that feel heavy will fall rapidly if people let go of them. The same kind of point can be made about powers. Further, in explaining unusual phenomena (for example, books

flying about the room, as in the example discussed at the end of Chapter 3), we seek to attribute them, in so far as we can, to the agency of one person (one poltergeist) rather than many. We seek, so far as we can, to postulate intentions, beliefs, and powers that are similar, or at any rate similarly formed, to those of other people, and that thus fit with our background knowledge. Within the limits of personal explanation, we then seek a complete explanation of phenomena in terms of the agency of as few persons with the most general intentions, beliefs, and powers that fit with other postulated intentions, beliefs, and powers into as simple a picture of the agents as we can get that leads us to expect the phenomena that we find and that we would have no other grounds for expecting. Clearly for human agents we often reach a situation where we have every reason to suppose that we have reached the end of the explanatory road: intentions than which the agent has no wider or simpler intentions, beliefs, and basic powers not derivable from wider or simpler ones.

Since, within scientific and personal explanation, explanations explain other explanations and are rightly judged so to do in so far as they satisfy the criteria of prior probability (determined by simplicity and fit with background knowledge), and of explanatory power, and since, as we saw in Chapter 2, scientific and personal explanations are on a level—that is, are rivals for the explanation of phenomena—it would seem to follow that a scientific explanation could explain a personal one, and conversely; and that the criteria that it does so are any gain of prior probability and explanatory power that would result from supposing that it does. By a scientific explanation explaining a personal explanation, I do not mean the one being analysed in terms of the other—we saw in Chapter 2 that a personal explanation cannot be analysed in terms of a scientific explanation, and it is surely equally plausible to suppose that scientific explanation cannot be analysed in terms of personal explanation. What, rather, I do mean by a scientific explanation explaining a personal explanation is the existence and operation of the factors involved in a personal explanation being explained by the existence and operation of factors involved in a scientific explanation. A scientific explanation might be given of how people come to exist, and to have the intentions, beliefs, and basic powers that they have. It is the programme of physicalism to effect a reduction of just this kind. The theist who tries to explain why the world is and works as it does is attempting the reverse programme—to give a personal explanation

in terms of the action of God, of the existence and operation of the factors involved in scientific explanation.

From our detailed consideration of the criteria operative within scientific or personal explanation for supposing that objects (substances, events, states, etc.) or reasons do not have a further explanation in terms of factors acting at the time, and so that any explanation in terms of the former is a complete one, we may reasonably conclude that the criteria for supposing that factors have no further explanation (scientific or personal) in terms of factors acting at the time and so that any explanation is a complete explanation overall (not just a complete explanation within scientific or within personal explanation) are that any attempt to go beyond the factors that we have would result in no gain of explanatory power or prior probability. You reach a theory such that, if you attempt to explain the existence and operation of the factors involved in it, you always reach a theory that explains nothing further and has no greater prior probability (in particular, is no simpler) than the theory that you already have; or, if it does have more of one of these factors, it has less of the other.

Let us now bring out the significance of these points by expressing them in our symbolic notation. We have as our evidence phenomena e . e will include some things e_1 that, it is probable, explain other things e_2 , and perhaps also some things e_3 that have not been explained. We postulate explanations in terms of new causes and reasons only for things that do not have them already in terms of the phenomena that form our evidence. Hence we do not postulate new things to account for e_2 , only to account for e_1 (and e_3). We may postulate h as an explanation of e_1 . h must be such as to lead us to expect e_1 (and so e_2)— $P(e_1|h \& k)$ must have a significant degree of probability (k being tautological background knowledge). There will be a gain of explanatory power in so far as h leads us to expect e_3 as well— $P(e|h \& k) > P(e|e_1 \& k)$. There will be gain of prior probability in so far as postulating the new explanation leads to a simpler world view— $P(h|k) > P(e_1 \& e_3 \& \sim h|k)$. In so far as there is reason to suppose that there is no h that will lead to an increase either of explanatory power or of prior probability, it is probable that ($e_1 \& e_3$) constitute ultimate brute facts. But, in so far as there is an h that results in a gain of either explanatory power or prior probability without a corresponding loss of the other, we must postulate that h for which the gain is greatest. Such an h is probably true.

APPENDIX

Aquinas and Scotus on Regress of Explanation

It may be helpful at this stage to contrast the terminology that I have introduced and the results that I have reached with those of Aquinas and Duns Scotus, who are among the few philosophers of the past who devoted much thought to this matter of explaining one explanation by another, and the latter in turn by another one. To start with, Aquinas and Scotus are concerned only with the ‘what’, causes; and not with the ‘why’, reasons; and by causes they mean not events, but objects (substances), the states of which are the events. They then consider cases where we can order causes in series. *A* being caused by *B*, *B* by *C*, *C* by *D*, and so on. They distinguish series of two kinds—essentially ordered causes, and accidentally ordered causes. Scotus explains that ‘in essentially ordered causes the second depends on the first precisely in its act of causation. In accidentally ordered causes this is not the case, although the second may depend upon the first for its existence, or in some other way.’⁸ Aquinas and Scotus give as an example of a series of essentially ordered causes—stone, stick, hand, when (the motion of) the hand makes the stick move the stone. The stick depends on the hand not for its existence, but for its operation in making the stone move. They give as an example of accidentally ordered causes the series of ancestors—son, father, grandfather, great-grandfather, etc. Here each member depends on the last for his coming into existence, but not for his operation in generating—the grandfather does not make the father generate the son. Then the causes that occur in any series of full explanations will be a series of essentially ordered causes, and conversely. A series of accidentally ordered causes in explaining beginnings of existence does not fully explain present existence or operation. Scotus claimed that any series of essentially ordered causes must be a series of simultaneous causes. But this is not at all intuitively obvious. Why cannot there be causes that act at a temporal distance? Why cannot *A* make *B* do something two seconds later? Aquinas claimed that

⁸ See Duns Scotus, *Opus Oxoniense*, I, Dist. II, Q1, trans. in *Duns Scotus: Philosophical Writings*, ed. A. Wolter (Hackett, 1987), 40–1. I am indebted for this quotation and for a careful analysis of the scholastic terms that I discuss in this paragraph to P. Brown, ‘Infinite Causal Regression’, *Philosophical Review*, 75 (1966), 510–25.

natural reason could not show that there cannot be an infinite regress of accidentally ordered causes, and so could not show that the universe had a beginning at a time. I agree with him that there can be no good deductive argument to this effect, but there might nevertheless be a good inductive one.⁹ However, our concern, as was Aquinas's, in arguments to the existence of God, is with series of essentially ordered causes (and in this case with series in which other members depend on the first member, God, not merely for some operation but for their existence).

Aquinas claimed that there cannot be an infinite regress of essentially ordered causes. The present operation of *A* may be due to the operation of *B*, which may in turn be due to the operation of *C*; but this series cannot go on, he held, *ad infinitum*. Presumably Aquinas held, like Scotus, that any causes fully responsible for *E* are contemporaneous with *E*. He did not, as I have said, much consider series of reasons, but we could hold a similar thesis with respect to them. In that case we can put all this in a neat form in the form of a thesis to which, I believe, he would have assented, as the thesis that every phenomenon that has a full explanation has a complete explanation. The thesis is that, if there is a full explanation of *E* by *C* and *R*, then, if there are any factors responsible for the current operation of *C* and *R*, you can find a set of such factors such that they themselves have no explanation in terms of contemporaneous factors.

Aquinas claimed to be able to prove his thesis that there cannot be an infinite regress of essentially ordered causes on *a priori* grounds, but it is not altogether clear just what his argument is. Patterson Brown¹⁰ claims that it is as follows. Essentially ordered causation¹¹ is transitive. If *W* causes *X*, and *X* causes *Y*, then *W* causes *Y*. If *V* in turn causes *W*, then *V* causes *Y*, and so on. Brown suggests that the argument is that, as long as this series *V*, *W*, *X*, *Y*, continues backwards, we have not found the real cause of *E*. Unless we reach a first cause, we have not found the 'Aristotelian explanation' of *E*. If this is Aquinas's argument, it seems to suffer from the completist fallacy to which I alluded earlier. Surely if *C* causes *E*, *C* really does explain the

⁹ For discussion of this issue, see my *Space and Time* (2nd edn., MacMillan, 1981), ch. 15.

¹⁰ Brown, 'Infinite Causal Regression', 522.

¹¹ Brown considers in detail only the case of 'moving'—that is, 'causing to move'—but claims that his account applies to other kinds of causation. Hence I phrase my account of Brown in terms of 'causes' rather than 'moves'.

occurrence of *E*, even if *C* itself needs explanation. Consider a long railway train in which each truck makes the next truck move. The motion of the last truck is certainly fully explained by the motion of the last truck but one, even if there are other things to be explained.

I know of no good *a priori* argument for Aquinas's thesis that there cannot be an infinite regress of essentially ordered causes, and so for the thesis that any phenomenon that has a full explanation has a complete explanation. The thesis may nevertheless be true, but until it is shown to be true we should not assume that it is. The infinitely long railway train, in which each truck by its own motion simultaneously makes the next truck move, *seems* a coherent supposition. Yet, although there may be no good *a priori* argument to show that phenomena that have full explanations always have complete explanations, it may be possible to show in particular cases that it is probable that they do. This chapter has been devoted to arguing this later thesis and to setting out the grounds for judging in a particular case that some explanation is a complete explanation.

The Intrinsic Probability of Theism

We have seen that the theist claims that the various phenomena that he cites in evidence of theism require explanation, and he claims that theism allows us to understand why these occur, and is itself a much more natural stopping point for explanation than are the original phenomena. The argument of the last chapter has shown that the correctness of this claim depends on how great are the prior probability of theism and its explanatory power with respect to the phenomena. In this chapter I shall consider the prior probability of theism. If we assume that all our empirical data are among the things to be explained, then our background knowledge will be mere tautological evidence; and so our concern will be with the intrinsic probability of theism, and that, we have seen, is basically a matter of how simple a hypothesis it is.

I shall now set out the hypothesis of theism in much greater detail than I did in Chapter 1, and shall examine just how simple a hypothesis it is.¹

The Nature of God

The definition of theism given on p. 7 involves the following. There exists now, and always has existed and will exist, God, a spirit, that is, a non-embodied person who is omnipresent. I considered at the end of Chapter 2 what is meant by saying of God that he is non-embodied. In essence, to say that God is not embodied is to deny that there is any volume of matter such that by his basic actions he can

¹ This chapter largely summarizes points made more fully and with greater rigour in *The Coherence of Theism* (Clarendon Press, 1993), and in *The Christian God* (Clarendon Press, 1994), chs. 6, 7.

control only it and such that he knows of goings-on elsewhere only by their effects on it. By contrast, to say that God is an omnipresent spirit is to say that he knows about goings-on everywhere without being dependent for that knowledge on anything, and can control by basic actions all states of affairs everywhere (in this or any other universe) without being dependent for that power on anything. God is creator of all things in that for all logically contingent things that exist (apart from himself) he himself brings about, or makes or permits other beings to bring about, their existence. He is, that is, the source of the being and power of all other substances. He is, for example, responsible for the past, present, and future existence of material objects and of the natural laws that they follow, of persons and their powers. And whatever else logically contingent there may be—devils and angels, and other universes—he makes them exist and behave as they do, or sustains in other beings the power so to do. Some thinkers have held that God created the world at a first moment of its history and imposed upon it then the laws of its future operation and thereafter left it to itself. This is the view of the deist. By contrast, in developing the theist's position, I postulate the more orthodox view that God is at each moment of the world's history responsible for its operation at that moment of its history. Of course the more orthodox theist does hold that, if the universe or anything else had a beginning of existence, God it was who brought that beginning about or permitted some other being so to do. God is perfectly free in the sense (which I introduce by definition) that nothing in any way causally influences his choices. Which choices he makes, that is, which intentions he adopts, depends on himself at the moment of choice alone (though he may form a particular intention—to cure your cancer, in order to fulfil another intention—to answer my prayers).

God is omnipotent in the sense (roughly) that he can do whatever it is logically possible that he do. The qualification in the last clause is important. There are some apparent states of affairs, the description of which involves a logical contradiction—for example, me existing and not existing at the same time. God cannot bring about such apparent states, not because he is weak, but because the description 'me existing and not existing at the same time' does not really describe a state of affairs at all, in the sense of something that it is coherent to suppose could occur. There are also states of affairs that it is coherent to suppose could occur, but that it is not coherent to suppose God

could bring about, because the very description of him bringing them about does not really describe an action. An example would be 'an uncaused state of affairs'. It is logically possible that such a state occur, but it is not coherent to suppose that God could bring about, that is cause, an uncaused state.² He is omniscient, at any rate in the sense that he knows at any time whatever it is logically possible that he know at that time. (It *may* be that there are true propositions that it is not logically possible that a person know at some time *t*—for example, propositions about some other person's future free actions. Then to claim that God is omniscient is not to claim that at *t* he knows these propositions.) He is perfectly good. I understand by this (roughly—in a way to be made precise shortly) that he is a being who always does a morally best action (where there is one) and does no morally bad action.

The theist holds that God possesses the properties described in some sense necessarily, and he is in some sense a necessary being. That is to say, God could not suddenly cease to be (for example) omnipotent. While God is God, he is omnipotent; nor could he cease to be God while remaining the same individual (as, for example, the Prime Minister can cease to be Prime Minister while remaining the same person). Further, while other things exist by chance or because of the action of yet other beings, God could not not exist. His existence is not dependent on any other being. Nor is it a matter of chance. But what sort of 'could not' is this? What sort of 'necessity' is involved? It seems to me that a theist, if he is to worship a God worthy of worship, must hold that God's necessity is necessity of the strongest kind that the being described so far could possess. My account of what this amounts to is as follows. To say that some being necessarily or essentially has certain properties is to say that without these properties he could not exist. The reason why God has necessarily the properties that I have just described is that having those properties is essential to being the kind of being that God is. Let us say that φ is an essential kind if an individual who is φ cannot cease to be φ while continuing to be. To use Kripke's well-discussed example,³ a person is an essential kind. If John is a person, he could not be anything else; because, if John ceases to be a person,

² This is a very inadequate account of a difficult concept. For a more adequate account, see *The Coherence of Theism*, ch. 9.

³ See (e.g.) his 'Naming and Necessity', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (D. Reidel, 1972).

he ceases to be. Let us call a person who is omnipotent, omniscient, perfectly free, perfectly good, and creator of all things a divine being. The theist must claim that God is a being who belongs to the essential kind of divine being.⁴ He could not cease to be divine without ceasing to be God. There is no obvious incoherence in supposing that a kind within the kind of person is also an essential kind in the sense in which person is. So much for the necessity of the divine properties belonging to God. A somewhat different account has to be given of the necessity of God's existence.

To say that 'God exists' is necessary is, I believe, to say that the existence of God is a brute fact that is inexplicable—not in the sense that we do not know its explanation, but in the sense that it does not have one. As we saw in Chapter 4, any terminus to explanation of things logically contingent must be itself something logically contingent. However, as we also saw there, there are two ways in which God's existence being an inexplicable brute fact can be spelt out. The first position is to say that God's essence is an eternal essence. God is a being of a kind such that if he exists at any time he exists at all times; his existence at all remains the one logically contingent fact. The alternative position is to say that the divine essence is a temporal essence; the ultimate brute fact is not God's existing as such, but his existing for a period of time without beginning. His subsequent existence would be due to his intentional choice at each moment of time to continue to exist subsequently. Theism has traditionally taken the former position, and I shall argue in favour of it shortly. In that case God will have the strongest kind of necessity compatible with his being a logically contingent being. Such necessary existence we may term factually necessary existence (in contrast to logically necessary existence). I argued in the previous chapter that, if God's essence is an eternal essence, any complete explanation of phenomena in terms of God bringing them about is also an ultimate explanation.

The Simplicity of Theism

Such is the hypothesis of theism, as I understand it. How simple a hypothesis is it? I propose to argue that it is a very simple hypothesis

⁴ I use the expression 'divine being' to replace the clumsy expression 'personal ground of being' used in earlier editions, and in *The Coherence of Theism*. 'Divine' is now used in a slightly different sense from the sense in which it is used in *The Coherence of Theism*.

indeed. I shall begin to do this by showing how the divine properties that I have outlined fit together. A theistic explanation is a personal explanation. It explains phenomena in terms of the action of a person. Personal explanation explains phenomena as the results of the action of a person brought about in virtue of his basic powers, beliefs, and intentions. Theism postulates God as a person with intentions, beliefs, and basic powers, but ones of a very simple kind, so simple that it postulates the simplest kind of person that there could be.

To start with, theism postulates a God who is just one person,⁵ not many. To postulate one substance is to make a very simple postulation. He is infinitely powerful, omnipotent. This is a simpler hypothesis than the hypothesis that there is a God who has such-and-such limited power (for example, the power to rearrange matter, but not the power to create it). It is simpler in just the same way that the hypothesis that some particle has zero mass, or infinite velocity is simpler than the hypothesis that it has a mass of 0.34127 of some unit, or a velocity of 301,000 km/sec. A finite limitation cries out for an explanation of why there is just that particular limit, in a way that limitlessness does not. As I noted in Chapter 3, scientists have always preferred hypotheses of infinite velocity to hypotheses of very large finite velocity, when both were equally compatible with the data. And they have always preferred hypotheses that some particle had zero mass to hypotheses that it had some very small mass, when both were equally compatible with the data. There is a neatness about zero and infinity that particular finite numbers lack. Yet a person with zero powers would not be a person at all. So in postulating a person with infinite power the theist is postulating a person with the simplest kind of power possible.

God's beliefs have a similar infinite quality. Human persons have some few finite beliefs, some true, some false, some justified, some not. In so far as they are true and justified (or at any rate justified in a certain way), beliefs amount to knowledge. It would seem most consonant with his omnipotence that an omnipotent being have beliefs that amount to knowledge. For, without true beliefs about the consequences of your actions, you may fail to realize your

⁵ If the existence of that person entails the existence of other divine persons, a possibility discussed in Additional Note 1, the original hypothesis is no less simple for that. A simple hypothesis is none the less simple for entailing complicated consequences. But a hypothesis of three independent divine beings would be much more complicated than theism.

intentions. True beliefs fail to amount to knowledge only if they are true by accident. But, if the divine properties are possessed necessarily, God's beliefs could not be false, and so could not be true by accident. And, if an omnipotent being has knowledge, the simplest such supposition is to postulate that the omnipotent being is limited in his knowledge, as in his power, only by logic. In that case he would have all the knowledge that it is logically possible that a person have—that is, he would be omniscient.

For a person to act, he has to have intentions. A person could be omnipotent in the sense that whatever (logically possible) action he formed the intention to do, he would succeed in doing, and also omniscient so that he knew what were all the (logically possible) actions available to an omnipotent being in his situation, and yet be predetermined to form certain intentions. His intentions might be determined by causal factors outside his control, or at any rate, as are those of humans, greatly influenced by them. But, if a person is predetermined (or has an inbuilt probabilistic tendency) to act in certain specific ways, this means that a tendency to act in a particular way is built into him. But a person with an inbuilt detailed specification of how to act is a much more complex person than one whose actions are determined only by his uncaused choice at the moment of choice. Such a being I call a perfectly free being. Theism in postulating that God is perfectly free makes the simplest supposition about his choice of intentions.

A substance who is essentially omnipotent, omniscient, and perfectly free is necessarily a terminus of complete explanation. For, if some state of affairs *E* is explained as brought about by God in virtue of his powers and beliefs and intentions to bring about *E*, how can the action be further explained? God's powers derive from his omnipotence, his beliefs from his omniscience, and his intention, if it derives from anything, can derive only from some wider intention of his. God's widest intention has no explanation except that he chose this intention—it follows from his perfect freedom that reason alone influences him to make that choice. God's being omnipotent, omniscient, and perfectly free is involved in his existing, given that, as we have supposed, these qualities belong to the divine essence. But his existing cannot be due to any contemporaneous factor that makes him exist or allows him to exist. For, if his existence depended on some factor apart from himself, that factor could not depend for its existence on himself (for one cannot have

causation in a circle).⁶ But, if this factor did not depend on God, then God would not have been able to make it exist or not exist, and so would not be omnipotent. He is necessarily a terminus of complete explanation. It is clearly simpler to suppose that the ultimate principle of explanation, the final source of things, has always been the same rather than to suppose that only, for example, in 4004 BC did God come to be and reign—and so to suppose that God has existed eternally. But, unless God's essence is an eternal essence, God's eternity will arise, not just from God existing, but from his existing throughout a period of beginningless time. This is a more complicated supposition—it brings when God exists into the explanation of phenomena. But, if God's essence is an eternal essence, then any complete explanation of phenomena in terms of God's agency is also an ultimate explanation. For God's existence at a time is entailed by his existing at all, and does not require to be explained in terms of his previous existence and previous choices. So the simplest kind of God is a factually necessary one, in the sense defined earlier.

I argue next that God's possession of the other properties ascribed to him—being an omnipresent spirit, being creator of all things, and (given a certain highly plausible assumption) being perfectly good all follow from his being omnipotent, omniscient, and perfectly free. His possession of the first two properties is easy enough to show. If God is omnipotent, then he must be able to control by basic actions all states of affairs everywhere. If God is omniscient, he must know what is going on everywhere. If he depended for this knowledge on the operation of nerves or eyes, then, if they were to behave in unusual ways, he would lack knowledge. But since, *ex hypothesi*, God's omniscience belongs to his essence, this could not happen. Hence God is an omnipresent spirit. Since God is omnipotent, then he could prevent anything from happening if he so chose. So whatever happens happens because he makes it or permits it to happen. Hence he is the creator of all things in the sense that I delineated.

Further, if one takes a certain view about the status of moral judgements, God's perfect goodness follows deductively from his omniscience and his perfect freedom. The view in question is the view that moral judgements to the effect that this action is morally good and that one is morally bad are propositions that are true or false. The truth of this view is, of course, a contentious philosophical

⁶ This argument is subject to a qualification discussed in Additional Note 1.

issue,⁷ but it is highly plausible. Surely the person who says that there was nothing morally wrong in Hitler's exterminating the Jews is saying something false. For reasons of space I shall assume rather than argue for the view that moral judgements have truth values. But if they do not have truth values, it would be misleading to call perfect goodness a *property* of God, for it would be neither true nor false to say of him that (for example) he does no morally bad acts. If my view is correct, it follows that an omniscient being will know the truth value of all moral judgements—that is, will know of all moral judgements whether or not they are true or false. I now proceed to argue further that necessarily an agent who is perfectly free (that is free in the sense that nothing in any way causally influences which choices he makes) will do what he believes to be the morally best action or one of equal morally best actions, and will do no action that he believes to be morally bad. Thence it will follow that, if this agent is also omniscient, he will do the morally best action (if there is one) or one of the equal morally best actions (if there are such), and no morally bad action—for necessarily his beliefs about their status will be true ones.

To do an action an agent has to have a reason for acting. A movement brought about by an agent would not be an action unless the agent had some reason for bringing it about. The reason may be simply just to do that action, but normally an agent will have some further purpose in doing an action. Having a reason for an action consists in regarding some state of affairs as a good thing, and the doing of the action as a means to forwarding that state, and hence itself a good thing. If my reason for going to Oxford was to give a lecture, I must regard it as in some way a good thing that I give the lecture, and so a good thing that I go to Oxford. If I regarded it as in no way a good thing that I give the lecture, if I thought that giving the lecture was an event that would serve no useful function at all, giving the lecture could not have been my reason for going to Oxford. The point that to do an action I must (of logical necessity) see my performance of it as in some way a good thing is a very old one due to Aristotle, emphasized by Aquinas, and re-emphasized in our day by, among others, Stuart Hampshire.⁸ God, like man, cannot

⁷ I argue for it in *The Coherence of Theism*, ch. 11.

⁸ 'A man cannot be sincere in accepting the conclusion that some course of action is entirely mistaken, if he at the same time deliberately commits himself to this course of action' (S. Hampshire, *Freedom of the Individual* (Harper & Row, 1965), 7).

just act. He must act for a purpose and see his action as in some way a good thing. Hence he cannot do what he does not regard as in some way a good thing. This is not a physical constraint, but a logical limit. Nothing would count as an action of God unless God in some way saw the doing of it as a good thing.

Now for many actions there are reasons for doing them and reasons for not doing them; in some ways it is good that the agent should do them, and in some ways it is good that he should refrain from doing them. It is good that I should watch the television, because I would enjoy doing so; yet bad because it will stop me reading a book. It is good for governments to lower taxes, because that will give people more money to spend; and yet bad because lowering taxes will promote inflation and social inequality. Frequently, perhaps normally, there is no objective scale in which competing reasons can be weighed; one cannot say that doing *A* is on balance better than refraining from doing *A*, or conversely. In such a case a person who does *A* need be no less sensitive to objective values than one who refrains from doing *A*. But sometimes competing reasons can be compared objectively; clearly sometimes doing *A* is overall better than refraining from doing *A* (better when all reasons are taken into account), or conversely. I understand by one action being morally better than another that it is overall better. I understand by an action being morally good that it is overall better to do it than to refrain, that there are overriding reasons for doing it; and by an action being morally bad that it is overall better to refrain from doing it than to do it, that there are overriding reasons for refraining. Given that sometimes a balance of reasons makes it better to do one action than to refrain, better to refrain from another action than to do it, better to do a third action than a fourth, there will be truths about which actions are morally good, morally bad, or better than others. (Since all my subsequent discussion is concerned with moral goodness and badness, I shall often omit the 'morally' in future.) Sometimes there is one action that is the best action to do. Today it might be the best thing for me to do to go for a walk, both because I will enjoy it and because it will make my subsequent work of better quality. Sometimes the best action is one that is also morally obligatory, one that I would be culpable for not doing. Keeping promises and telling the truth are (at least under normal circumstances) obligatory; and normally it is the best action to fulfil any

obligation.⁹ But the best action need not be an obligatory one. It may be the best action for a soldier to do, to give his life to save his comrade; but he has no obligation to do this—it is a supererogatory act. Agents are meritorious for doing supererogatory acts, but not culpable for not doing them. Sometimes there is no one best action. It may be a best action to give £1,000 (all that I can properly spare) to an educational charity, but an equal best action to give the money to a health charity; and I cannot give it to both. An agent whose power, unlike ours, is in some respect unlimited may often be in a situation where not merely is there no best or equal best action for him to do, but, for every good action that he could do, there is an incompatible better action that he could do instead. Consider an artist who can create as many great paintings as he chooses. Plausibly, however many he creates, it would be better if he created more. Even if he creates an infinite number, he could still create more. All that he can do in this situation is a good action, even though there is a better one.

But it may be that the infinite series of possible actions fall into kinds, such that it is better that the agent do some action of one kind (perhaps beyond some minimal level of goodness) than that he do any action of any other incompatible kind. Suppose that the painter can compose symphonies as well as paint paintings, and with symphonies as with paintings the more the better. And suppose also that it would be better for the painter to create at least some symphonies as well as paintings, rather than merely any number of paintings. In that case he can do a best kind of action, create both symphonies and paintings, even if there is not a best action of that kind. Or it may be that there are two or more kinds of action (neither better than the other) such that it is better that the agent do an action of one of these kinds (perhaps beyond some minimal level of goodness) than that he do any action of any other incompatible kind. Then the former kinds will be equal best kinds of action, of either of which kinds there may be no best action. Suppose that the painter can paint

⁹ Occasionally, an agent may be subject to conflicting obligations, only one of which he can fulfil. Thus he may be obliged to repay two debts when he has the money only to repay one. In these circumstances there will be an obligation to fulfil which will not be the best action, for it will be equally good or better to fulfil the other obligation. For discussion of this situation, see my *Responsibility and Atonement* (Clarendon Press, 1989), 35–7. God, however, will never be in this situation, because obligations are obligations to other beings; and God, as the source of the existence of all other beings, will (as a best act) ensure that he never puts himself in a situation when he is under an obligation to other beings that he cannot fulfil.

as many paintings as he chooses, and *either* compose as many symphonies as he chooses *or* write as many novels as he chooses, but cannot write novels as well as composing symphonies. Then he has a choice of three kinds of action—just painting paintings, painting paintings and composing symphonies, and painting paintings and writing novels. Suppose now that it is better that he do some action of one of the two latter kinds than any action of the former kind; but it is never better to do some action of one of the latter kinds than to do any action of the other latter kind. In that case there will be two equal best kinds of action, but no best of any kind.

We have seen that an agent has to have some reason if he is to do an action *A*, to see doing *A* as in some way a good thing. Can an agent still do action *A* even if he judges that on balance it would be better to refrain from doing *A*? What are we to make of the suggestion that someone might see doing *A* as a good thing in one way (for example, by its giving sensual pleasure to himself), refraining from doing *A* as a good thing in another way (for example, by its contributing to the lifelong peace of mind of someone else), and refraining from doing *A* as overall a better thing than doing *A*, but nevertheless do *A*. When it is suggested that a case is of this sort, we may well suspect that it is not, that the agent did not really see refraining from doing *A* as overall a better thing than doing *A*. Yet we are sometimes prepared to allow that a situation is of this kind. We do seem to allow the possibility that someone might do an action that he regarded as a good thing only in some respect, but on balance a bad thing. But, although we allow this possibility, we do feel that some further explanation is called for. If someone really does accept that to refrain from doing *A* would on balance be better than to do *A*, he recognizes that he has adequate reason for refraining from doing *A*, but inadequate reason for doing *A*. Rational considerations point clearly in one direction, and yet the agent goes in the other direction. Yet to say that someone recognizes that he has a reason for doing something is to say that, if there are no equally good reasons for not doing that thing and if no factors other than reasons influence him, he will do that thing. We would not understand an agent who claimed to recognize 'overriding reason' for refraining from doing *A* rather than doing *A* and also claimed to be uninfluenced by anything other than the reasons that he acknowledged, and yet did *A*. For, if the latter claim is taken at its face value, what on earth can the agent have meant when he said that he recognized 'overriding reasons' for

refraining from doing *A*? Not what we normally mean, for normally to recognize a reason for doing something involves acknowledging an inclination *ceteris paribus* to do that thing. So to say of someone that he recognizes that he has overriding reasons for refraining from doing (or for doing) action *A* is to say that, in so far as no factors other than reasons influence him, he will refrain from *A* (or do *A*, as the case may be). If you said that you recognized that overall it would be better for you to go home rather than to go to the cinema, and then you went to the cinema, we should have to suppose either that you were lying or had changed your mind, or that factors other than reasons influenced what you did. An explanation of your behaviour is needed, not only in terms of what you believed about the relative merits of the actions, in terms, that is, of reasons, but also in terms of desires, passive inclinations to act that led you to do what you did not recognize adequate reason for doing. If someone has strong desires, it makes sense to suppose that he recognizes refraining from *A* as overall better than doing *A*, but nevertheless intentionally does *A*. Such non-rational factors over which the agent does not have control explain 'weakness of will', a person acting 'against his better judgement'. But the suggestion that someone might see refraining from *A* as overall better than doing *A*, be subject to no non-rational influences inclining him in the direction of doing *A*, and nevertheless do *A* is incoherent.

It follows from all this that an agent subject to no non-rational influences, that is, a perfectly free agent, can never do an action if he judges that overall it would be worse to do the action than to refrain from doing it. Hence he can never do an action that he judges to be overall a bad action, and especially one that he judges to be a morally wrong action. A perfectly free agent will always do any action that he believes to be the best action available to him. If he believes that there are a number of equally good incompatible actions open to him, all better than any other incompatible actions that he could do, he will do one of the former. They are actions that he believes to be equal best actions. But when the agent has before him an infinite number of possible actions, of each of which he believes that it is less good than another, but he believes that there is no best or equal best, his perfect freedom does not entail which of these he will do. However, it may be that he believes that the incompatible actions in this infinite series fall into kinds, such that it is better that he do any action of some one kind (at least beyond some minimum level of goodness) even though

there is no best action of that kind than that he do any action at all of any incompatible kind. In that case I suggest that he will do some action of that best kind beyond the minimum level, for he has more reason for doing such an action than for doing any action of any kind incompatible therewith. Doing an action of the former kind would instantiate a qualitatively superior kind of goodness; it would not be better than other actions merely in a quantitative respect. If there is, he believes, no best kind of action, but two or more kinds of action, such that it is not better that he do some action of one of these kinds (beyond a minimum level) than that he do any action of another of these kinds; but that it is better that he do some action of one of these kinds (beyond a minimum level) than that he do any incompatible action of any other kind, then I suggest—for the same reasons—that he will do some action of one of the former (equal best) kinds (beyond any minimum level). All of this constitutes the logical limits on which actions a perfectly free agent (in the sense in which I have defined this term) can do. If it is suggested that a really free agent could do what he regarded as evil just as well as what he regarded as good, the answer must be that in that case what he did would not be an intentional action, would not be something he did and meant to do for reason but was simply a causeless reaction. In writing earlier of God's intentions as not having an explanation, I must be understood as saying that they have no causal explanation—no other entity in any way influences his choice, nor does any earlier state of himself; but still God will be guided by the merits of possible choices—that is, by reason. Yet so often, perhaps almost always, reason provides no unique best action for God to do.

Given that moral judgements have truth values, an omniscient person will know them. His judgements about which actions are morally bad and which actions are morally good will be true judgements. Hence a perfectly free and omniscient being can never do actions that are morally bad, and will always do the best action, or an equal best action, or a best kind or an equal best kind of action (if there are these). But otherwise all that follows from his nature is that he will do a good action. Since there could not be a morally better being than a being of this kind, it is surely right to call such a being a perfectly good being. So God's perfect goodness (in the sense that I have now made precise) follows from his omniscience and his perfect freedom. I conclude that theism postulates one person of a very simple kind—a person who is essentially omnipotent,

omniscient, and perfectly free and eternal. Such a being will necessarily be an omnipresent spirit, creator of all things, and perfectly good.

The hypothesis of theism postulates not merely the simplest starting point of a personal explanation there could be (simpler than many gods or weak gods), but the simplest starting point of explanation for the existence of the universe with all the characteristics that I shall be analysing. We shall see this in detail in subsequent chapters, but the basic point is this. A scientific explanation, will have to postulate as a starting point of explanation a substance or substances that caused or still cause the universe and its characteristics. To postulate many or extended such substances (an always existing universe; or an extended volume of matter-energy from which, uncaused by God, all began) is to postulate more entities than theism. The simplest scientific starting point would be an unextended point. This, however, would have to have some finite amount or other of power or liability to exercise it (since what it will create would not be constrained by rational considerations), and so it would not possess the simplicity of infinity.

Furthermore, if some actual or postulated entity other than God is to provide a complete (or ultimate) explanation of phenomena, it needs to have added to it (in the case of a person) specific powers, beliefs, and intentions, or (in the case of an inanimate substance) specific powers and liabilities to exercise them. We need both the 'what' that causes, and the 'why' it causes. The advantage of theism is that the mere existence of God provides most of that extra 'why'. The powers and beliefs of a God are part of his simple nature. And his perfect goodness constrains the intentions that he will form—he will, as we have seen, always do the best or equal best action or kind of action in so far as there are such, and no bad action. God chooses to bring about what he does in virtue of seeing the goodness of things; and, in so far as that still gives him an enormous choice of what to bring about, he chooses by a 'mental toss up'. Thus for the theist, explanation stops at what, intuitively, is the most natural kind of stopping place for explanation—the choice of an agent. We ourselves make choices, and it seems to us as we do so that we are the source of one state of affairs coming about rather than another. Of course there *may* be some explanation of why we make the choices that we do. But we understand what is happening without having to make that supposition. Hence we have a familiar concept of an agent's bringing

about through his choice the diversity of things, which it is natural to use in this context. It follows that the very existence of God entails most of the other elements involved in a full personal explanation of phenomena, requiring only the addition of his intention at the time (the limits to his possible intentions being set by his existence as a necessarily perfectly good being). Such a full explanation will also, we have seen, be an ultimate explanation.

In the case of any actual or postulated inanimate substance, there is no reason to expect it to have the liability to exercise any powers it might have in this way or that way. That is, there is no reason to expect an always existing universe or a universe-creating entity to be of any one kind rather than any other, create a universe of any one kind rather than any other. By contrast, God's goodness (which follows from his other properties) will lead us to expect to find a universe of one kind rather than another. And if, as I shall be arguing, the actual universe is of a kind we would expect to find, then theism will have considerable explanatory power. To postulate a rival hypothesis that has the same explanatory power, we would have to complicate the bare hypotheses that I have been discussing (an always existing universe, or an unextended point from which all began) by supposing that they had the requisite extra properties (always being of a certain kind, or creating a universe of a certain kind) for no very good reason—that's just how it is. So, even if some rival actual or postulated substance was as such as simple as theism, it would have to be made a lot more complicated in order to have as much explanatory power as theism.

Note the nature of the connection that the theist postulates between God's personal causation and scientific causation. God is omnipotent. His power is not dependent on brain or nerves. His intentions are immediately operative—because that is how things ultimately are. Hence the existence of matter and the operation of natural laws.¹⁰ There is a simple connection between the factors cited

¹⁰ Theism is certainly not committed to the (to my mind manifestly false) view of occasionalism that physical conditions or substances never produce effects, but that all effects are brought about by rational agents, normally God. The occasionalist claims that the ignition of gunpowder does not really cause explosions; it is just that, when people ignite gunpowder, God reliably causes explosions. The far more plausible view advocated in the text is that God causes the operation of scientific causality; he reliably causes the ignition of gunpowder to cause explosions. Not merely does occasionalism seem to fly in the face of an obvious datum of experience—that physical objects (or their states) often cause events; but it would be self-defeating for a theist who wishes to argue to God from the

in this personal explanation and what they effect—God’s intention to bring about φ is followed by φ . Among the things that God brings about are the existence of human persons with brains, nerves, muscles, etc. and the natural laws that determine when, by what route, and within what range their intentions are efficacious. There is a similar simple connection between God’s knowledge and the world. If p is a true proposition that it is logically possible that a person know, then God knows p . Although certain physical conditions of the brain need to occur if human agents are to have knowledge of the external world and intentions that are efficacious, the human model suggests a simpler model in which such limitations are removed.

One final feature of great importance about the hypothesis of theism is this. What is at stake in the various arguments that we shall be considering is whether we ought to go beyond various phenomena to postulate a God who brings them about. This is a matter of whether the hypothesis of theism has sufficient prior probability and explanatory power. But, if it has, there is no similar issue of whether we ought to go beyond theism in order to provide a complete explanation. For, if theism is true, then, of logical necessity, God’s action provides a complete and ultimate explanation of what it explains. For, as we saw earlier, it follows from God’s omnipotence and perfect freedom that all things depend on him whereas he depends on nothing.¹¹ If God features at all in explanation of the world, then explanation clearly ends with God.

So then theism has very considerable simplicity. Simplicity is the major determinant of intrinsic probability. We saw that the other determinant is narrowness of scope. It is not quite clear how we are to assess theism on this criterion. Like its rival physicalism, theism is a theory of very wide scope purporting to explain the universe and all its characteristics. However, I argued earlier that it is clear from

physical universe and its characteristics to deny that physical objects often cause events. For we derive our understanding of what is evidence that some object causes some event by extrapolation from innumerable mundane situations in which apparently some physical object causes some event. If there really is no causation of the sort we suppose in these situations, our criteria of what is evidence for causation are misleading, and so we would be in error to use them to conclude that God is the cause of the existence of the universe or anything else.

¹¹ Except perhaps something that in turn depends on him—see Additional Note 1 on the Trinity.

examples that the great simplicity of a wide hypothesis outweighs by far its wideness of scope in determining intrinsic probability. Perhaps it seems *a priori* vastly improbable, if one thinks about it, that there should exist anything at all logically contingent. But, given that there does exist something, the simple is more likely to exist than the complex. Hence with k as mere tautological evidence and h as the hypothesis of theism, even if $P(h|k)$, the intrinsic probability of theism, is low, it will not be nearly as low of $P(h_n|k)$ for many other hypotheses h_n about what there is. The intrinsic probability of theism is, relative to other hypotheses about what there is, very high, because of the great simplicity of the hypothesis of theism.

The Explanatory Power of Theism: General Considerations

Summary of the Argument so far

I argued in Chapter 3 that the probability of a hypothesis h on evidence e and background knowledge k is a function of its prior probability ($P(h|k)$) and its explanatory power $\left(\frac{P(e|h \& k)}{P(e|k)}\right)$. By Bayes's theorem:

$$P(h|e \& k) = \frac{P(h|k)P(e|h \& k)}{P(e|k)}.$$

Now let h be the hypothesis of theism, that there is a God; let k be mere tautological evidence (and so $P(h|k)$ is the intrinsic probability of theism), and e the evidence cited in arguments for and against theism.

We saw in Chapter 3 that by the 'relevance' criterion an argument from e to h is a good C-inductive argument if (and only if) $P(e|h \& k) > P(e|k)$ and that this will be the case if (and only if) $P(e|h \& k) > P(e|\sim h \& k)$. Hence the occurrence of certain phenomena will confirm—that is, raise the probability of—the existence of God if and only if it is more probable that those phenomena will occur if there is a God than if there is not. How probable the phenomena will make the existence of God will depend on just how high or low are the probabilities (on the right-hand side of Bayes's theorem) that we will be discussing. The force of the arguments from e to h will depend, as well as on the constant factor ($P(h|k)$), on the explanatory power of theism with regard to those phenomena; how much more likely does the existence of God make the occurrence of those phenomena than it would be if we do not assume the existence of God.

We saw in Chapter 3 that the major determinant of the intrinsic probability of h was its simplicity. We saw in the last chapter that, although $P(h|k)$ may be low, it is significantly greater than that of so many alternative fillings for h . Theism postulates a God of infinite power, knowledge, and freedom. Hence theism forms a natural stopping point for explanation, a natural candidate that is for a brute fact that explains other things, but itself has no explanation. A stopping point for explanation is, of course, a highly mysterious thing. That there should be anything (logically contingent) at all is overwhelming strange, when we think about it. But there is something logically contingent. The issue is whether the world and its operations are the stopping point, or whether we must go beyond the world to find the stopping point of explanation; or whether, although we go beyond the world to find an explanation of the world, there is no stopping point for explanation to be found. By the argument of the last two chapters, the God of theism is a good candidate for a stopping point, in so far as the world and its operations (e), if not sustained by God, have a significantly lower intrinsic probability ($P(e \& \sim h|k)$) than the existence of God ($P(h|k)$), so that the former is much more in need of explanation than the latter, and, if the hypothesis of theism makes it more likely that e will occur than it would be if theism were false, ($P(e|h \& k)$ exceeds $P(e|\sim h \& k)$).

More formally— $P(e|k)$, the intrinsic probability of e , is the sum of two other probabilities:

$$P(e|k) = P(e|h \& k)P(h|k) + P(e|\sim h \& k)P(\sim h|k).$$

The first conjunct on the right-hand side of this equation ($P(e|h \& k)P(h|k)$) simply repeats the top line of Bayes's theorem. The second conjunct ($P(e|\sim h \& k)P(\sim h|k)$) is the probability that there is no God and nevertheless e occurs. If this second term has some value greater than 0, the bottom line of the right-hand side of Bayes's theorem will inevitably be greater in value than the top line, and so the whole (top line divided by bottom line) will be less than 1—that is, it will not be certain on e that h . Whether $P(h|e \& k)$ is very large (very close to 1) or very small (0 or very close to 0) will depend on whether the second conjunct is very small or very large relative to the first conjunct. It will be very large if both of the terms ($P(e|\sim h \& k)$ and $P(\sim h|k)$) are large, and smaller in so far as they are smaller. $P(\sim h|k)$, the intrinsic probability that there is no God,

is $1 - P(h|k)$. If, as I suggested, $P(h|k)$ might be very small (so improbable is it *a priori* that there should exist anything at all), $P(\sim h|k)$ may be very large. So the value of $P(e|k)$ will turn crucially on $P(e|\sim h \& k)$, the probability that, if there is no God, we will have the evidence that we do. I shall be considering the value of that probability with respect to the different e in subsequent chapters. e could occur even if $\sim h$, if either some being or beings inanimate or personal yet less great than God brought e about, or e occurred uncaused as a brute fact. The earlier equation expressing the prior probability of e can be rephrased to express the point that that probability is the sum of the prior probabilities of the various theories in the field, each multiplied by the probability of e given that theory. So:

$$P(e|k) = P(e|h \& k)P(h|k) + P(e|h_1 \& k)P(h_1|k) \\ + P(e|h_2 \& k)P(h_2|k) + \dots$$

where h, h_1, h_2 , etc. are such that one and only one such theory must be true. Among these theories will be the theory that e has no cause. I shall be arguing in subsequent chapters, with respect to different e , that there is a very low probability that e would occur uncaused, or caused by some lesser being or beings (some initial singularity in space-time, or some committee of lesser gods). The grounds in each case will be those discussed in Chapter 4—both would be too complicated (too un-simple) to form (at all probably) a stopping point for explanation. And the greater for this reason of simplicity is $P(h|k)$, relative to the prior probability of rival hypotheses that allow the existence of logically contingent things, the greater will be the posterior probability of theism— $P(h|e \& k)$. Everything turns on how much more probable it is intrinsically that if there exists anything at all God exists rather than no God and something else instead.

Which Worlds God is Likely to Create

The remaining term is $P(e|h \& k)$, the probability that if there is a God we would find the evidence we do. Now, if there is a God, e may occur either because God brings e about or because God creates some creature whose behaviour is not predetermined and has the power to

bring *e* about. The obvious case of the latter is where that creature has libertarian free will, that is, the freedom to choose whether or not to bring about some effect (such as *e*), where the totality of causes that influence him (making it harder or easier for him to make a particular choice) do not totally determine how he will choose. God is by definition omnipotent; he can bring about any state of affairs that it is logically possible for him to bring about. Among states of affairs that it is not logically possible for God to bring about is that humans always freely do what is good, when they have the free will at the moment of choice to choose between good and evil independently of the causes influencing them. Whether God brings about any state of affairs that it is logically possible for him to bring about depends on whether he chooses to do so. It follows from the account of God that I developed from the definition in the last chapter that he will do any action that is the best action to do, one of any set of alternative equal best actions, or—if there are no best or equal best actions—some action of a best kind or equal best kind; or—if there are none of these—some good action and no bad action.

Our understanding of what is good and bad is very limited. Some actions may be good or bad because of intrinsic qualities that they possess to which we as morally imperfect beings are totally insensitive. Some actions may be good or bad because of consequences that they have but of which we, as beings of very limited knowledge and intelligence, have not the slightest notion. Yet clearly most of us have some understanding of moral values. When we judge that it is good for us to feed the starving and help the weak, wrong to tell lies and break promises (all of this at any rate under normal circumstances), we make true moral judgements. And we are able to judge to some extent whether these actions would be good or wrong for us to do, as the case may be, if we were beings of different kinds—if we were very powerful or had created the people who are now starving. We reach judgements of general moral principle by reflecting on particular cases, and considering the grounds on which we judge that this action was bad and that one supererogatory. And then we can see whether the goodness (or whatever) of the action depends on certain features of the circumstances of the agent and time of his action, or whether the action would be good for any agent to do at any time. We can see, for example, that, although it might be good for me and perhaps also for you to punish a child of mine for breaking your window, it would not be good for a mere stranger to take upon

himself that task. We can see too that my obligation to keep a promise I have made to you would be entirely unaffected by how powerful I was; the obligation would remain even if I was omnipotent. Even on individual such matters we could be mistaken. Our understanding of most other things discussed in this book, and in most books about most things, is very limited and prone to error, but is such that we can grow in it. We have to make tentative judgements in the light of our understanding at the time of our investigation—in this matter as in all matters—bearing in mind the possibility of future revision. But it is wildly implausible to suppose that our understanding of what is morally good and bad is totally in error. And, if it was, I cannot see that we would have a concept of *moral* goodness at all.¹

So, given some idea of moral goodness, we have some idea of the kinds of world that God, if there is a God, would be likely to bring about. If there is in some situation a unique best action, he will do it. God will therefore, I suggest, always keep his promises and tell the truth. But I suggest that it is not generally the case that there is before God a unique best action, or a set of incompatible equal best actions. If there could be a unique best of all the possible worlds that (it is logically possible) God could create, it would be a unique best act to create it. But, contrary to Leibniz,² there could not be such a world.³ For suppose that there is such a world, *W*. *W* will presumably contain a finite or infinite number of conscious beings. Would a world be a worse world if, instead of one of these conscious beings, it contained

¹ See my *Responsibility and Atonement* (Clarendon Press, 1989), ch 1.

² 'Now this supreme wisdom, united to a goodness that is no less infinite, cannot but have chosen the best . . . If there were not the best among all possible worlds, God would not have produced any. I call "world" the whole succession and the whole agglomeration of all existing things . . . There is an infinitude of possible worlds among which God must needs have chosen the best' (G. W. Leibniz, *Theodicy*, trans. E. M. Huggard (Routledge & Kegan Paul, 1951), 128). Like Leibniz, I have been using the word 'world' in this very general sense. In this sense, God is part of the world. So my and Leibniz's talk of God making this or that world is to be read as God bringing it about that the world in which he exists is a certain way in other respects. By contrast, I am using the word 'universe' as synonymous with 'physical universe' in the narrower sense defined on p. 133.

³ Leibniz assumes (mistakenly in my view) that any best of all possible worlds would be one that it is logically possible for God to create (e.g. that its goodness does not depend on the uncaused good choices of free creatures). That there is a better world than the world that God has created was the view of Aquinas. See *Summa Theologiae*, Ia.25.6 ad 3: 'God could make other things, or add other things to those he has made, and there would be another and better universe.'

another with the same properties—if, instead of Swinburne, it contained a counterpart of Swinburne who wrote an exactly similar book and in other ways had exactly similar properties and did exactly similar actions? Surely not. But then there will be no unique best of all possible worlds that God could create. If there could be *a* best of all possible worlds that God could create, that is, a world such that no world is better than it (although other worlds may be equally good), then it would be an equal best act to create such a world. But it seems almost equally implausible to suppose that there could be such a world. For again take any world *W*. Presumably the goodness of such a world, as I shall argue in more detail later, will consist in part in it containing a finite or infinite number of conscious beings who will enjoy it. But, if the enjoyment of the world by each is a valuable thing, surely a world with a few more conscious beings in it would be a yet more valuable world—for there would be no reason why the existence of the latter should detract from the enjoyment of the world by others—they could always be put some considerable distance away from others, so that there was no mutual interference. I conclude that it is not, for conceptual reasons, plausible to suppose that there could be a best or equal best of all possible worlds that God could create, and in consequence God could not in creating a world be doing a best or equal best action.⁴ But it is highly implausible to suppose that merely for that reason a God would not have created anything at all.

⁴ Even if there were a best of all possible worlds that God could create, God would not be under any obligation to create one. See R. M. Adams, 'Must God Create the Best?', *Philosophical Review*, 81 (1972), 317–32. For God need not wrong any one by creating a world that was not a best. There might be certain beings in a best of all possible worlds that do not exist in the world that God actually makes. Yet, Adams points out, God would not do wrong to those beings in not creating them—for you cannot do any wrong to a being that never has existed and never will exist. There might be other beings in the world that God actually makes who would not exist in a best of all possible worlds—God would hardly do them any wrong by creating them, so long as he created them so as to live a life that was better to live than to live no life at all. Finally, there might be some beings who are less perfect in the world that God actually makes than they would be in a best of all possible worlds. Would God wrong them by creating them thus? So long as God has created them now in such a condition that it is better for them to exist than not to exist, in what way does he wrong them if he makes them less perfect than he could have made them? For he still gives them a reasonable existence that they would not otherwise have. Even if we suppose that a given embryo, if interfered with sufficiently early, could grow into a rabbit rather than a goldfish, and that this technique is available to a breeder, there is surely nothing wrong in breeding goldfish. Although, *ex hypothesi*, the breeder could have made those goldfish into rabbits, the breeder does the goldfish no wrong in not having made rabbits out of them.

We can also conclude that God will not do any action that is overall bad. If he brings about suffering, or permits other agents to do so, it must be that bringing about or permitting that suffering serves a greater good that could not be achieved without it, and God must have the right to impose that suffering on the sufferer. I shall be arguing in Chapter 11 that suffering does sometimes in this way serve a greater good, and that God does have limited rights to impose or permit suffering. But I shall also claim that he does not have the right to impose or permit unlimited suffering (for example, endless suffering) on anyone contrary to his or her choice.

While often there may be no best action that God can do, it may sometimes be that there is a best kind of action or an equal best kind of action in the senses defined in the previous chapter. In that case for the reason given in the previous chapter God's perfect goodness will require him to do some action of the best or an equal best kind. Often as with actions, so with kinds of actions, there is no best or equal best for God to do. Suppose, for example, God has before him the choice of creating many different species of higher animal. He can create only lions, or only lions and tigers. Plausibly, however many lions he creates, it would be better if he creates, as well as any number of lions, some minimum number of tigers. He could also create only lions, tigers, and pumas. And, plausibly again, it would be better if he creates some minimum number of pumas, as well as some minimum numbers of lions and tigers, rather than any number of lions and tigers. And so, plausibly *ad infinitum*, for plausibly there are an infinite number of possible species.⁵ However many species God makes (with any number of members), it would be better if he made more species rather than more members of the existing species—for there is no best kind of act of creating members of just these or just this number of species. God's perfect goodness has no consequences for how many species of animals he will create.

Is there any way of dividing up actions into kinds of incompatible actions open to God, such that there is a best kind of action? If so, he

⁵ Plato held that there was a world of forms of all 'intelligible living creatures' after which this universe was copied (*Timaeus* 30c): and was understood by Aristotle in consequence to claim that every possible kind of thing existed (*Metaphysics* 990^b). Aristotle reasonably denied this claim: 'it is not necessary that everything that is possible should exist in actuality' (*Metaphysics* 1003^a). For the development of Plato's view, which Lovejoy calls 'The Principle of Plenitude', in subsequent centuries, see A. O. Lovejoy, *The Great Chain of Being* (Harvard University Press, 1936).

will do an action of that kind. Plausibly, it is better for God to bring about the existence of something beyond himself rather than to do any action of a kind incompatible therewith (that is, to refrain from bringing about anything). So the kind of action of causing the existence of something else is a best kind. God must bring about the existence of other things. This is affirmed by a principle that Aquinas often invokes, and sometimes attributes to Dionysius, that 'Goodness is by its very nature diffusive of itself and (thereby) of being'. Norman Kretzmann⁶ has spelled out and justified this principle, understanding it as the principle that a good being will inevitably try to make other good things; and so a good God to whose power there is no limit will inevitably go on making more good things. Hence God must inevitably bring about the existence of things apart from himself, a consequence from which Aquinas backs away when it becomes explicit, in view of his wish to defend the normal Christian view, which is that God did not have to create anything apart from himself.

Can we say anything further about what God must bring about? That depends on whether there is any way of dividing substances into a finite number of types that, unlike the animal species I considered earlier, are such that creating members of certain of those types beyond a certain level of goodness (for example, at least a certain number of members, or members of a certain degree of goodness) is better than or equally good as creating any number of members of types with any level of goodness of a set that does not include all those types. I think that such a division can be made.

Substances are either inanimate or animate (that is, sometimes conscious). Animate substances differ according to whether they can have only the most primitive sensations, or whether they can have desires, beliefs, thoughts, and intentions of various degrees of sophistication. While having any intention at all to do some action involves regarding that action as good to do, animals may not be able to make many true judgements of comparative value. They may not be able to compare the values of different actions, or have any conception of the true value of helping members of other species. So it is a further characteristic for substances to have significant true

⁶ See Norman Kretzmann, *The Metaphysics of Theism* (Clarendon Press, 1997), esp. 223–5, for this justification, and for references to Aquinas's use of the principle.

moral beliefs, which I shall call moral awareness. And they may also have libertarian free will of different degrees of freedom. Free will is a matter of degree; agents can be totally immune from non-rational influences, as is God, or subject to influences (desires) of different strengths deterring them from the pursuit of the good. Finally, they may have different degrees of power and knowledge. Persons, as I defined them in Chapter 1 n. 16, have at least moderate degrees of all the qualities that I have just been setting out (with the possible exception of free will). Now this range of possible substances will include substances of four important types—inanimate; animate without moral awareness or free will; animate with moral awareness and limited free will, power, and knowledge; and animate with unlimited knowledge, power, and freedom. The animate substances without moral awareness and free will are the animals. Those with moral awareness and limited free will, power, and knowledge are what I shall call ‘humanly free agents’. I call them this because, as I shall argue in due course, probably human beings are substances of this kind. A substance with unlimited knowledge, power, and freedom is a divine substance. A conscious life is a good thing. Animate substances are substances of a better type than inanimate ones. It is a good thing that it should be up to the individual substance to make a free choice in the light of moral beliefs about the worth of different actions. So humanly free agents are substances of a better type than animals. But, since the freedom of humanly free agents is limited and they are subject to irrational influences, they may choose to do what is bad. Best of all is a divine being, not subject to this limitation and able to control all things for good. Ranking the types of substance in the order inanimate, animal, humanly free, and divine, we may say that any substance of a later type is better than any substance of an earlier type.

God must bring about something. Can he bring about other divine beings? I discuss this issue in Additional Note 1. In that case the inevitability of God bringing about something could be satisfied, on the Christian doctrine of the Trinity, by a first divine being (‘The Father’) bringing about ‘from all eternity’—which I understand as ‘at each moment of everlasting time’—the Son and the Spirit. In that case there would be no need for Aquinas to reject the Dionysian principle. The need for continuing creation could also be satisfied by the divine beings continuously keeping each other in

existence.⁷ And, if God can make other divine beings, he must surely do so. A solitary God would be a bad state of affairs. God needs to share, to interact, to love, and he can do so most fully with equals. But, if he cannot bring about other divine beings, he must create more limited conscious beings with whom to interact in love—maybe semi-divine beings of limited power but perfect freedom and knowledge or maybe just humanly free agents. But, if God can create as many semi-divine beings as he chooses, must he also create humanly free agents?⁸

Agents who have moral awareness with limited power and freedom, will in virtue of their limited freedom be subject to non-rational influences, temptations to do other than the good. Hence they will have significant free choice in the sense of a free choice that can make real differences to things for good or ill. The goodness of significant free choice is, I hope, evident. We think it a good gift to our own children that they choose their own path in life for good or ill, and influence the kinds of persons (with what kinds of character and powers) they and others are to be. We do this in the belief, or at any rate the hope, that nothing else will cause them to make their choices; that they will make them to some extent independently of the influences upon them. We want them to choose freely (in this sense) to show us love. Humanly free agents, I shall claim in the next chapter, are creatures capable of loving God, and so able to choose freely to show him love and so to form their character that they come to love him naturally. That there should be such creatures is a very good thing. But it involves them being free also to reject God. And, good though the free will to choose between good and evil is, there is

⁷ The Father bringing about the Son and Spirit is not normally called 'creating' them. Christian theologians have usually understood the word 'create' to designate the freely chosen bringing about the existence of a finite thing, not out of previously existing matter or other stuff. See Additional Note 1 (and the proper discussion of this matter in my book *The Christian God* (Clarendon Press, 1994), ch. 8) for argument that God the Father can bring about only two additional divine substances—that is, divine persons.

⁸ If God is not triune, and so must create finite conscious beings other than himself, they would not need to have the free will to reject him. Angels created as essentially perfectly good from the moment of their first existence would satisfy the need to have finite conscious beings with whom to interact. So there would be no inevitability in God creating humanly free agents. If there are angels as traditionally depicted, they are not—at any rate any more—humanly free agents. For their characters are fixed for good or ill (on the traditional view as the result of one initial choice at the first moment of their creation).

the risk that those who have it will make bad choices, form bad characters for themselves, hurt others, and influence their characters for evil. For this reason I suggest that it would not be a good action to create beings with unlimited power to use this limited freedom to form bad characters for themselves, hurt others, and influence their characters for evil. For they might choose to exercise their power to cause evil to the full; and, if they did, the resulting state of affairs would be so terrible that God could not take the risk that it might occur. And, as we shall see in Chapters 10 and 11, certain bad states not caused by humans are necessary if humans are to have deeply significant choices at all. As our creator and the source of so much good for us, God has, I shall argue, the right to require from us some suffering to make possible good for others, and further good for ourselves. But there is surely a maximum to the amount of suffering that a good God would allow anyone to endure (against their will) for these reasons. If God creates beings with the freedom to choose between good and evil, they must be finite, limited creatures. Even so, as we well know, much evil is to be expected if creatures are given this dangerous choice; and so I suggest that God would not inevitably bring about such creatures. Still, there is great value in a world containing such creatures. For it would contain a very special kind of goodness. God would be creating a kind of freedom of choice that he does not possess himself—God can do no evil. Hence I am inclined to suggest that it would not be worse (or better) for God to create some humanly free agents (with some minimum number and degree of their freedom, power, etc.) than to create instead any number of other divine or semi-divine beings, animals, or inanimate things (and any degree thereof). To create some humanly free agents and not to create some humanly free agents (in addition to whatever else he brings about) would be acts of equal best kinds.

God has reason also to create animals, beings simpler than humanly free agents, ones that spontaneously do good without having the free will to choose between good and evil; conscious beings who want (that is, have desires) to have various sensations and do various actions; and so get pleasure or enjoyment from having their wants satisfied. It would be good that they should learn what is to their benefit or harm, and use their knowledge and capacities spontaneously (not through free choice) to care for themselves and to prolong their life; and to care for each other and especially their young. No doubt snakes and fish get pleasurable sensations from

food and sex. And birds and rabbits rejoice in controlling their bodies to fly and run. They learn where food is to be had, and danger avoided, and through effort often get the food and avoid the danger. As we move up the evolutionary scale, we find animals whose actions are less a matter of instinct, and more a matter of learning and so of knowledge. It is good that there should be beings who make a difference to the world through their knowledge of how it is and their power over it, and that they should seek knowledge (for example, of where food is to be had) in order to know how to make a difference. But the kind of goodness possessed by animals is simply goodness of a lower degree than the kind of goodness possessed by humanly free agents and indeed God himself. They are conscious, like God; and have good desires, like God; but of course lack so much that God has. Unlike humanly free agents, they do not have a kind of goodness that God lacks; but they do not (like humanly free agents) suffer from the possibility of freely doing what is wrong. Nevertheless clearly it is good that they exist. God has good reason to create animals. I shall argue in Chapter 11 that for the higher animals (cats and monkeys, as opposed to worms and insects), who alone can do significant actions, some suffering is necessary. And so with animals, as with humanly free agents, there must be a limit to the amount of suffering that God is justified in allowing to occur for such good purposes; and a limit lower than the limit for humanly free agents whose life in general has so much more good in it.

And finally, of course, God has reason to create a beautiful inanimate world—that is, a beautiful physical universe. Whatever God creates will be a good product; and so any physical universe that he creates will be beautiful, as are humans and animals. Consider the stars and planets moving in orderly ways, and plants growing from seed into colourful flowers and reproducing themselves. Even if no one apart from God himself sees such a world, it is good that it exists. And maybe the kind of beauty that it exemplifies is different from that of finite beings, or of God himself. But I do not find it obvious that it exemplifies in respect of its beauty a kind of goodness not possessed by God himself or possessable by creatures of other kinds, or that it possesses any other kind of goodness peculiar to the inanimate. I shall, however, be arguing shortly that humanly free agents (and animals) need to have bodies and to live in a wider physical universe, if they are to have the properties for which they

are valuable to any significant degree, and so, if God creates them, he will create a beautiful physical universe.

If God does create a type of substance that has its own special value, although there may be no best number of such substances or best degree in which they can instantiate what is valuable in them, there must be a high probability that he will instantiate significant numbers or degrees or extent of them—since there are so many other possible numbers, extent, and degrees. In general it will always be a better act to create more—more humanly free agents, and animals, and a larger physical universe (or more physical universes). But, since there is no maximum to the more, God's perfect goodness does not require him to create any particular amount of such beings. But in one respect there is a maximum. As I claimed earlier, there must be a maximum to the degree of power that God gives to humanly free agents because of the evil they may do; and a maximum to the amount of evil produced by any natural processes that are required in order to make possible the exercise of significant free will by such agents.

I conclude that, while any act of creating that did not include humanly free agents would leave the world without a certain kind of goodness, it is no way obvious that not creating animals or (if there are no humanly free agents or animals) not creating a physical universe would leave the world without a certain kind of goodness. For the kind of goodness that animals or the inanimate have is possessed by God himself or could be possessed by other divine or semi-divine beings. Hence it will not be a best or equal best kind of act to create animals, or a physical universe without animate beings, since it would always be better to create more substances of other types.⁹ (God could always more than compensate for the absence of animals by creating more angels.) But that does not mean that it is not a good act to create animals, since, for any act of creating other things, the act of creating those same things plus the act of creating animals as well will be a better act. However, with animals there will always also be some better act of creating a world without them. Not so for humanly free agents.

⁹ It does not, of course, follow that he will create more substances of other types; since, for any number of substances of any such type that he creates, it would always be better if he created more.

The probability that, if there is a God, he will do the best action or an action of the best kind (when there is one) is 1. The probability that he will do a bad action is 0. If there is before God a choice of n equal best incompatible actions (or kinds of action), the probability that he will do any particular one of these is $\frac{1}{n}$. Hence the probability that, if there is a God, there will be rational beings other than a single divine person is 1. The probability that there will exist humanly free agents (and so a physical universe) is $\frac{1}{2}$, subject to a maximum of power, provided by the maximum of the amount of suffering that God would be justified in allowing those agents or others to suffer in consequence. And, while the goodness of animals leads us to expect that God might create them, I do not think—for the reason given above—that we can associate quite such a high probability to this. However, the moral intuitions on these matters that I am commending to my readers must inevitably be somewhat tentative, and, even if they are correct, very precise numerical values may not capture the resulting probabilities about what God will create. My arguments do not depend on giving very precise values to the probabilities involved. All I am claiming is that it is fairly probable that God will create humanly free agents (of power up to a certain limit, which will be considered further in Chapter 11), and so (for reasons yet to be given) a beautiful physical universe; and not very improbable that he will create animals also.

Humanly Free Agents Need Bodies

I defined ‘humanly free agents’ as animate beings with moral awareness and limited freedom, power, and knowledge. They are persons of a limited kind. I claimed also that God would be unlikely to instantiate the kind of goodness that they possess in a minimum degree. If he makes humanly free agents, he will give them a significant amount of freedom, power, and knowledge. If their limited freedom is to be greatly valuable, it will be a freedom to choose between good and evil in the exercise of power to make deeply significant differences to themselves, each other, and the physical world by their choices (including the power to increase their powers and freedom of choice). They need to be able to cause in themselves and others pleasant and unpleasant sensations; investigate the world and acquire knowledge, and tell others about it. But significant

responsibility involves also a power for longer-term influence over those powers themselves. They must be able through choice to influence the powers of themselves and others to acquire these beliefs and cause sensations, and to influence what they find pleasant or unpleasant; and to influence the ways (for good and evil) in which they are naturally inclined to use their powers. They must be able to help each other to grow in knowledge, factual and moral; in their power to influence things; and in the desire to use their powers and knowledge for good. And they must also, in order to have significant responsibility, be able either through negligence or through deliberate choice to restrict their own and each other's knowledge, powers, and desire for good.

So these creatures must start life with (or acquire by natural processes) limited unchosen power and knowledge and desires for good and bad, and the choice of whether to extend that power and knowledge and improve those desires, or not to bother. And, if that choice is to be a serious one, it must involve some difficulty—time, effort, and no guarantee of success must be involved in the search for new knowledge and power, and improved desires. So creatures need to have at their disposal an initial range of basic actions. We may call the kinds of effects that a creature can (at some time) intentionally bring about by her basic actions her region of basic control. Creatures need an initial region of basic control. And, knowledge being a great good and necessary for control, creatures need also an initial range within which they can acquire largely true beliefs about what is the case. Let us call the kinds of such beliefs that a creature can acquire his region of basic perception. The region of basic control will have to lie within the region of basic perception. For creatures cannot bring about effects intentionally unless they know which effects they are bringing about. For us humans, certain states of our bodies are our region of basic control, and our region of basic perception consists of the observable states of things in a wide region encompassing our bodies.

Extending our region of control beyond the basic region will involve discovering (that is, acquiring true beliefs about) which basic actions have further effects. For the possibility of a large extension of our region of control, it needs to be the case that our basic actions have different effects beyond the basic region that vary with the circumstances in which they are done. What these circumstances are must themselves be alterable by our basic actions; and, if

we are to affect the region of control of others, we must be able to alter the circumstances in which those others are to be found. Effects 'beyond' the basic region means in some sense effects more 'distant' than it; and altering 'the circumstances' involves in some sense 'movement'. We can learn what effects we have when we change circumstances if our region of basic perception moves with our region of basic control—though that may not always be necessary if the former region is much larger than the latter region. We can learn how to produce some effect in another room, by moving into the room, and, when we are there (but not here), we can see the effects of our actions there—our region of basic perception has moved with our region of basic control. By going to see where our bullet lands when we fire our gun at different angles, we can learn the distant effects of firing a gun at different angles and, in this way again, extend our region of control. But we can learn how to hit some distant person with a stone without altering our region of basic perception, for it is large enough for us to discover without moving the effects of throwing stones in different ways. We can affect the region of control of others by moving those others by our basic actions—we can take someone and put him or her in prison or on a train. The region of our perception may be increased by discovering (through previous movement) which basic perceptions are normally evidence of more distant phenomena. We can learn to see things far away through a telescope where we have discovered (through going to see) the correlation of things a little way away with their images in the telescope, and extrapolating from that to a similar connection between their images in the telescope and things at a great distance. Control may be widened so as to include events well in the future; and perception may be widened so as to include events well in the past.

So, in order to have significant freedom and responsibility, creatures need to be situated in a 'space' in which there is a region of basic control and perception and a wider region into which they can extend their perception and control by learning which of their basic actions have which more distant effects when they are stationary, and which of their basic actions cause movement into which part of the wider region; and which of their basic perceptions are caused by which more remote events. If they are to be able to perform mediated actions—that is, their basic actions are intentionally to have distant effects (including which ones move them

into parts of a wider region), and distant events are to have basically perceptible effects—the spatial world (I shall be arguing more fully in Chapter 8) must be governed by laws of nature. And, if creatures are not simply to have true beliefs about what these effects are, but to learn by rational inference from observation and to have a choice of whether to acquire such beliefs by rational inquiry, those laws of nature must include laws governing properties that they can observe, and the laws must be sufficiently simple for creatures to understand.

Further, if creatures are not merely to find themselves with beliefs about each other's beliefs and purposes (which they will need to do if they are to be able to influence them), but to be able to choose to learn about each other's beliefs and purposes and to communicate with them in the public way needed for cooperative action and cooperative rational discussion (which will involve language), they need to manifest their beliefs and purposes in a public way—that is, through their regions of basic control, which must therefore be physical regions. These regions need to behave in such a way that the simplest explanation of that behaviour is in terms of the beliefs and purposes (that is intentions) of creatures whose regions they are. In consequence, for example, we must be able to attribute to each other (on grounds of being the simplest explanation of the behaviour of others) beliefs sensitive to input—for example, to attribute to someone a belief that some object is present when light comes from some object on to their eyes; and purposes that—although not fully determined by brain states—do show some constancy. We can, for example, come to understand the language of another human only if we assume that he normally seeks to tell the truth and has some language constant over time by which he expresses his beliefs, and that his beliefs are often sensitive to incoming stimuli in ways similar to our own. Thus we may notice that often when it is raining and his eyes are pointed in the direction of outdoors, or he is outdoors and raindrops fall on his face, he says 'google'. If we suppose his beliefs sensitive to stimuli in the same ways as our own, and that he has the purpose of truth-telling, it is a simple hypothesis explaining his speech behaviour to suppose that he means by 'google' that it is raining. Or, at any rate, it is a simple initial hypothesis that would need to be combined with many other hypotheses about what he means by other words to form a simple overall theory of his linguistic behaviour, which explained it in terms of a system of beliefs

and purposes that explained all his behaviour (linguistic and non-linguistic).

It would be good that creatures should have the power not merely to extend their regions of control and perception beyond the basic, but that they should have the power to extend or restrict (or prevent being restricted by others or by natural processes) the regions of basic perception and control (including the ability to move) of themselves and others, and to extend or restrict the range of pleasant or unpleasant sensations and the desires to do this or that which they have. There need to be basic actions that creatures can do, or non-basic actions that they can learn to do, that under various circumstances will make differences to their capacities for basic action and perception, and to their sensations and desires. That involves there being natural processes that they can discover and so affect, that enable them to perform their basic actions and acquire and retain in memory basic perceptions, and diminish or increase pain or pleasure. And, if these processes are to be manipulable not merely by the human whose they are, but by other humans as well, they must be public processes.

For this latter to be possible there needs to be a public place at which we or others can act to interfere with and so improve or damage the quality of our sensations and desires, and the extent of our capacity for basic actions and perceptions. We need not merely a region of basic control and a region of basic perception, but what I shall call a machine room. This is a public place where our intentions are translated into basic actions, and incoming stimuli are translated into sensations and beliefs, and processes give rise to desires and thoughts. When there is such a physical object, we and others can damage or improve these processes. In all these ways we need a controllable public region where we are. And so the existence of humanly free agents with significant freedom requires the existence of a physical universe.

All this means that, in order to have the kind of limited freedom and responsibility that I have analysed, creatures need bodies, in the sense in which this notion was analysed in Chapter 2. They need to perform basic actions through some public object. Their abilities, sensations, and beliefs need to depend on the functioning of a public object. So too do their basic perceptions, which will be caused by events outside that object. And, in order to have a region of perception wider than a region of basic perception (that is, in order to be

able, by learning which events beyond the region of basic perception are connected with which events within that region, to extend their perceptual knowledge) there will need to be a centre—the region of basic perception, from where they look out on the world.

In human bodies all these regions fit together. The body of a human consists of just one spatially extended and connected public object, many of whose states form that human's region of basic control, and many of whose parts form that human's machine room. Many states of the machine room are caused by events outside the region of basic control, and thus cause the human to have a region of basic perception. For us our region of basic control is what we can do with our limbs, mouths, and tongues—just like that, not by doing anything else. The region of basic control varies with the age of a human; and increases and then decreases again with time largely unaided by other humans—how fast we can move our arms and legs does not depend too much on learning or help from others. But we can discover or be taught how to increase that region of control in many respects—above all, how to influence others by uttering sentences of a language. And we have a range of basic perception, increasing or decreasing with the age of a human independently of any intentional action. Recognizing inanimate objects of many kinds is a perceptual capacity that develops without much help; learning to understand people's words needs more by way of help from others. We learn by our basic actions to hurt or benefit others, to use tools, build houses, or cut down trees. We utilize principles of what is evidence for what to detect the previous presence of others from footprints and remains of fires, and the passage of elementary particles from tracks in cloud chambers. Through our growth of knowledge and control, we learn how to cause pleasure and pain, to give knowledge and control to others or to refuse to do so. We can allow ourselves to get into situations where it is difficult to do good, and so fall into bad habits and naturally develop bad desires—or, alternatively, prevent this happening. And, through learning, we can acquire the ability to influence the ways in which others desire to use their powers—we can educate them morally or immorally.

But, as well as learning how to extend the region of control and perception (of ourselves and others) beyond the basic, we can also learn how to extend or restrict (or prevent being restricted) the region of basic control and perception itself. By starving ourselves or others, we can restrict basic powers and perceptual abilities; as we

can by cutting off arms or tongues or eyes. In ways unintended by ourselves or others, our powers may diminish through disease, but we can learn to prevent the effects of disease by medicine and surgery; or we can not bother to take the trouble to discover how to do so. And our present abilities to affect our basic powers and perceptual abilities seem fairly small compared with what medical science will surely provide for us in the course of the next thousand years. Medical intervention will surely enable us even within the next hundred years to grow new limbs and sense organs, and to slow down memory decay. Our bodies being spatially extended bodies thus involves the surface of mind–brain interaction lying within the body (that is, in the brain); and it also involves events elsewhere within the body affecting what we basically perceive and how we can act basically.

There could be kinds of body very different from human or animal bodies. The body of some rational creature in another world might consist of two or more parts not spatially connected. One part might be his machine room, and another part that chunk of matter whose states form his region of basic control. To improve the latter, you would need to tamper with the former, which might be a mile or two away. Or both parts might be ones that the creature could control, and that were also parts of his machine room. Our bodies do not merely have a spatial location, but are spatially extended. It would also be possible for creatures' basic capacities for perception and action to depend on public processes if they had a 'particle-body' as their machine room. The public processes would then need to consist of temporally extended input to a spatially unextended object. The latter would be like a totally impenetrable black box. Creatures could discover how to improve or damage their sight, or weaken or strengthen their memories by giving a certain input to the box over a long period of time. But the box could not be opened; indeed, it would have no spatial extension. Their memories would not then depend on a brain; they would depend on input over time, which affected memory by action at a temporal distance. This would provide a less immediate kind of embodiment—for the dependence of the mental on the physical would not be instantaneous. But it might seem, nevertheless, to be an alternative way in which creatures could have the ability to affect their and each other's basic capacities without the normal kind of embodiment.

In one or other of these ways, if they are to have the great good of being able to affect themselves and each other greatly for good or ill, finite creatures need bodies. Angels traditionally are finite creatures, but we cannot blind them or embrace them, because there is no place at which to direct our activity; we cannot capture them in order to affect them. Mere telepathic communication with individual spirits does not allow for public discussion with many humans and spirits. And, if you begin to add to these situations features that do make for the ability to capture angels or have public discussion with spirits, you will find that you are beginning to give them bodies in my sense. If God is to create creatures with limited free choice to make deeply significant differences to themselves, each other, and the physical world for good or ill, he must make them embodied. Humanly free agents need bodies, and thus to be placed in a physical universe that God has made. If humans have free will (an issue that I will discuss briefly in due course), humans are thus evidently one kind of humanly free agents.

Since humanly free agents have desires and moral awareness, they will be capable of love and gratitude; and, since they are capable of significant growth in knowledge, they will be able to develop the metaphysical concepts that allow them to have the concept of God and so to love him if they come to believe that he exists. If God does create such creatures capable of loving him, he must make himself known to them, and help them to deal with their problems—especially if they ask him, and they have got themselves into difficulties too great for them to sort out for themselves. Or at least he must do this, except in so far as making himself known and helping them makes it more difficult for them to make significant free moral choices themselves. I shall argue in Chapter 11 that there are such limits to the goodness of God making himself known to humans and helping us. But, barring these limits, my view is that God has an obligation to make himself known. (Fathers who fail to take evident interest in their children are bad fathers.) Hence there is a probability of 1 that he will do so. But again nothing turns on the exact value. So let us just say that there is a high probability that a God will make himself known to creatures capable of loving him, except in so far as that will curtail their freedom. Humans are capable of loving God, for they can have the concept of God and the ability to love. But, in view of the disadvantages of his revealing himself to humans, I shall argue that there is also

a high probability that his self-manifestation to them will be limited.

For the same reasons as humanly free agents, animals need bodies in order to have a location in space where they can influence each other, learn from each other, and be affected by each other.

Conclusion to the Chapter

I have argued in this chapter that there is a modest probability intermediate between 1 and 0, to which I will give the artificially precise value of $\frac{1}{2}$, that a God will create humanly free agents located in a beautiful physical universe, containing perhaps also animals. Subsequently I will discuss various particular features of the universe that are necessary for the existence of humanly free agents, and show that it is most improbable that they would occur unless God brings them about. Further chapters will discuss further features with which God might be expected to endow those creatures and their universe, and which also there is no reason to expect if God does not bring them about. The considerations of this chapter and the two previous chapters now enable me to meet an important objection to probabilistic arguments, which is put in a fairly precise form in an article by D. H. Mellor.¹⁰ Mellor's argument is directed against the argument from design, but it could be directed equally well against almost any argument for the existence of God. I will give Mellor's argument, which is only one of a number of interconnected arguments in his paper, a somewhat more precise form than it has there. I hope that I shall not represent him unfairly.

Mellor imagines that we find the cards in a bridge hand that has been just dealt in a certain kind of order (e), and he considers the worth of the argument from this order to a cheat's having put them in order (h) (that is, to the shuffling having been rigged). He rightly says that we can assess the probability of this hypothesis only if we could have written down in advance of looking at the pack in what order (or orders) a cheat would be likely to have arranged the cards. In our terminology, in order to assess $P(h|e \& k)$, we need to know, in advance of observing e , $P(e_i|h \& k)$ and $P(e_i|k)$ for different e_i . This part of the argument seems to me

¹⁰ 'God and Probability', *Religious Studies*, 5 (1969), 223–34.

indubitably correct. Mellor then implies that we cannot do this where h is the hypothesis of theism, and e_i are different possible worlds, since we have no idea what are the intentions of God, if he exists. The answers to Mellor's objection should now be clear from what has been said earlier in this chapter and in Chapters 4 and 5. First, contrary to Mellor, we do have some idea of what kinds of world God is likely to create and hence an idea of how $P(e_i|h \& k)$ will differ for different e_i ; and, secondly, unlike in the bridge analogy, not all e_i are equally likely *a priori*. Some e_i , some worlds, have such complexity that *a priori* they are not be expected—probably only the power and choice of God can bring them about. So $P(e_i|h \& k)$ may be known to exceed $P(e_i|k)$ not only because of the known character of God (because he is known to be more likely to bring about some e_i than others) but also because of the known *a priori* improbability of certain world states occurring uncaused.

The Cosmological Argument

The previous chapters have been concerned with elucidating the general principles for assessing the worth of arguments from experienced phenomena to the existence of God. With this chapter I begin to apply them to the consideration of particular arguments, starting with the argument that has the most general premiss of all—the cosmological argument.

The Nature of Cosmological Arguments

Kant defined a cosmological argument as one that starts from ‘experience that is purely indeterminate’ or ‘experience of existence in general’. Let us say, more precisely, that it is one that starts from the existence of a finite object—that is, an object of limited power, knowledge, or freedom—that is, any object other than God. However, other arguments called cosmological have in effect started from something rather more specific, the existence of a complex physical universe; and I shall confine my discussion mainly to these. I understand by a physical universe a physical object consisting of physical objects spatially related to each other and to no other physical object. (By ‘spatially related to each other’, I understand ‘at some distance in some direction from each other’.) Our physical universe, *the* universe, is the physical object that consists of all physical objects including the earth, things on them, and gases between them. The universe is the only physical universe of which we have certain knowledge, but I define it in such a way as not to rule out the logical possibility of other physical universes,¹ or of objects

¹ Another physical universe would be a physical object consisting of physical objects, spatially related to each other, but not to the objects of our universe such as the earth.

that are not part of any physical universe (for example, God or some finite spirit, neither of which is a physical object). By a complex physical universe I understand one consisting of many physical objects of diverse volume, shape, mass, etc. On the assumption that our physical universe is the only one, I phrase the cosmological argument as an argument from that universe; but, if there is more than one universe, the argument should be treated as an argument from all the physical universes there are (the 'multiverse', as it is sometimes called, which, in view of the complexity of our universe, will itself evidently be complex). For the purposes of this chapter, nothing depends on how many physical universes there are; but we shall need to take the possibility of there being more than one universe a bit more seriously in the next chapter.

From time to time various writers² have told us that we cannot reach any conclusions about the origin or development of the universe, since it is the only one of which we have knowledge, and rational inquiry can reach conclusions only about objects that belong to kinds, for example, it can reach a conclusion about what will happen to this bit of iron only because there are other bits of iron, the behaviour of which can be studied. This objection has the surprising, and to most of these writers unwelcome, consequence, that physical cosmology could not reach justified conclusions about such matters as the size, age, rate of expansion, and density of the universe as a whole (because it is the only one of which we have knowledge); and also that physical anthropology could not reach conclusions about the origin and development of the human race (because, as far as our knowledge goes, it is the only one of its kind). The implausibility of these consequences leads us to doubt the original objection, which is indeed totally misguided.

Uniqueness is relative to description. Every physical object is unique under some description, if you allow descriptions that locate an object by its spatial position—that is, by its distance and direction from named objects. Thus my desk is the one and only desk in such and such an apartment; and that apartment is the penultimate one on the left in a certain row. And, even if you allow only descriptions

² Including, for example, Hume. In Hume's *Dialogues concerning Natural Religion* (first published 1779, ed. H. D. Aiken (Hafner, 1948), 23), Philo objects to arguments to the cause of the universe as an object that is 'single, individual, without parallel or specific resemblance'. See also Hume's *Enquiry concerning Human Understanding* (first published 1748, ed. L. A. Selby-Bigge, 2nd edn. (Clarendon Press, 1902)), 115.

in qualitative terms—for example, the one and only existing desk of such-and-such a shape, such-and-such a weight, with such-and-such carvings on its legs, and scratches on its top situated in an apartment that is the penultimate one in a row of apartments—it is still plausible to suppose that most physical objects have a unique description.³ In the first respect, the universe is, like all physical objects, pickable out by a unique description—‘the physical object consisting of all physical objects including the Earth spatially related to each other and to no other physical object’. In the second respect, too, the universe may well be describable by a unique description—for example, ‘the physical object consisting of physical objects that are all spatially related to each other and to no other physical object, governed by laws of nature *L* beginning from initial conditions *I*’ (where *L* and *I* are specified in detail). In all this the universe is no more ‘unique’ than the objects that it contains. Yet all objects within the universe are characterized by certain properties, which are common to more than one object. My desk has in common with various other objects that it is a desk; and with various different objects, that it weighs less than a ton, and so on. The same applies to the universe itself. It is, for example, like objects within it such as the solar system, a system of material bodies distributed in empty space. It is a physical object and, like other physical objects, has density and mass. The objection fails to make any crucial distinction between the universe and other objects; and so it fails in its attempt to prevent at the outset a rational inquiry into the issue of whether the universe has some origin outside itself.

So then, to return to the main thread, a cosmological argument is an argument to the existence of God from the existence of some finite object or, more specifically, a complex physical universe. There have been many versions of the cosmological argument given over the past two-and-a-half millennia; the most quoted are the second and third of Aquinas’s five ways to show the existence of God.⁴ However,

³ The claim that this is necessarily so for all objects is one version of the principle of the identity of indiscernibles. I do not rely on this principle, but only on the plausible empirical claim made in the text.

⁴ See St Thomas Aquinas, *Summa Theologiae*, Ia2.3. Aquinas’s first way is sometimes said to be a version of the cosmological argument, but it does not count as one on my definition of a cosmological argument, since it argues not from the existence of physical objects, but from change in them. It claims in effect that, given that there are physical objects, change in them is so surprising that we need to invoke God as its source. I cannot see that change in them is so surprising that we need to invoke God as its source. Given the

Aquinas's 'five ways', or rather the first four of his five ways, seem to me to be one of his least successful pieces of philosophy.⁵ In my view the two most persuasive and interesting versions of the cosmological argument are that given by Leibniz in his paper 'On the Ultimate Origination of Things', and that given by his contemporary Samuel Clarke in his Boyle Lectures for 1704 and published under the title *A Demonstration of the Being and Attributes to God*.⁶ The former seems to be the argument criticized by Kant in the *Critique of Pure Reason* and the latter the argument criticized by Hume in the *Dialogues*. In so far as I consider one detailed example of a cosmological argument, I shall consider Leibniz's version, but most of my remarks will apply to most versions of the argument.

The starting points of cosmological arguments are evident facets of experience. There is no doubt about the truth of statements that report that they hold. It seems to me equally evident that no argument from any of such starting points to the existence of God is deductively valid. For, if an argument from, for example, the existence of a complex physical universe to the existence of God were deductively valid, then it would be incoherent to assert that a complex physical universe exists and God does not exist. There would be a hidden contradiction buried in such co-assertions. Now, the only way to prove a proposition to be incoherent is to deduce from it an *obviously* incoherent proposition (for example, a self-contradictory proposition),⁷ but, notoriously, attempts to derive obviously incoherent propositions from such co-assertions have failed through the commission of some elementary logical error. Furthermore, it seems easy enough to spell out in an obviously coherent way one way in

existence of physical objects, it seems to me no more surprising that they should change than that they should remain changeless. Aquinas's supposition to the contrary arises from the Aristotelian physics that is so closely meshed with his philosophy. It is more plausible to suppose that the existence of *orderly* change is surprising, but the argument from orderly change is Aquinas's fifth way and is a teleological argument that I shall discuss in the next chapter.

⁵ For detailed criticism of Aquinas's five ways, see the full and careful discussion in A. Kenny, *The Five Ways* (Schocken Books, 1969).

⁶ Clarke's argument, treated as a deductive argument, has received very full and interesting treatment in W. L. Rowe, *The Cosmological Argument* (Princeton University Press, 1975).

⁷ I argue for this claim about how incoherence is to be proved in *The Coherence of Theism* (Clarendon Press, 1993), ch. 3. I argue also that the main way to prove the coherence of a claim is to spell out in an obviously coherent way one way in which it could be true—that is, one obviously coherent conjunction of propositions that entail the claim.

which such a co-assertion would be true. There would be a complex physical universe and no God, if there had always been matter rearranging itself in various combinations, and the only persons had been embodied persons; if there never was a person who knew everything, or could do everything, etc. Atheism does seem to be a supposition consistent with the existence of a complex physical universe, such as our universe. Of course things may not be as they seem, but, in the absence of any worthwhile argument to the contrary known to me, I shall assume that the non-existence of God is logically compatible with the existence of the universe, and so that the cosmological argument is not a valid, and so not a good, deductive argument. Our primary concern is however to investigate whether it is a good C-inductive or P-inductive argument, and just how much force it has.

The Scientific Inexplicability of the Universe

Now there could be a universe today for whose existence today there was no scientific explanation at all. But, of course, there is a full scientific explanation of the existence of our universe today in terms of it existing in a certain state yesterday (for example, having approximately the same amount of matter-energy as it does today), and laws of nature (including the law of the conservation of matter-energy) operating on its state yesterday to produce a universe today. I express this explanation for the moment in terms of the amended Hempelian model of causes as prior states of affairs that together with laws explain subsequent events. The state of the universe yesterday clearly also has a full explanation in terms of its state the day before yesterday and the operation of the same laws of nature. And clearly we can go back in time in this way providing full explanations of the state (and so the existence) of the universe for many millions of years. Denoting a state of the universe at time n by S_n , where a larger n indicates an earlier period of time, by s_1 the state of the universe today, and by L the laws of nature, we get the following picture:

$$\begin{array}{cccc} L & L & L & L \\ \dots S_5 \neg > S_4 \neg > S_3 \neg > S_2 \neg > S_1 . \end{array}$$

This picture takes the 'states' as temporally extended states, states that last for a period of time and not merely for an instant.⁸ Between any two instants of time (for example, 2 p.m. and 3 p.m. today) there will be an infinite number of temporal periods (for example, the periods of $\frac{1}{2}$ hour, $\frac{1}{4}$ hour, $\frac{1}{8}$ hour, $\frac{1}{16}$ hour, etc.). The interesting question about whether the universe is of finite age, or of infinite age, is the question about whether there has been a universe only for no more than a finite number of periods of equal length (for example, a finite number of years) or whether it has lasted for an infinite number of such periods.⁹ If the series is finite, and we go backwards, we will reach a first state of the universe (as governed by laws that operate today); if the series is infinite, we go backwards for ever.¹⁰ We do not know whether the universe has a finite or an infinite age, but

⁸ I argue in 'The Beginning of the Universe and of Time' (*Canadian Journal of Philosophy*, 26 (1996), 69–89) that all talk of events happening at instants (e.g. 2 p.m.) is reducible to events happening over periods of time; that instants are merely the boundaries of periods (2 p.m. and 3 p.m. are the two boundaries of the hour between them); and that periods are not made up of instants. An object being brown at 2 p.m. is it being brown for a period that includes 2 p.m.

⁹ Talk of the universe being finitely or infinitely old makes sense only if we have a universe law-governed throughout its past, in which temporal intervals can be measured. See 'The Beginning of the Universe and of Time'. In a universe with a chaotic past, there would be no difference between it being infinitely old and it being finitely old. A universe being finitely old does, I argue there, entail it having a beginning (in the sense of there being empty time before the universe came into existence). But I argue there that it being infinitely old does not entail it having no beginning in this sense. And, if it was at some earlier period chaotic, although it would not then have been of either finite or infinite age, it could still have had a beginning or no beginning. My argument in the present book is an argument for God being the cause of the existence of the universe, whether it is of finite or infinite age. (It could be extended into an argument for God being the cause of a universe that was neither finite or infinite, because it had an earlier chaotic period. But, I argue, we have and can have no knowledge that it was like this.) We can ignore the issue of whether an infinitely old universe has a beginning. If my argument in the cited article is correct, someone who held that it is an item of revelation that the universe has a beginning can rationally do so independently of whether science suggests that it is of finite or infinite age.

¹⁰ In presenting his 'Kalam cosmological argument', William Craig claims (1) that a 'beginningless series of events in time cannot exist'. See e.g. his 'The Kalam Cosmological Argument', in W. L. Craig (ed.), *Philosophy of Religion* (Edinburgh University Press, 2002). His argument for (1) is that it follows from (2), 'an actually infinite number of things cannot exist'. His argument for (2) is that Cantor's system of mathematics for operating with infinite numbers includes the principle of correspondence that two sets (such as the infinite set of integers 1, 2, 3, 4, . . . and the set of even integers 2, 4, 6, 8 . . .) have the same number of members if (as in this example) they can be put into one-one correspondence with each other. This principle has paradoxical consequences such as that in a hotel with an infinite number of rooms, all occupied, you can accommodate a further

science may be able to show us which is the more probable. Science might show that (on the assumption that the same laws operate in the past as operate today) extrapolation backwards from S_1 via L eventually leads to a physically impossible state or a state with no matter, at a time t ; whence we could conclude that probably the universe came into existence at a time after t , and not as a result of the operation of scientific laws. My assessment of the present state of science is that that is what it does tend to show. It suggests the simplest explanation of the current mutual recession of the galaxies (the groups of stars that get further and further apart from each other) is that this is a consequence of fundamental laws operating on matter-energy produced by an enormous explosion, the Big Bang, fifteen billion years ago. As we go backwards in time, matter was more and more dense. But if, as it appears, it would have been a physically impossible state for matter to be packed into a point with infinite density, the matter must have come into existence and the Bang caused its recession when it was packed very densely but not infinitely densely. However, new scientific data or further reflection on existing data might lead scientists to conclude that the best (that is, the simplest) explanation of the laws that operate in today's relatively spaced-out universe is in terms of a more fundamental law that has the consequence that quite different less fundamental laws would have operated in an earlier and denser universe. Extrapolating backwards in accordance with these laws might lead to the conclusion that any Big Bang would have occurred in a very dense state produced by a previous contraction of the

infinite number of guests by moving each existing guest to the room with the number twice that of his present room, and then proceeding to fill the (infinite number of) vacant odd-numbered rooms. But I suggest that we can allow what seems to me the obvious logical possibility of there being an infinite number of things (e.g. stars), without adopting Cantor's mathematics, or this kind of way of applying it. Further (1) above follows from (2) only on the assumption that events that are all now past are in some sense actual. But in that case all the members of the infinite series of periods of unequal length, of $\frac{1}{2}$ hour, $\frac{1}{4}$ hour, $\frac{1}{8}$ hour, etc., which have already occurred during the past hour, are also now actual, which by Craig's (2) is not possible. So either (2) is false, or (1) does not follow. Craig's Kalam argument to God requires not merely (1) but also 'whatever begins to exist has a cause'. But it seems to me, for reasons given in this chapter, this, like 'the universe began to exist', can be given only an inductive justification. Kant also affirmed (1)—see his *Critique of Pure Reason*, B454, trans. N. Kemp-Smith (Macmillan, 1964). His argument is that an infinite series cannot have a last member; and that a beginningless and so infinite series terminating in a present event would have a last member. But Kant's claim that an infinite series cannot have a last member holds only of infinite series with a first member—which a beginningless series would not have.

universe. But we can have no evidence of the operation of quite different laws in the past, unless their operation is a consequence of the simplest explanation of what is happening in the present. In so far as science shows that the fundamental laws¹¹ of nature operating today are *L*, and that extrapolating *L* backwards leads to a physically impossible state, we have to conclude that there was a beginning to the universe-governed-by-today's-laws and that we can have no knowledge of anything earlier than that. There might have been a physical universe governed by quite different laws, or there might have been no universe at all. But it is always simpler to postulate nothing rather than something; and so, in the absence of observable data made probable by the hypothesis that quite different non-fundamental laws were operating in the past, the hypothesis that the universe came into existence a finite time ago will remain the more probable hypothesis. But it is certainly possible that science might come to show that the fundamental laws governing our universe are such that we can extrapolate backwards for ever from the present state of the universe. Then evidence would support the claim that the universe is infinitely old.

If we confine ourselves to scientific explanation, it will now follow that the existence of the universe (for as long as it has existed, whether a finite or an infinite time) has no explanation. It has often been assumed and sometimes argued by philosophers, including Hume, that, if we have a scientific explanation of each of a collection of states, then we have an explanation of the whole collection. Thus, Hume:

In . . . a chain . . . or succession of objects, each part is caused by that part which preceded it, and causes that which succeeded it. Where then is the difficulty? But the *whole*, you say, wants a cause. I answer that the uniting of several parts into a whole, like the uniting of several distinct countries into a kingdom, or several distinct members into one body, is performed merely by an arbitrary act of the mind, and has no influence on the nature of things. Did I show you the particular causes of each individual in a collection of

¹¹ See p. 158 n. 3 for the difficulty that there may be no most fundamental laws of nature, that every law may be explicable by the operation of a yet wider law (Galileo's law by Newton's laws, Newton's laws by Einstein's laws, and there be no end to this process). I shall argue there that, if this were so, the situation will be equivalent to one in which there is an infinite series of equally fundamental laws, each applicable to regions of the universe in different physical conditions. In that case, *L* in the text of this chapter should be taken as such an infinite series.

twenty particles of matter, I should think it very unreasonable, should you afterwards ask me what was the cause of the whole twenty. This is sufficiently explained in explaining the cause of the parts.¹²

To assess the worth of Hume's claim we need to develop general principles concerning the relation of causes of parts to causes of wholes.

One principle that might be proposed in this connection is that a cause of the occurrence of a collection of states is any collection of the causes of each. More particularly, a full cause of the occurrence of a collection of states is any collection of full causes of each. This principle clearly holds for any finite set of effects, where none of the causes of any member of the collection of effects is itself a member of the collection of effects. If a full cause of a is a' , of b is b' , of c is c' , and of d is d' , a , b , c , d , a' , b' , c' , and d' being distinct states, then a full cause of $a + b + c + d$ is $a' + b' + c' + d'$. If a full cause of one lamp's lighting up is its being connected to a battery, and a full cause of a second lamp's lighting up is its being connected to a different battery, then a full cause of the two lamps' lighting up is the connection of the two to batteries. This principle seems also to hold where the collection of effects is infinite, and none of the causes of any member of the collection of effects is itself a member of the collection of effects. If a full cause of the existence of every double-star system in the universe is the breaking-up of a single star, then a full cause of the existence of the double-star systems is still the breaking-up of single stars, even if the number of double-star systems is infinite.


However, the principle must be modified if it is to take account of cases where the cause of some member of a collection of effects is itself a member of that collection. For, when b is the cause of a , and c is the cause of b , the cause of $a + b$ is c , not $b + c$. If c is the lighting of a fuse, b is an explosion caused by c , and a an explosion caused by b , then the cause of $a + b$ is just c . To take account of this point, the previous principle must be expressed more generally as follows: a (full) cause of the occurrence of a collection of states is any collection of (full) causes of each, which are not members of the former collection. Hence, if a full cause of a is b , of b is b' , of c is d , and of d is d' , then a full cause of $a + b + c + d$ is $b' + d'$. If a full cause of a is b , of b is c , of c is d , and of d is e , then a full cause of $a + b + c + d$ is

¹² Hume, *Dialogues*, 59–60.

e. In so far as some member of the collection does not have a cause, to that extent the collection of states does not have a cause. If *a* has no cause, but *c* is a full cause of *b*, then there is no full cause of *a + b*, but *c* is a partial cause. Hence, in so far as a finite collection of states has a cause, it has its cause outside the collection.

Given all this, it follows from the stated principle that, if the universe is of finite age, and so its existence over time consists of no more than the occurrence of a finite collection of past states, each lasting for the same finite time (for example, a finite number of years), and the only causes of these past states are prior past states (that is, scientific causality alone operates), the whole collection of past states will have no cause and so no explanation. The same result follows if the universe is infinitely old (and so its history consists of the occurrence of an infinite collection of past states each lasting for the same finite time). The whole infinite series will have no full explanation at all, for there will be no causes of members of the series lying outside the series. In that case, the existence of the universe over infinite time will be an inexplicable brute fact. There will be an explanation (in terms of laws) of why, once existent, it continues to exist. But what will be inexplicable is its existence at all throughout infinite time. The existence of a complex physical universe over finite or infinite time is something 'too big' for science to explain.

But even for the infinite case the result follows only on the assumption that scientific causality alone operates. It does not follow merely from the fact that each state of the universe has a full explanation of a scientific kind. For that leaves open the possibility that there may be a deeper explanation, one in which the explanatory factors themselves are explained by a personal cause acting at the time when they act. In particular, the operation of the laws of nature may be due to an external cause whose action, together with the previous state of affairs, provides a complete explanation of any given state. This will be the case if a person *G* brings it about by a basic action at each period of time that the laws of nature *L* operate; and so brings it about that S_{n+1} brings about S_n . We thus get this picture:



G then ensures by an intention continuing over infinite time that the whole infinite series of states exists. If the universe is only of finite age and so has a first state, then there is the possibility that a person G causes this first state as a basic action, and, by subsequently continuously making the laws of nature to operate (again, by a basic action), keeps the universe in being. In either case G by his intention over the whole time while the universe has existed would—given that his own existence and intentions at each moment of time are uncaused—provide not merely a complete, but an ultimate explanation of the existence of the universe.

In his exposition of the cosmological argument Leibniz makes the point that a backwardly infinite series of states of universe, each explained by an earlier state without any cause acting from outside, would leave the existence of the universe at all totally unexplained.

Neither in any one single thing, nor in the whole aggregate and series of things, can there be found the sufficient reason of existence. Let us suppose the book of the elements of geometry to have been eternal, one copy always having been written down from an earlier one; it is evident that, even though a reason can be given for the present book out of a past one, nevertheless out of any number of books taken in order going backwards we shall never come upon a full reason; though we might well always wonder why there should have been such books from all time—why there were books at all, and why they were written in this manner. What is true of the books is true also of the different states of the world: for what follows is in some way copied from what precedes (even though there are certain laws of change). And so, however far you go back to earlier states, you will never find in those states a full reason why there should be any world rather than none, and why it should be such as it is. Indeed, even if you suppose the world eternal, as you will be supposing nothing but a succession of states and will not in any of them find a sufficient reason, nor however many states you assume will you advance one step forward giving a reason . . .¹³

Like Leibniz, I conclude that the existence of the universe over finite or infinite time would be, if only scientific explanation is allowed, a brute inexplicable fact. And, as Leibniz also did, I conclude that there is the possibility of an explanation of that existence in personal terms. The existence of the (physical) universe over time comes into my category of things too big for science to explain. If the existence of the universe is to be explained, personal explanation must be brought in,

¹³ G. W. Leibniz, *On the Ultimate Origination of Things*, trans. M. Morris, in *The Philosophical Writings of Leibniz* (Everyman edn., J. M. Dent and Sons, 1934), 31–2.

and an explanation given in terms of a person who is not part of the universe acting from without.

This argument will need to be recast if it is to be put in terms of the substances-powers-and-liabilities account of scientific explanation, in which causes are objects (substances) with powers and liabilities. The principle about the cause of a collection of states then becomes: 'a (full) cause of the occurrence of a collection of states is any collection of (full) causes of each, whose states as they cause are not members of the former collection'. But the S-P-L account has a different understanding from the amended Hempelian model of what are the states of objects. For the S-P-L account the 'why' in explanation is constituted by the powers and liabilities of substances (objects) that are properties of substances, and not by laws independent of the substances that they govern. So a full explanation of the existence of a substance will invoke a substance (either the same substance or another one) and its powers and liabilities; and a full explanation of the latter will also involve a substance, its powers, and liabilities. And the very factors that explain the existence of substances, the 'why' of explanation (in this model, powers and liabilities), now become explicable in the same terms as the 'what'.

It still follows, as before, that, if every state of the universe has a full cause in the universe at some earlier time causing it, then there will be no explanation within the scientific pattern of why there is a universe throughout all history, only of why it exists at any particular moment. But the question again arises as to whether the operations of each full cause itself depends on a more fundamental cause. As before, the latter cannot be a physical cause, for there are no physical causes apart from the universe itself and parts thereof. So the issue is whether a personal cause acting from outside the universe causes the causes within the universe to cause what they do. More precisely, the issue is whether the power of the universe to continue its existence into the next moment, and its liability to exercise that power, have no explanation at the time in question, or whether their existence and operation depend on a person who keeps them in existence and operating. Is the scientific explanation not merely a full, but a complete explanation; or does it itself have an explanation in terms of a person *G* who chooses to use the universe itself to keep the universe in being (as well as to bring it into existence, if it had a beginning)? If so, *G* by his continuing intention is the ultimate cause of there being a universe over infinite time.

So, either way, there is the possibility of a person *G* being the ultimate cause of there being a universe at all; and being the complete cause of its existence at any particular moment. The issue is in effect the same on the two models of explanation (whether we think of laws as separate from, or in a way properties of, the objects that they govern); but it will be simpler, for purposes of exposition, initially to pose the questions in terms of the amended Hempelian model; and then revert subsequently to the S–P–L model.

The Argument to God

For the reasons given in Chapter 3, the simplest explanation is, other things being equal, the most probable. Hence it is more probable, if there is such a *G*, that there is the simplest kind of *G*; and that—for the reasons discussed in Chapter 5—is a *G* of infinite power, knowledge, and freedom—that is, God. To postulate a *G* of very great but finite power, much but not all knowledge, etc., would raise the inevitable questions of why he has just that amount of power and knowledge, and what stops him from having more, questions that do not arise with the postulation of God. And even less simple, and so less probable, is polytheism, the supposition that the universe was created and is conserved by a committee of gods of limited power.

In the *Dialogues* Hume suggested that polytheism was at least as good an explanation of the existence of the universe as is theism: ‘A great number of men join in building a house or a ship, in rearing a city, in framing a commonwealth: why may not several deities combine in contriving and framing a world?’¹⁴ Hume again is aware of the obvious counter-objection to his suggestion. ‘To multiply causes without necessity is . . . contrary to true philosophy.’¹⁵ He claims, however, that the counter-objection does not apply here, because (in my terminology), although the supposition that there is one cause is as such a simpler supposition than the supposition that there are many, in postulating many gods of limited powers to be responsible for the order of the universe we are postulating gods more similar to humans in their power and knowledge than is the God of monotheism—that is, we are putting forward a hypothesis that fits in better

¹⁴ Hume, *Dialogues*, 39.

¹⁵ *Ibid.* 40.

with our background knowledge of what there is in the world. But whether this latter point has any force depends on just how we construe the argument. In putting forward the cosmological argument, I have assumed that we have no contingent background knowledge (we have mere 'tautological' knowledge) and are therefore looking only for the simplest explanation of the data. We must, of course, take into account all the evidence we have, including the evidence of what humans are like. But I bring in this evidence as data to be explained by the hypothesis—when we come on to consider (in Chapters 8 and 9) whether, if there is a God, it is probable that there would exist finite embodied human beings. And that is the proper place to consider this evidence—since theism purports to explain why there are such beings; according to theism their existence is not part of the framework of the world that would exist whether or not there was a God. To repeat a point made in Chapter 3—it is no objection to some theory of physics postulating fundamental particles and purporting thereby to explain the physical and thereby chemical behaviour of medium-sized objects that it postulates particles quite unlike those medium-sized objects (in having very small diameters etc). The issue is only whether what it postulates is relatively simple, and whether what is postulated makes it probable that we will observe what we do.

However, it may still be argued that, although any particular hypothesis of a finite god or many finite gods is less probable than theism, still there are so many possible hypotheses of finite gods with different degrees of power, and knowledge and different numbers of gods, that surely it is more probable that one of these is correct than that theism is correct.¹⁶ But, as I claimed in Chapter 3, consideration of the weight we give to simplicity in other areas of inductive inquiry suggests that we normally give it such weight that a really simple hypothesis is intrinsically more probable than a disjunction of many more complex hypotheses. For example—would a detective *really* think that it is more probable that some hypothesis or other to the effect that Jones cooperated with at least one of the few billion inhabitants of the earth to do some crime more probable than the hypothesis that Jones did it alone when both hypotheses give equal probability to the observed data?

¹⁶ For this objection, see Mark Wynn, 'Some Reflections on Richard Swinburne's Argument from Design', *Religious Studies*, 29 (1993), 325–35.

And, to jump ahead a bit, there are two further problems with polytheism as an explanation of the existence of not merely a universe but a universe governed throughout space and time by the same natural laws (a phenomenon that I shall discuss in the next chapter). If this order in the world is to be explained by many gods, then some explanation is required for how and why they cooperate in producing the same patterns of order throughout the universe. This becomes a new datum requiring explanation for the same reason as the fact of order itself. The need for further explanation ends when we postulate one being who is the cause of the existence of all others, and the simplest conceivable such—I urge—is God. And, further, the power of polytheism to explain this order in the world is perhaps not as great as that of theism. If there were more than one deity responsible for the order of the universe, we would expect to see characteristic marks of the handiwork of different deities in different parts of the universe, just as we see different kinds of workmanship in the different houses of a city. We would expect to find an inverse square of law of gravitation obeyed in one part of the universe, and in another part a law that was just short of being an inverse square law—without the difference being explicable in terms of a more general law.¹⁷ For these reasons, which apply to all the positive arguments for the existence of God that we shall be considering, the probability of theism is significantly greater than that of a hypothesis of one or more limited gods, and so in discussing these arguments in future I shall largely ignore the latter hypothesis.

So the choice is between the universe as stopping point and God as stopping point. In the latter case God's existence and intention over all the universe's history will provide not merely a full explanation but a complete and ultimate explanation of the existence of the universe.

Can we rest with the universe as a brute, inexplicable fact? Leibniz cannot, because the Principle of Sufficient Reason pushes him further. He writes:

¹⁷ Kant saw that, in so far as the 'physico-theological' argument, as he called the teleological argument, worked, it pointed to one deity and not many. He had no criticism of this part of the argument. 'The unity of this cause may be inferred from the unity of the reciprocal relations existing between the parts of the world, as members of an artfully arranged structure—inferred with certainty in so far as our observation suffices for its verification, and beyond these limits with probability in accordance with the principles of analogy' (*Critique of Pure Reason*, A625–6).

The reasons of the world then lie in something extra-mundane, different from the chain of states, or series of things, whose aggregate constitutes the world. And so we must pass from physical or hypothetical necessity, which determines the subsequent things of the world by the earlier, to something which is of absolute or metaphysical necessity, for which itself no reason can be given. For the present world is necessary physically or hypothetically, but not absolutely or metaphysically. In other words, when once it is determined that it shall be such and such, it follows that such-and-such things will come into being. Since then the ultimate root must be in something which is of metaphysical necessity, and since there is no reason of any existent thing except in an existent thing, it follows that there must exist some one Being of metaphysical necessity, that is, from whose essence existence springs; and so there must exist something different from the plurality of beings, that is the world, which, as we have allowed and have shown, is not of metaphysical necessity.¹⁸

Leibniz has here deployed the Principle of Sufficient Reason as a 'metaphysically necessary' truth. The principle boils down to the claim that everything not 'metaphysically necessary' has an explanation in something 'metaphysically necessary'. A being has 'metaphysical necessity' according to Leibniz, if from its 'essence existence springs'—that is, if it could not but exist. Whether this 'could not' is a logical 'could not' is, however, unclear. If the 'could not' is a logical 'could not', then the metaphysically necessary being with whom explanation is supposed to end is a logically necessary being. But against this claim that explanation ends with a logically necessary being I urge the point that I made in Chapter 4—the logically necessary cannot explain the logically contingent. Further, as I argued at the beginning of this chapter, it seems coherent to suppose that there exists a complex physical universe but no God, from which it follows that it is coherent to suppose that there exists no God, from which in turn it follows that God is not a logically necessary being. There can be no 'absolute explanation' of the existence of the universe.¹⁹

¹⁸ Leibniz, *On the Ultimate Origination of Things*, 33.

¹⁹ If this is correct, then, of course, the traditional ontological argument (see p. 8 n. 4), which attempts to prove that God exists of logical necessity, fails utterly. Kant accused the cosmological argument of being the ontological argument in disguise. His criticisms have force only if we suppose that the necessary being to which the cosmological argument purports to argue is a logically necessary being. See Kant, *Critique of Pure Reason*, sect. 5: The Ideal of Pure Reason. For detailed discussion of Kant's treatment of the cosmological argument, see J. Bennett, *Kant's Dialectic* (Cambridge University Press, 1974), ch. 11, and

If, however, Leibniz's metaphysically necessary being is not a logically necessary being, but the supreme brute fact, then his principle boils down to the simple claim that there is a terminus to explanation, that everything that has a full explanation has an ultimate, or at least a complete, explanation. In an earlier chapter we saw no general reason for accepting this claim. Whether it is rational to suppose that phenomena have complete explanations is a matter of whether we have potential explanations for them of great simplicity and explanatory power. Leibniz claims that the universe is not metaphysically necessary, and so that its existence needs explanation. He may be right, but I cannot see how you can argue for this claim except in terms of the relatively greater simplicity and explanatory power of a potential explanatory hypothesis. Leibniz does not provide such an argument. So I need to consider the force of an argument along these lines for the existence of God.

We saw earlier that $P(h|e \& k)$ will exceed $P(h|k)$ and so there be a correct C-inductive argument from e to h if (and only if) $P(e|h \& k)$ exceeds $P(e|k)$ —that is, if e is more to be expected given h than it would be otherwise. Let h be the hypothesis of theism, and k be mere tautological evidence. Let e be the existence over time of a complex physical universe. e could not, as we have seen, have a scientific explanation. Either e occurs unexplained, or it is due to the action of a person, the most likely person being God. It is now reasonable to ignore the alternative that we have shown to be *a priori* much less probable, that e was brought about by a person or persons of very large but finite power, very considerable but limited knowledge, etc. Hence we may regard $P(e| \sim h \& k)$ as the probability that there be a physical universe without anything else having brought it about.

references contained there. Kant holds that all necessity is either logical necessity or, more widely, necessity for human thought. Either way, for Kant, there is no necessity in things, only in our thought about them. ('The concept of necessity is only to be found in our reason, as a formal condition of thought; it does not allow of being hypostasized as a material condition of existence' (*Critique of Pure Reason*, A620).) This is another of the general principles of Kant's philosophy that he brings to the philosophy of religion, and that spoils his treatment of it. Kant's principle is quite mistaken. Clearly there is a necessity in the conformity of material bodies to laws of nature that exists in things and not in our thought. See my Chapter 2. That suggests that there may be other kinds of necessity in things, and so led me to suggest in Chapter 5 that God's necessity amounts to his being the ultimate brute fact. For more extended discussion of kinds of necessity, see *The Coherence of Theism* (Clarendon Press, 1993), ch. 13.

A complex physical universe (existing over endless time or beginning to exist at some finite time) is indeed a rather complex thing. We need to look at our universe and meditate about it, and the complexity should be apparent. There are lots and lots of separate chunks of it. The chunks each have a different finite and not very natural volume, shape, mass,²⁰ etc.—consider the vast diversity of the galaxies, stars, and planets, and pebbles on the seashore. Matter is inert and has no powers that it can choose to exert; it does what it *has* to do. There is just a certain finite amount, or at any rate finite density of it, manifested in the particular bits; and a certain finite amount, or at any rate finite density of energy, momentum, spin, etc. There is a complexity, particularity, and finitude about the universe that cries out for explanation, which God does not have. Or, rather, that is the case with the universe as it now stands. But, as I discussed earlier, all the evidence suggests that the universe evolved from a much simpler state in accord with laws of nature ensuring that such a universe would develop into a large complex universe. But the earlier state still needed to have had some complexity to it—either it had to consist of extended matter-energy or of an unextended point with laws of nature that were such as to give rise to a complex material universe. Put in terms of the S–P–L model—either each bit of the extended matter-energy had to have powers to keep itself in existence, or the unextended point had to have powers to produce a large number of chunks of matter, with (in each case) the liability to continue to exercise these powers. Either way, this starting point would be a finite thing with certain ways of developing built into it and no reason why those particular ways of developing should be built into it, rather than any other ways. There could have been no laws of nature and so complete chaos, or laws that soon ensured the complete elimination of the universe. When the universe began, it had to have a certain kind of complexity built into it if there was to result a complex physical universe. And, if the universe did not begin to exist, it would have had always to have a certain kind of complexity (its matter-energy would have had always to have the relevant powers).

As we saw at some length in Chapter 5, the supposition that there is a God is an extremely simple supposition; the postulation of a God

²⁰ I shall come in Chapter 8 to the point that the building blocks of which the chunks are made—i.e. the fundamental particles—are of only a few distinct kinds.

of infinite power, knowledge, and freedom is the postulation of the simplest kind of person that there could be. God is an unextended object, the divine properties fit together, and they are properties of infinite degree; we saw in Chapter 3 that infinite degrees of a property have a simplicity lacked by large finite degrees of the same property. *A priori* the existence of anything at all logically contingent, even God, may seem vastly improbable, or at least not very probable. (Hence 'the mystery of existence'.) Yet, whether this is so or not, the existence of the universe is less simple, and so less to be expected *a priori* than the existence of God. Hence, if there is no God, the existence of a complex physical universe is not much to be expected; it is not *a priori* very probable at all—both because (it may well seem) it is vastly improbable *a priori* that there would be anything at all; and because, if there is anything, it is more likely to be God than an uncaused complex physical universe. $P(e|\sim h \& k)$ is low.

Yet, if there is a God, clearly he can create a universe; and he will do so in so far as his perfect goodness makes it probable that he will. I argued in Chapter 6 that God has good reason to create humanly free agents—that is, creatures with limited free choice between good and evil and limited powers to make deeply significant differences to themselves, each other, and the physical world by those choices, and also (because of the evil they may produce) reason not to create such creatures. I argued that it would be an equal best act to create or not to create such creatures, and so we should suppose the logical probability that God would create such creatures to be $1/2$. I argued that these creatures would need to have bodies, and so there would need to be a physical world. So for this reason alone the probability that a God will create a physical world will be no less than $1/2$. And God might create a physical universe, even if he does not create humanly free agents. And so perhaps the probability that God would create a physical world might be greater than $1/2$. Yet, of course, our judgements as to what a perfectly good God might do may be in error, because our views of moral goodness are limited; and, as I emphasize throughout this book, we cannot often give more than the roughest of values to the probabilities that I discuss. Nevertheless, if the moral intuitions about what a perfectly good God would do that I am commending to my readers are in any way close to the truth, we must conclude that the logical probability that, if there is a God, there will be a physical universe is quite high. Since $P(e|h \& k) > P(e|\sim h \& k)$ and so $> P(e|k)$, by the

relevance criterion $P(h|e \& k) > P(h|k)$, and so the argument from the existence of a complex physical universe to God is a good C-inductive argument.

The argument of the last few pages can now be put in simple words as follows. There is quite a chance that, if there is a God, he will make something of the finitude and complexity of a universe. It is very unlikely that a universe would exist uncaused, but rather more likely that God would exist uncaused. Hence the argument from the existence of the universe to the existence of God is a good C-inductive argument.

Teleological Arguments

I understand by an argument from design one that argues from some general pattern of order in the universe or provision for the needs of conscious beings to a God responsible for these phenomena. An argument from a general pattern of order I shall call a teleological argument. (The name 'teleological argument' has usually been used to characterize much the same arguments as 'argument from design'. I am giving 'teleological argument' a narrower use.) I shall deal with teleological arguments in this chapter. I shall deal in Chapter 10 with the argument from the occurrence of provision for the needs of conscious beings, and I shall call such an argument an argument from providence. In the definition of 'teleological argument' I emphasize the words 'general pattern'; I shall not count an argument to the existence of God from some particular pattern of order manifested on a unique occasion as a teleological argument.

I begin with the distinction between spatial order and temporal order, between what I shall call regularities of co-presence and regularities of succession. An example of a regularity of co-presence would be a town with all its roads at right angles to each other, or a section of books in a library arranged in alphabetical order of authors. Regularities of succession are simple patterns of behaviour of objects such as someone moving his or her legs in accord with the standard movements of a dance. In both these cases the regularities are produced by humans. The universe is characterized by regularities of both kinds not produced by humans or other embodied beings. There is first the temporal order of the regular successions of events, codified in laws of nature. In books of physics, chemistry, and biology we can learn how almost everything in the world behaves. The laws of their behaviour can be set out by relatively simple formulae that humans can understand and by means of which they

can successfully predict the future. The orderliness of the universe to which I draw attention here is its conformity to formulae, to simple, formulable, scientific laws. The orderliness of the universe in this respect is a very striking fact about it. The universe might so naturally have been chaotic, but it is not—it is very orderly. And then there is the spatial order of the intricate arrangement of parts in human (and animal) bodies. We have limbs, liver, heart, kidneys, stomach, sense organs, etc. of such a kind that, given the regularities of temporal order, our bodies are suitable vehicles to provide us with an enormous amount of knowledge of the world and to execute an enormous variety of purposes in it (as described more fully in Chapter 6). This is similar to the way in which parts of machines are arranged so as to produce an overall result from the operation of the machine; though—so far—machines intentionally constructed by humans are far less intricate than human bodies.

A teleological argument, whether from temporal or spatial order, is, I believe, a codification by philosophers of a reaction to the world deeply embedded in the human consciousness. Humans see the comprehensibility of the world as evidence of a comprehending creator. The prophet Jeremiah lived in an age in which the existence of a creator-god or gods of some sort was taken for granted. What was at stake was whether there was only one god, and the extent of his goodness, knowledge, and power. Jeremiah argued from the order of the world that there was one powerful and reliable god, and that god was God. He argued to the power of the creator from the extent of the creation: ‘The host of heaven cannot be numbered, neither the sand of the sea measured’; he argued that its regular behaviour showed the reliability of the creator, and he spoke of the ‘covenant of the day and night’ whereby they follow each other regularly, and ‘the ordinances of heaven and earth’,¹ and he used their existence as an argument for the trustworthiness of the Jewish God. The argument from temporal order has been with us ever since.

The Datum of Temporal Order

We find the argument from temporal order also in Aquinas’s fifth way, which runs as follows:

¹ Jer. 33: 20 and 25.

The fifth way is based on the guidedness of nature. Some things lacking awareness seek a goal—which is apparent from the fact that always or most usually they behave in the same way which leads to the best result. From this it is evident that it is not by chance but by intention that they reach their goal. Nothing, however, that lacks awareness tends to a goal, except under the direction of someone with awareness and with understanding; the arrow, for example, requires an archer. Everything in nature, therefore is directed to its goal by someone with understanding and this we call ‘God’.²

Aquinas argues that the regular behaviour of each inanimate thing shows that some animate being is directing it (making it move, so as to achieve some purpose, attain some goal); and from that he comes—rather quickly—to the conclusion that one ‘being with understanding’ is responsible for the regular behaviour of all inanimate things (apart, maybe, from the behaviour for which humans and animals are responsible).

It seems to me fairly clear that no argument from temporal order, whether Aquinas’s fifth way or any other argument, can be a good deductive argument. For, although the premiss is undoubtedly correct—a vast pervasive order characterizes the world—the step from premiss to conclusion is not a valid deductive one. Although the existence of order may be good evidence of a designer, it is surely compatible with the non-existence of one—it is hardly a logically necessary truth that all order is brought about by a person. And although, as I have urged, the supposition that one person is responsible for the orderliness of the world is much simpler and so more probable than the supposition that many persons are thus responsible, nevertheless the latter supposition seems logically compatible with the data. So we must turn to the more substantial issue of whether the argument from the temporal order of the world to God is a good (C- or P-)inductive argument. Since just the same kind of considerations apply to all other claims that some argument from an observable feature of the world to the existence of God is a valid deductive argument, I shall not in future bother to repeat them where we come to new arguments. I shall assume that no such argument is deductively valid.

But, before considering whether the argument from temporal order is a good inductive argument, I must deal with three preliminary matters. First, there is an objection that this temporal order is

² St Thomas Aquinas, *Summa Theologiae*, Ia2.3, my translation.

not an objective feature of the world but a mere human artefact; the order that we seem to see in the world is order that we impose on it, and is not there independently of our imposition. Put another way, all that this temporal order amounts to, it might be said, is a coincidence between how things have been so far in the world and the patterns that humans can recognize and describe. In fact, however, the temporal order of the world is something deeper than that. We rightly explain our observations so far in terms of laws of nature involving a physical necessity in nature (as I analysed on pp. 28–9), which determine how things behave and allow us to predict how they will behave in future. It is the operation of these simple natural laws that this argument seeks to explain.

Then there is the objection that there is nothing to be explained in the fact that we find an orderly universe—for we could not possibly find anything else. For, unless the universe were an orderly place, we would not be here to comment on the fact. (If there were no natural laws, there would be no regularly functioning organisms, and so no human beings.) Hence there is nothing surprising in the fact that humans find order—we could not possibly find anything else. This conclusion is clearly a little too strong. There would need to be quite a bit of order in and around our bodies if we are to exist and observe the world, but there could be chaos outside the earth, so long as the earth was largely unaffected by that chaos. There is a great deal more order in the world than is necessary for the existence of humans. So humans could still be here to comment on the fact of order even if the world were a much less orderly place than it is. But, quite apart from this minor consideration, the argument still fails totally for a different reason, which can be brought out by an analogy. Suppose that a madman kidnaps a victim and shuts him in a room with a card-shuffling machine. The machine shuffles ten packs of cards simultaneously and then draws a card from each pack and exhibits simultaneously the ten cards. The kidnapper tells the victim that he will shortly set the machine to work and it will exhibit its first draw, but that, unless the draw consists of an ace of hearts from each pack, the machine will simultaneously set off an explosion that will kill the victim, in consequence of which he will not see which cards the machine drew. The machine is then set to work, and to the amazement and relief of the victim the machine exhibits an ace of hearts drawn from each pack. The victim thinks that this extraordinary fact needs an explanation in terms of the machine having been

rigged in some way. But the kidnapper, who now reappears, casts doubt on this suggestion. 'You ought not to be surprised', he says, 'that the machine draws only aces of hearts. You could not possibly see anything else. For you would not be here to see anything at all, if any other cards had been drawn.' But of course the victim is right and the kidnapper is wrong. There is indeed something extraordinary in need of explanation in ten aces of hearts being drawn. The fact that this peculiar order is a necessary condition of the draw being perceived at all makes what is perceived no less extraordinary and in need of explanation. The teleologist's starting point is not that we perceive order rather than disorder, but that order rather than disorder is there. Maybe only if order is there can we know what is there, but that makes what is there no less extraordinary and in need of explanation.

The third preliminary matter is to note the kinds of regularities to which the argument appeals. The regularities of temporal succession in our universe are of two kinds. There are the phenomenal regularities that are very rough probabilistic laws about what happens on perhaps 97 per cent of occasions; and there are the fundamental regularities that explain these. The phenomenal regularities are the macroscopic ones by which humans (and the higher animals) guide their daily life, ones evident to people without much scientific education. They include such regularities as that seeds when watered often grow into plants, people who do not eat or drink for a month or two die, mushrooms nourish but toadstools poison, arrows shot quickly penetrate human skin, day is followed by night and night by day at approximately similar intervals (as measured by pendulum clocks), and so on. But scientists have discovered that these phenomenal regularities are brought about by more fundamental regularities. The phenomenal regularities arise from laws of chemistry about possible combinations of atoms into molecules and the resulting stability of solid objects; and these are brought about by laws of physics governing the electrons, protons, and neutrons of which atoms are made; and these are brought about by the laws governing the quarks of which the protons and neutrons are made and so on. These laws at the latter level include the laws of the four forces (gravity, electromagnetism, strong force, and weak force) constrained by the general requirements of Quantum Theory and the General Theory of Relativity. Probably the laws of electromagnetism and the weak force derive from the more general laws of an 'electroweak'

theory; and there is some reason to suppose that, in due course, physicists will discover a ‘theory of everything’, whose laws have (within science) no more complete explanation,³ and which explain all physical phenomena. So the physical world is such that it is governed by relatively simple fundamental laws (deterministic, or—more likely—probabilistic), concerning the tiny unobservable building blocks of the world, of such a kind that they often lead to laws about the observable behaviour of medium-sized entities. Not all behaviour of physical objects at the phenomenal level is governed by simple regularities—the behaviour of the pendulum is, the behaviour of the weather is not. And phenomenal regularities do not concern what always happens, only what happens almost always; and in that way they are very reliable, but not totally reliable. Buildings, bridges, and trees normally remain immovable, but just occasionally they collapse. Peanuts normally nourish, but very occasionally they poison.

It is the (almost entirely reliable) phenomenal regularities that we observe and then use in order to bring about our chosen goals. Regularities (fundamental or phenomenal) have to exist if we are to be able to make a difference to things beyond our bodies. If we are to grow plants, it has to be the case that certain basic actions of ours will have this result. But, unless we are to be non-rational creatures, we need to be able to observe phenomenal regularities and learn from them. Humans (and often too higher animals) can

³ It may be urged that we have no reason to suppose that there is a most fundamental law of nature. Maybe law L operates in circumstances C because it follows from L' that it does; and L' operates in circumstances C' (which include C) because it follows from L'' that it does; L'' operates in circumstances C'' (which include C') because it follows from L''' that it does, and so on *ad infinitum*. This difficulty can, however, be avoided as follows. Either such a series ends with a most fundamental law that holds in all circumstances, or it does not. In the latter case let us represent the more fundamental laws as conjunctions of laws that hold without exception under specifiable circumstances such as C . Thus to say L' holds in C' will be to say that L holds in C and L_1 in C_1 ; to say that L'' holds in C'' will be to say as well as this that L_2 holds in C_2 . Then the claim that there is an infinite series, $L, L', L'',$ etc. is the claim that there exists an infinite series of non-fundamental laws $L, L_1, L_2,$ etc. that hold without exception in circumstances $C, C_1, C_2,$ etc.; and that, although there is an explanation of the operation of any finite subseries, there is no explanation of the operation of the whole series. That the whole series operates will then be the starting point for a teleological argument; that it operates shows a conformity of the world to order similar to that shown by the conformity of the world to storable most fundamental laws of nature, which form the starting point for the simpler argument. I shall henceforward deal with the simpler argument on the not implausible assumption that there are most fundamental laws of nature.

observe seeds being watered, or day being followed by night, and extrapolate to the regularities described as the simplest account of what they observe (that is, they can infer that what holds for the seeds that they have observed holds for seeds generally, and so on.) They can then rely on these regularities to produce effects. They can water seeds and grow plants. Wanting to travel a long distance easily, they can travel by day and not by night (since they know that day will come again soon). Knowledge of such regularities gives to humans choices. Discovering that toadstools poison, they can choose to poison someone by encouraging him or her to eat toadstools; or they can prevent accidental poisoning by toadstools by uprooting toadstools and telling people that they are poisonous; or they can choose not to bother.

Now in order for any of its inhabitants to attain their goals, the universe needs to evince regularity at some level or other. Then those inhabitants will need to have sense organs sensitive to how things are at that level if they are to detect and use the regularities. In our universe both the fundamental regularities and many less fundamental regularities are relatively simple. I have called the latter regularities the phenomenal regularities, because—as far as we know—these are the only regularities to which the rational creatures of our universe are sensitive. But, if there are creatures whose senses inform them without the help of apparatus or inference of the location of individual atoms, then they can use the regularities in the behaviour of these atoms to attain their goals. There is, however, no guarantee that regularity at a fundamental level (in the behaviour of fundamental particles) will ensure usable regularity at a less fundamental level. Whether it does will depend on what the less fundamental laws are and on the boundary conditions of the universe. Even in our universe, if the temperature never became low enough for atoms to combine to form medium-sized solid objects, there would be only clouds of gas or liquid that do not as such behave in very simple ways. It is also possible to have laws of nature fairly complex on the small scale, producing many fairly simple regularities on the large scale. For it might be that the boundary conditions of the universe were such that fundamental particles were normally found only in states that easily allowed them to combine into larger objects, and the behaviour of the particles in combination was restricted to some very simple patterns. But clearly total chaos at a fundamental level will lead to chaos at any other level.

The argument from temporal order is an argument from regularity at some level or other. And, while the operation of non-fundamental laws may be explained by the operation of fundamental laws, that these are fundamental laws of nature is, like the very existence of a complex physical universe, where science starts from in order to explain other things. It is something 'too big' for science itself to explain.

The Probability of Temporal Order in a Godless Universe

So how probable is it intrinsically that in a Godless universe there will be laws of nature at some level guaranteeing that things behave in very largely predictable ways? The answer to this question depends to some extent on what laws of nature are. I discussed in Chapter 2 three theories of this. There is first the immensely implausible Humean account developed by Lewis—that the conformity of all objects to laws of nature is just the fact that they do so conform; there is no more fundamental explanation of this conformity. It is just a brute fact that (both at a fundamental level and at a phenomenal level) objects (substances) fall into kinds (electrons, positrons, pendula, seeds) in such a way that the simplest extrapolation from their past behaviour leads to generalizations that predict their future behaviour more or less correctly. In the near past, as in the more remote past, every positron has continued to attract every electron with exactly the same force inversely proportional to the square of their distances apart. There are innumerable other ways in which objects could have behaved, almost all of them such that the simplest extrapolation from their past behaviour would not have correctly predicted their subsequent behaviour. It is only if there is a common explanatory cause of the behaviour of objects that there is any reason to suppose that they will behave in the same way. And in a Godless universe on the Humean theory of laws of nature there is no more fundamental explanation of the coincidence in the ways in which objects behave. On this view 'laws' do not really explain the behaviour of objects, they merely describe it.

Alternative accounts of laws of nature represent talk of 'laws' as talk about a feature of the world additional to the mere succession of events, a feature of physical necessity that is part of the world. As we saw in Chapter 2, this feature of physical necessity may be thought of

either as separate from the objects (substances) that are governed by it, or as a constitutive aspect of those objects. The former approach leads to a picture of the world as consisting of events (constituted perhaps by substances with their properties), on the one hand, and laws of nature, on the other hand; the most common version of this view claims that laws of nature are logically contingent relations between universals. The conformity of all objects to simple laws of nature consists on this account of the instantiation of quite a few universals each connected in simple ways to one or two other universals. If, despite the difficulties raised in Chapter 2, we adopt this account, the first question is why should there be universals connected to each other before they are instantiated, and why—if there is a universe, and so some universals must be instantiated—should quite a few universals be instantiated in such a way as to form a whole system of laws of nature. There might be many universals that were instantiated without bringing any other universals with them, so that there was no predictable effect of the instantiation. But on this account virtually all universals are connected to other universals. And there might be universals, but only ones of kinds instantiated once or twice in the history of the universe, rather than ones like ‘photon’ or ‘copper’ that are instantiated often and so can be used for useful prediction. And, again, the mathematical connections between the universals—for example, between the masses of bodies, their distance apart, and the gravitational attraction between them—might be of such complexity as never to be inferable from the past behaviour of objects.

Now I suggest that a universe without connections between universals would be simpler than one with connections; and one with simpler patterns of connection would be simpler than one with such complicated patterns of connection that rational beings would not be able to infer the future behaviour of objects by means of the simplest extrapolation from their past behaviour. Among theories of the universe as a whole (which will thus have equal scope), simplicity is the sole indicator of intrinsic probability. It then follows that, if we give it the weight that I have urged that we should (so that a very simple theory is more probable than a disjunction of many more complex theories), it would be very probable that there would be no connections between universals at all—that the universe would be chaotic. But note that, if we give simplicity much less weight and suppose that a simpler theory is merely somewhat more probable

than a more complex theory, it might be that it is more probable that one of a disjunction of alternative sets of fairly simple connections between universals holds rather than no connections at all. But in that case, since there are a very large number of complex ways in which universals could be associated, and we are giving simplicity only a moderate weight, then it will be at least as probable that one of the complex connections between universals will hold as that one of the simple connections will hold—there being so many more (infinitely many more) of the former. Either way, it is going to be improbable that in a Godless universe there will be simple connections between universals, and so simple laws of nature.

The same issues arise on the substances-powers-and-liabilities account of laws of nature. On this account, powers and liabilities are among the properties of substances. Laws of nature are then just contingent regularities—not of mere temporal succession (as with Hume), but of causal succession, regularities in the causal powers (manifested and unmanifested) of substances of various kinds. The conformity of all objects to simple laws of nature consists on this account in all substances falling into very few kinds with the same powers and liabilities as each other. Why does this happen? The S–P–L model has an initial answer to this question; it can provide an explanation of this fact in a way that the other two models cannot provide an explanation of the corresponding fact. The S–P–L model's answer is an answer in terms of the causal ancestry of substances. A substance has the powers and liabilities it does because it was produced by another substance exercising (in virtue of some liability to do so) its power to produce a substance with just those powers and liabilities. If a proton is produced (together with an electron and an neutrino) by the decay of a neutron, then the proton's powers and liabilities are caused by the neutron, in virtue of its powers and liabilities. There are then different ways in which it could have come about that all substances fall into a small number of kinds in the way described, according to whether this process had a beginning and of what kind that beginning was.

Suppose, first, that the universe did have a beginning. There are two different kinds of theories of a beginning. The first state might have been a spatially extended state, or a spatially pointlike state. In the first case, we would still have a lot of substances, but perhaps crammed into a very small space. In terms of the Big Bang model,

there would not have been literally a singularity; it would just have been that, as you approach the first instant in the temporally backward direction, you would find denser and denser states; but it really all started in a very but not infinitely dense state. If that state was to give rise to our present universe of very few kinds of substance, it must itself have consisted of a very large number of substances of very few kinds. The alternative first state would be a literally pointlike one. At the first instant of the universe's history, on this theory, there was an unextended point, endowed with the power to decay into innumerable substances of very few kinds, and a liability to exercise that power at some time or other. Suppose, secondly, that the universe had an infinite age. The properties (of powers and liabilities) of every substance are then caused by those of a preceding substance. So there can be substances with exactly the same such properties (including the powers to produce substances of the existing kinds) only if there always have been.

Study of the present data of physics and cosmology will allow us to say roughly how probable on those data are the three different theories—on the basis of how probable it is that we would find those data, given each of the theories, and of how simple are the different theories. My assessment of the present state of cosmology is that a beginning is more probable than an infinite age; and that evolution from a very dense state is more probable than evolution from an infinitely dense state. (All matter-energy occupying an unextended point is, as I suggested in the previous chapter, not a possibility allowed by the current theory of matter-energy, which would require considerable complication in order to allow for this while continuing to make the present data probable.) But, of course, new data could change the probabilities.

The issue for us, however, is not what are posterior probabilities on the physical data that the different theories are true, but how probable is it *a priori* if there is no God that the true theory will be such as to lead to substances of a very few kinds. This will depend solely on the relative simplicity of the three theories (since they are of equal scope), and the probability on each of these theories that substances of such kinds would result. Simplicity is the sole relevant *a priori* criterion. There is no doubt that the theory that the universe began at a point is simpler and so intrinsically more probable than any particular theory that it began with many substances, so much simpler that I suggest that it is more probable than the disjunction of all

theories claiming that it began from or always contained many substances.

But, if it did begin from an unextended point, the simplest theory of such a beginning would seem to be that that point would have had no power to produce extended substances. If it did have such a power, it would seem simpler to suppose that it would have the power to produce just one extended substance. The theory that it would have the power to produce extended substances falling into few kinds, themselves having the power to produce more such substances, all with the liability to exercise these powers from time to time, seems just one of a number of equally simple theories, less simple than that the theory that the unextended point had no power or only the power to produce just one extended substance. But any theory that at a beginning or always there were many substances, which fall into kinds with identical powers and liabilities, is again a theory of a very improbable coincidence. Such a coincidence cries out for explanation in terms of some single common source with the power to produce it. Just as we would seek to explain all the coins of the realm having an identical pattern in terms of their origin from a common mould, or all of many pictures having a common style in terms of their being painted by the same painter, so we should seek to explain all physical objects having the same powers in terms of their deriving them from a common source. So again, on the S-P-L account of laws of nature, as on the universals account (and *a fortiori* on the Humean account), it is very improbable that there would be in a Godless universe laws of nature sufficiently simple for rational beings to extrapolate from past to future with normal success.

The Probability of Temporal Order, given Theism

Theism leads us to expect (with significant probability) that God will bring about humanly free agents, as described in Chapter 6. They will be embodied creatures, and will start with limited power and knowledge. If they are to extend their power, they must discover which of their basic actions will have which more remote effects in which circumstances—for example, which sequence of basic actions done in which circumstances will lead to a house being built, and which sequence will lead to a bomb being built. Only with this knowledge will they have a choice of whether to build houses or bombs. For

there to be such recipes for producing different effects that creatures can discover, there need to be simple regularities in the behaviour of things that creatures can note and use. It has to be the case that this brick put on top of cement that is on top of another brick will stay there and resist much pressure, and so on and so forth. Theism leads us to expect a world at some phenomenal level, simple and reliable. It leads us to expect that God will bring about an initial singularity of the right kind, or an initial arrangement of substances with the same powers and liabilities of the right kind and keep them in existence, or that he will always have kept in existence such substances. (Or, on the universals model of laws of nature, theism leads us to expect that God will produce the right kinds of connections between universals. On the Humean model of laws of nature, theism leads us to expect that God will make things behave in simple regular ways.)

I have been assuming so far that there is only one universe. But there may be many universes. If there were actually existing all possible universes, some of them will be law governed and it might be expected that we would find ourselves in such a universe. However, it would be the height of irrationality to postulate innumerable universes just to explain the particular features of our universe, when we can do so by postulating just one additional entity—God. Science requires us to postulate the simplest explanation of the data, and one entity is simpler than a trillion. In order rationally to postulate other universes we would need to find new data in our universe best explained by postulating that there are also other universes. In particular, we would need to have data such that extrapolating back from the present state of our universe in accord with the mathematically simplest supposition about what are its laws that would explain these data would lead us to a state at which there was a universe split, a state in which those laws will have dictated that another universe would ‘bud off’ from our universe. But in that case the other universe would be governed by the same fundamental laws as govern our universe, and so we can consider the two universes (or however many universes we learn about) as one multiverse, and the whole preceding structure of argument gives the same results as before. So it does not affect the issue of why things are law-governed if we suppose (on good evidence) that there is more than one universe.⁴ And I have

⁴ Hume, *Dialogues concerning Natural Religion* (first published 1779, ed. H. D. Aiken (Hafner, 1948), 53–5) suggested a temporal version of many universes—that perhaps this ordered universe is a mere accident among the chance arrangements of eternal matter. In

argued that whether talk of ‘laws’ is talk of regular successions of events, of concrete entities determining the behaviour of substances, or of the powers and liabilities of substances, it is *a priori* improbable that a Godless universe would be governed by simple laws, but there is quite a significant probability that a God-created universe would be governed by simple laws. Hence the operation of laws of nature is evidence—one strand of a cumulative argument—for the existence of God.

Let us represent by e this conformity of the world to order, and let h be the hypothesis of theism. It is not possible to treat a teleological argument in complete isolation from the cosmological argument. We cannot ask how probable the premiss of the teleological argument makes theism, independently of the premiss of the cosmological argument, for the premiss of the teleological argument entails the premiss of the cosmological argument. That there is order of the kind described entails that there is a complex physical universe. So let k be now, not mere tautological evidence, but the existence of a complex physical universe (the premiss of the version of the cosmological argument to which I devoted attention). Let us ask how much more probable does the orderliness of such a universe make the existence of God than does the mere existence of the universe. As we have seen, $P(h|e \& k)$ will exceed $P(h|k)$ if and only if $P(e|h \& k) > P(e|\sim h \& k)$. Put in words with our current fillings for h , e , and k , the existence of order in the world confirms the existence of God if and only if the existence of this order in the world is more probable if there is a God than if there is not. The arguments of the previous pages have sought to show just this; and indeed that the probability of order of the right kind is very much greater if there is a God, and so that the existence of such order adds greatly to the probability that there is a God.

the course of eternity matter arranges itself in all kinds of ways. We just happen to live in a period when it is characterized by order, and mistakenly conclude that matter is always ordered. Hume is certainly right in claiming that this is a logical possibility, but the point made above remains—that it is irrational to postulate other universes, unless features of this universe give support to laws that have the consequences that there are other universes; from which it would follow that they are governed by the same fundamental laws as our universe.

The Argument from Spatial Order

Those who marvel at the order of the universe may marvel at either or both of the regularities of co-presence and of succession. The thinkers of the eighteenth century to whom a teleological argument appealed so strongly were struck almost exclusively by the regularities of co-presence. They took the regularities of succession largely for granted. What struck them was the subtle and coherent arrangement of parts in human and animal bodies and in plants, enabling humans and animals to acquire an enormous amount of knowledge and to execute an enormous variety of purposes, including to reproduce their kind, and enabling plants to grow and multiply (without choosing to do so). Paley's *Natural Theology* dwells mainly on details of comparative anatomy, on eyes and ears and muscles and bones arranged with minute precision so as to operate with high efficiency, and in the *Dialogues* Hume's Cleanthes produces the same kind of examples: 'Consider, anatomize the eye, survey its structure and contrivance, and tell me from your own feeling, if the idea of a contriver does not immediately flow in upon you with a force like that of sensation.'⁵

The eighteenth-century argument from spatial order seems to go as follows. Humans, animals, and plants have the power to reproduce their kind, and so, given their past existence, their present existence is to be expected. But what is vastly surprising is the existence of humans, animals, and plants at all. By natural processes they can come into being only through generation from organisms of the same species. But, it was claimed, the world has not been going on for ever, and so the great puzzle is the existence of the first humans, animals, and plants in 4004 BC or whenever exactly it was that they began to exist.⁶ Since they could not have come into existence by natural scientific processes, and since they are very similar to the machines that certain rational agents—namely, humans—make, it is very probable that they were

⁵ Hume, *Dialogues*, 28.

⁶ Even if they had supposed that the world had been going on for ever and had contained humans, animals, and plants for ever, the thinkers of the eighteenth century could still have constructed an argument from the eternal existence of a universe containing humans, animals, and plants rather than a universe not containing humans, animals, and plants; but the argument would have to have been a more subtle one than the one that we are considering.

made by a rational agent—only clearly one much more powerful and knowledgeable than humans.

Hume's objections in the *Dialogues* (through the mouth of Philo) to a teleological argument are directed against the argument from spatial order, although—if they worked—some of them would also have had force against the argument from temporal order and so I have considered them in that connection. Despite Hume's objections, the argument is, I think, a very plausible one—given its premisses. But one of its premisses was shown by Darwin and his successors to be clearly false. Humans can be produced through generation from complex animals, and complex animals and plants can be produced through generation from less complex animals and plants—species are not eternally distinct; and simple animals and plants can be produced by natural processes from inorganic matter. This discovery led to the virtual disappearance of the teleological argument from popular apologetic—mistakenly, I think, since it can easily be reconstructed in a form that does not rely on the premisses shown to be false by Darwin. The basic mistake of those who regarded Darwin's discoveries as destructive of the argument from spatial order is that they ignored the fact that only certain processes acting on a certain initial kind of inorganic matter would have produced human bodies (and animals and plants); and that it is *a priori* improbable that the processes and initial matter would be of the right kind, but that this is to be expected if theism is true.

The argument is, I think, best treated not as an argument from analogy (the way typical of the eighteenth century) but in the way in which other arguments are treated in this book, as an argument from evidence that it would be probable would occur if theism is true, but not otherwise. The argument must be construed as an argument from human bodies, not humans. The argument from human bodies being connected to a mental life is a separate argument, to be analysed in the next chapter. We also have the evidence of animal bodies and plants. To make the exposition simpler I shall largely ignore these latter until towards the end of this chapter.

The argument from human bodies must be construed as an argument from the existence of bodies having certain features characteristic of human bodies. They are features that the body of a humanly free agent, as defined in Chapter 6, would need to have. To be the body of a humanly free agent, a body needs to be suited for the acquisition of

true beliefs about the environment, the formation of purposes in the light of desires, and the expression of them via chosen basic actions designed to affect the agent, others, and the world for good or ill. To do this job a body needs: (1) sense organs with an enormous variety of possible states varying with an enormous variety of different inputs caused by different distant world states; (2) an information processor that can turn the states of sense organs into brain states that give rise to beliefs of moral or prudential importance; (3) a memory bank, to file states correlated with past experiences (we could not consciously reason about anything unless we could recall our past experiences and what others have told us); (4) brain states that give rise to desires, good and evil (desires to eat and drink, to care for others or to hurt them; and to discover whether or not there is a God); (5) brain states caused by many different purposes that we have; (6) a processor to turn these states into limb and other voluntary movements (to turn, for example, my purpose of telling you that today is Friday into those twists of tongue and lip that will produce an English sentence with that meaning); and (7) brain states that are not fully determined by other physical states. (As far as physical laws are concerned, there needs to be a certain amount of indeterminism in the brain if undetermined human choices are to determine what happens in the brain.)

Clearly human bodies have characteristics (1) to (6). Fairly clearly too there is a small amount of indeterminism in the brain, for, if the laws of Quantum Theory that govern matter on the smaller scale have no deeper deterministic explanation (as most physicists claim), then the behaviour of objects on the small scale is not fully determined. The laws of Quantum Theory merely tell us the physical probabilities of various outcomes. In general, indeterministic behaviour on the small scale averages out to produce virtually deterministic behaviour on the large scale. If each coin had a physical probability of $1/2$ of landing heads and $1/2$ of landing tails, there will be a very large probability close to 1 that the number of coins landing heads in 1,000 tosses will not differ very much from 500. So, even if there is a significant probability that individual atoms will behave in ways different from the norm, bricks and billiard balls are most unlikely to do so. But it is possible to have devices that multiply small-scale indeterminacies, so that a small variation in the behaviour of one atom can have a large-scale effect. The brain is an

extremely complicated system in which small differences cause large differences. But we do not yet know enough about the brain to know whether very small differences of the kind that Quantum Theory seems to tell us are not physically determined are such as to ensure that it is not physically determined which actions humans do. It is, however, evident that the brain is a physical system quite unlike any other physical system, in that it causes conscious events and its states are caused by conscious events; and so clearly laws of a very different kind govern the brain from those that govern all other physical states (as I argue more fully in Chapter 9). So it may well be that there is brain indeterminacy, sufficient for human free choices to produce physical effects, because of some feature of the brain other than the multiplication of indeterminacies within the Quantum limit. I conclude that, on present evidence, there is no good reason to suppose that the brain does not have characteristic (7).

In that case, human intentions will cause human behaviour without being caused to do so by physical events. Humans will therefore have libertarian free will unless something non-physical causes them to form the intentions they do. The most plausible such something is God. But, if I am right in my claim in Chapter 11 that a perfectly good God will allow humans to suffer at the hands of other humans to the extent to which they do only if they have libertarian free will, God will not form their intentions for them. So, if there is a God, humans will have libertarian free will, and so humans will be humanly free agents. So, whether or not there is a God, there seems no reason at present to deny that humans have free will. Yet it seems to each of us at the moment of choice that we are making our choice independently of the forces acting upon us (if we yield to some desire, we are choosing to yield to that desire; and, if we resist the desire, we are choosing to resist it), and we are justified in so believing in the absence of counter evidence.⁷ (This follows from the Principle of Credulity that I discuss in Chapter 13.) Hence (in the absence of some new evidence from neurophysiology) humans are probably humanly free agents, and I shall so assume in future.

We know that human bodies have evolved by natural processes from inorganic matter. But clearly the evolution can have taken place only given certain special physical laws. These are, first, the

⁷ For fuller argument that humans do have libertarian free will that makes a difference to their behaviour in the physical world, see my *The Evolution of the Soul* (2nd edn., Clarendon Press, 1997), ch. 13.

chemical laws stating how under certain circumstances inorganic molecules combine to make organic ones, and organic ones combine to make organisms. And, secondly, there are the biological laws of evolution stating how complex organisms evolve from simpler organisms. I have no wish to challenge Darwin's account of how this happens. Organisms have many offspring, some of which differ from their parents in small ways in respect of one or more characteristics—some offspring are a little taller, some a little shorter, some fatter, some thinner than their parents, some have a small extra growth, some have lost a small part and so on. The new characteristics in turn are passed on to the offspring's own offspring, which in turn also exhibit some variations of characteristics from those of their parents. Given predators and shortage of food, there will be competition for survival, and those organisms whose characteristics give them an advantage in the struggle for survival will survive. Among organisms very well fitted for survival (should they evolve) will be organisms who can see how their environment is changing in crucial ways (by the presence of predators, absence of prey and other food, etc.), and work out (using the criteria of what is evidence for what in more sophisticated ways than non-human animals) in the light of past experience how to survive and to help their young to survive. These organisms will have human bodies of the kind described. It was shown by work subsequent to Darwin that the main mechanism by which these small variations are caused are 'mutations' in genes on chromosomes; there is no regular pattern in mutations—they may occur at any time on any gene affecting any characteristic.⁸

So the question arises why the inorganic matter of which the Earth was made was of a kind that under the operation of laws of chemistry and biology could be converted into human bodies. As we have noted, physics tells us that there was a Big Bang some fifteen billion years ago, which produced matter-energy that condensed into the fundamental particles that came together to form the chemical elements that eventually condensed to form the inorganic matter at the beginning of the history of planet Earth. But why were there laws of physics that brought this about, and the laws of chemistry and

⁸ See Additional Note 2 for recent challenges to this Darwinian account, which—it is claimed—are evidence for various steps of the evolutionary process requiring the intervention of a 'designer'.

biology that led to the inorganic matter being formed into human bodies? Presumably because these laws followed from the fundamental laws of physics. So our question becomes—why are there not just any laws of nature, but laws of a particular kind such that together with the initial matter-energy at the time of the ‘Big Bang’ would lead to the evolution of human bodies. That there are the laws of nature that there are, and the boundary conditions of the universe were as they were, is again where scientific explanation starts: it is something ‘too big’ for science itself to explain. I shall argue that the laws and initial conditions being such as to lead to the evolution of human bodies is very improbable *a priori*, but fairly probable if there is a God who brought it about, and so we have a further substantial C-inductive argument for the existence of God.

Fine-Tuning

Not all initial conditions or laws of nature would lead to, or even permit, the existence of human bodies at some place or other at some time or other in the universe. So we may say that the universe is ‘tuned’ for the evolution of human bodies if the laws and initial conditions allow this to occur (in the sense that they fully cause this evolution if the laws are deterministic, or make it significantly probable if the laws are probabilistic). If only a very narrow range of laws and initial conditions allow such evolution, then we may say that the universe is ‘fine-tuned’ for this evolution. If the fundamental laws and initial conditions are, as we suppose, the laws of Quantum Theory and Relativity Theory with the four forces (strong force, weak force, electromagnetic force, and gravity) governing the basic array of fundamental particles (photons; leptons, including electrons; mesons; and baryons, including protons and neutrons) (what I shall call the standard theory), and the initial conditions are such conditions as the velocity, density, and degree of isotropy of the matter-energy of the universe immediately after the time of the Big Bang; and these are measured in normal ways, then—recent work has shown—the universe is fine-tuned. The constants of its laws and the variables of its initial conditions needed to lie within very narrow ranges if human bodies were ever to exist. One such set of narrow ranges are those centred on the actual values (as we believe them to be) of the constants of laws and variables of initial conditions. It is

worthwhile in this context to bring out how this happens,⁹ why any one constant or variable lying outside the range (while the others lie within the range) would prevent the evolution of human bodies. This section may not be fully comprehensible to those without some scientific background. I suggest, nevertheless, that such readers read through these pages; they will get the main message.

A life based on carbon, in combination with certain other elements, especially hydrogen, oxygen, and nitrogen, is well suited for the formation of bodies with the seven features listed above. With a valence of 4, carbon can enter into many different chemical combinations. Carbon compounds are stable over long periods of time; but are also metastable in that in certain particular situations they can easily be induced to interact with other compounds to produce new compounds. Hence, 'more information can be stored in carbon compounds than in those of any other elements.'¹⁰ Together with hydrogen, nitrogen, and oxygen, carbon can form long, complex chain molecules; and, together with calcium giving skeletal rigidity, such an information-processing system can be made a continuing independent component of the universe. Carbon-based life requires for its stability a moderate range of temperature and pressure, and—if the purposes of organisms are to make much difference to things—a solid planet on which to live.

Given the standard theory with constants and variables of initial conditions having their actual values, it is highly doubtful whether there could be any other kind of intelligent life. It has sometimes been suggested that silicon could replace carbon in its central role, but this seems doubtful in that silicon compounds do not have the stability of carbon compounds.¹¹ Another recent suggestion has been that intelligent systems of particles relying on the 'strong' interaction for their organization might exist inside neutron stars; but it seems doubtful whether they could have nearly as much information-processing capacity as does carbon-based life on Earth.¹² So let us

⁹ The original classical physical analysis of the extent of fine-tuning in the universe is J. D. Barrow and F. J. Tipler, *The Anthropic Cosmological Principle* (Clarendon Press, 1986). This has been carefully reanalysed and updated in Robin Collins, 'Evidence for Fine-Tuning', in N. A. Manson (ed.), *God and Design* (Routledge, 2003). I am much indebted to the latter article for its presentation of the latest state of relevant physics.

¹⁰ Barrow and Tipler, *Anthropic Cosmological Principle*, 547.

¹¹ *Ibid.* 545–8.

¹² Barrow and Tipler, *Anthropic Cosmological Principle*, 343–6.

suppose, plausibly enough, that carbon-based life is the only possible kind of life (given standard theory and the actual values of its constants and variables of initial conditions). If silicon-based life is possible, the argument below would not need much alteration (for the conditions necessary for the evolution of silicon-based life are very similar to those necessary for the evolution of carbon-based life), and neutron-star life is too speculative a suggestion to be taken into account. Given the four fundamental forces and the basic array of fundamental particles, the strengths of the forces and the masses of particles have to have ratios to each other within certain narrow bands if the larger chemical elements, including carbon and oxygen (needed for carbon-based life), are to occur at all; and the Pauli exclusion principle has to hold. This principle (applying to all fermions—for example, electrons and protons) says that in any one system (for example, one atom) only one particle of the same kind can be in a given quantum state. In consequence there are only a small number of possible energy states for the electrons of an atom, and only a small number of electrons can be in each energy state. While the basic laws of Quantum Theory ensure the stability of the atom—electrons do not collapse onto the nucleus—the Pauli principle leads to the electrons being arranged in ‘shells’. Hence atoms of a finite number of different kinds can be formed by different numbers of electrons surrounding the nucleus, and molecules can be formed by bonds between the electrons of different atoms. No exclusion principle, no chemistry. But not much chemistry unless there is plenty of possibility for different structures to be built up, to be relatively stable, to interact, and to form new structures. For that we need atoms to be large structures with plenty of empty space between well-defined central nuclei and electrons.

The build-up of the atoms required for carbon-based life requires the four forces to have certain strengths, relative to each other. If there are to be stable nuclei, the strong force that keeps the protons and neutrons together in the nucleus has to be strong enough to overcome the electromagnetic repulsion between the protons. A 50 per cent decrease in the strong force ‘would undercut the stability of all elements essential for carbon-based life, with a slightly larger decrease eliminating all elements except hydrogen.’¹³ But the process by which carbon and oxygen are built up from the actual initial

¹³ Collins, ‘Evidence for Fine-Tuning’, 183.

conditions of the universe requires a far greater degree of this tuning to lead to their production. An increase or decrease of more than 0.5 per cent in the strength of the strong force or more than 4 per cent in the strength of the electromagnetic force would lead to such small amounts of carbon and oxygen being produced as to make the production of intelligent life very unlikely.¹⁴ A thirty-fold decrease in the weak force would lead to stars being made almost entirely of helium and so having a short life (of about 300 million years) in no way conducive to the evolution of intelligent life.¹⁵ An increase in the strength of the gravitational force by 3,000 would lead to stars with lives of no more than a billion years (compared to the ten billion years of our sun's lifetime), which would make the development of intelligent life much less probable.¹⁶ The weakest of the four forces is the gravitational force, whose effects are significant only over large distances. Over the small distances when the strong force has significant influence, its strength is of the order of 10^{40} times that of the force of gravity. It follows that the kinds of increases or decreases in the strengths of the forces mentioned above (50 per cent, 4 per cent, etc.) compatible with the production of carbon-based life represent a very small range indeed of the values of the strengths of the forces involved within the range of actual values of any of the forces, and an infinitesimal range within the range of logically possible values of the forces. For example, G has to lie between 0 and $3000G$, which represents one part in 10^{36} of the range of values of the force constants. And so on for the other constants.¹⁷ The expansion of the universe is governed by the strength of the initial Big Bang, and the restraining effect of gravity possibly diminished or increased by the value (positive or negative) of the cosmological constant (Λ), which latter may be regarded as determining a fifth force. This needs to lie extremely close to zero if space is not to expand so rapidly that every object in the universe flies apart, or to collapse so rapidly that every object in the universe is crushed.¹⁸

Further, given the actual laws of nature or laws at all similar thereto, boundary conditions will have to lie within a narrow range of the present conditions if intelligent life is to evolve (or else they will have to lie well outside that range; this point will be discussed later). If the

¹⁴ Ibid. 183–6. ¹⁵ Ibid. 188–9.

¹⁶ Ibid. 189–90, 192–4.

¹⁷ Ibid. 190. ¹⁸ Ibid. 180–2.

universe had a beginning, the boundary conditions are the arrangements and properties of the matter-energy of the universe at the time that the universe started off. Present evidence suggests, as I have written earlier, that the universe began in a very dense state some fifteen billion years ago. For the formation of intelligent life in a universe expanding from such a state, conditions at the time of the Big Bang have to be (within the narrow ranges) just right. The initial rate of expansion is critical. If (for the actual value of the gravitational and cosmological constants) the initial velocity of expansion had been slightly greater than the actual initial velocity, the effect would have been the same as would result from a significant positive cosmological constant—stars and so the heavier elements would not form; if it had been slightly less, the effect would have been the same as would have been produced by a significant negative cosmological constant—the universe would collapse before it was cool enough for the elements to form.¹⁹ It has been calculated that (barring a possible qualification from ‘inflation theory’ to which we shall come shortly) a reduction in the rate of expansion of one part in a million would lead to premature collapse, and an increase by one part in a million would have prevented the evolution of stars and heavier elements.²⁰ Some initial inhomogeneity in the distribution of matter-energy is needed if galaxies, and so stars, are to be produced; too much would lead to black holes being formed before stars could form.²¹ In the beginning there was a slight excess of baryons over anti-baryons; all but the excess baryons became matter-energy. If the excess number had been even slighter, there would not have been enough matter for galaxies or stars to form. If it had been much greater, there would have been too much radiation for planets to form.²² And so on. The universe has to start with the right density and amount of inhomogeneity of radiation and velocity of expansion, and that means (within a very narrow range) the actual amount.

I have been setting out the generally agreed point that, if any one of the constants of the laws and variables of the initial conditions were to lie outside a narrow range, human bodies would not evolve.

¹⁹ Barrow and Tipler, *Anthropic Cosmological Principle*, 410–12.

²⁰ Papers by S. W. Hawking and by R. H. Dicke and P. J. E. Peebles cited in J. Leslie, *Universes* (Routledge, 1989), 29.

²¹ Barrow and Tipler, *Anthropic Cosmological Principle*, 414–19.

²² *Ibid.* 401–8.

Recent work has suggested²³ that, if a number of the constants and variables were all significantly different, each having a value within a different small range, human bodies could still evolve. That is, there are several small islands within the space of possible values of constants and variables within which human life could evolve. But this does not alter significantly the point that such islands are the exceptions and the tuning needs to be fine-tuning for this evolution.

If standard theory provides the ultimate explanation of the universe (and so God does not bring it about that standard theory operates), such fine-tuning is *a priori* very improbable. For the form in which any theory, including the standard theory, is stated by scientists in their books and articles is its simplest form—scientists do not try to complicate things for themselves and their readers unnecessarily. This form involves variables and constants being measured in the normal way. It is the form in which we judge the simplicity of the theory that determines (for theories of equal scope) the intrinsic probability of its truth. Versions of standard theory expressed in its simplest form will differ only in respect of the values of constants of laws and of variables of boundary conditions therein. Given all that, a version that claims that a constant or variable lies within one range will not differ greatly²⁴ in simplicity from theories that claim that it lies within another range of equal size; and so each such version will be approximately equally probable *a priori*. But since only a few versions of standard theory in which constants vary over a very small range are tuned for the evolution of human bodies, such evolution is *a priori* very improbable. In slightly more technical terms, the claim is that the probability density for constants and variables measured in the normal way is roughly constant (that is, the probability that these will lie close to a given

²³ For one example, see Collins, 'Evidence for Fine-Tuning', 185.

²⁴ It will differ a little if the simplest formulation of the theory yields a unique zero point for measurement of some variable or constant (a unique point at which some quantity has its lowest value), as, for example, does the Kelvin scale for temperature measurement (0°K being that temperature at which an ideal gas would exert no pressure, and no lower temperature is possible). For then it will be a non-arbitrary matter whether the value of the constant or variable lies within a lower range or a higher range of possible values. It will be a bit more probable that it lies within the former, for the reason that laws containing small integers are simpler than ones containing larger integers (see p. 54).

value is roughly constant for all values of the constants and variables of standard theory).²⁵

It is worth noting the effect of not choosing the simplest formulation of a theory, on the probability density of different constants and variables. I take a very easy example. Newton's law of gravitational attraction $F = G \frac{mm'}{r^2}$ could be expressed as $F = \frac{mm'}{d^3 r^2}$ where d is defined as $G^{-1/3}$. A constant probability density distribution for d (that is, the assumption that it is equally probable that d lies within any range of given size) will not yield a constant probability density distribution for G , and conversely. A constant probability distribution for d will yield the result that d is equally likely to lie between 1 and 0.5 as between 0.5 and 0, and so that G is equally likely to lie between 1 and 8 as between 8 and infinity (i.e. to have any value whatsoever above 8). Expressing the laws of our standard theory in very complicated forms, logically equivalent to their simplest forms, and assuming a constant probability density for the constants and variables of these forms, could have the consequence that much greater variation of these (far less 'fine-tuning') would be compatible with the universe being hospitable to human bodies. But laws are judged simpler and so to have greater prior probability in virtue of the features of their simplest forms. Since a constant is simpler than a constant to the power ($-1/3$), the traditional form of Newton's law is the simplest and so most fundamental form. And, more generally, insistence on the simplest form of a law should yield a unique probability density distribution for the constants and variables of laws of that kind (or, at most, if there are a number of equally simple forms of a law, a few different probability density distributions that

²⁵ The constants and variables of standard theory with which we are concerned do in general have unique zero points (see n. 24 above). In measuring the density of matter-energy, or the velocity of relative recession of the galaxies, for example, velocity and density have unique zero points on the simplest way of measuring them. Hence lower values of these are somewhat more probable than higher values. This has the consequence that, although there is an infinite range of possible values of these constants and variables, there can be a finite probability that some such constant will have a value lying within any given range. But, if the constant or variable having a value within a range of given size was the same throughout the infinite range (as would be the case for constants and variables without a unique zero point), the probability of it lying within any finite range would be infinitesimal. (For the need to use infinitesimals in assessing probabilities, see my *Epistemic Justification* (Clarendon Press, 2001), Additional Note G). So, either by ascribing higher intrinsic probabilities to lower values of the constants and variables, or by using infinitesimals, I avoid what is known as the 'normalizability problem'. (See e.g. Timothy McGrew *et al.*, 'Probabilities and the Fine-Tuning Argument', in N. A. Manson, *God and Design*.)

are not likely to make much difference to the extent of need for fine-tuning). So, given the standard theory, and no more fundamental explanation thereof (physical or theistic), tuning is *a priori* immensely improbable.

Physical cosmology is a very unstable branch of physics. New theories are produced each year. Changes are possible that would have the consequence that constants and variables can vary within a much wider range and yet life still evolve. One possible change, though in my amateur judgement a fairly unlikely one, is that it may become established that the boundary conditions are significantly different from what has been supposed—for example, that the universe evolved not from an initial singularity, but from a very dense state resulting, perhaps, from a prior collapse, or perhaps from a quantum mechanical fluctuation of the ‘vacuum’.²⁶ Such a change, probably going with the adoption of the view that the universe was infinitely old, would have the consequence that a far wider range of boundary conditions would give rise to life.

The role of ‘boundary conditions’ in a backwardly eternal (that is, infinitely old) universe may need clarification. Imagine a billiard table sealed under a glass cover in which the balls move in a vacuum (and that any energy transfer to or from the outside can be discounted). The laws of collision govern the interaction of the balls, which bounce off each other and off the walls for the indefinite future. It could have been that this process was started off by someone arranging the balls and giving them an initial push before the table was sealed. In that case the boundary conditions would be the initial conditions (arrangement and velocity of balls), and they together with the laws of collision would determine all the subsequent behaviour of the balls. Some initial conditions would allow the balls to arrange themselves in all the (logically) possible arrangements during the course of a subsequent infinite time. Yet some initial conditions (for example, the balls moving initially with velocities parallel to each other and to the walls) will ensure that the balls occupy only a few of the possible arrangements even in the course of infinite time. Suppose, now, that the process has been going on for ever (that is, is not merely forwardly but backwardly eternal). Then the set-up may still have certain features at some given time that would occur only if a narrow set of possible arrangements either ever have been or ever will

²⁶ Barrow and Tipler, *Anthropic Cosmological Principle*, 440–1.

be occupied (for example, this might also be the case if at a certain time the balls are moving parallel to each other and to the walls); or, much more likely, features that would occur only if in the course of infinite time backward and forward all possible arrangements of those balls occur. However, the sealing of the table still ensures that the only possible arrangements are arrangements of those balls—there cannot be more or fewer balls in the past or future. The ‘boundary conditions’ of an infinite universe are those features of its condition at any one time (for example, in a Newtonian universe, the quantity of energy) that (together with the laws governing it) restrict its possible future and past states.

Now, if the universe is backwardly and forwardly eternal, its present state may be such that we can infer that it must pass through such and such a range of states in the course of infinite time. These might include all the logically possible states of matter-energy; but that is not very likely, for some kind of principle of conservation of energy (within quantum limits) will ensure that past (and future) states are limited to rearrangements of the existing amount of energy. However, although all this would have to be worked out, it is highly plausible to suppose that (for given scientific laws) life is much more likely to evolve at some time in the course of the history of our universe if it has an infinite past than if it has only a finite past. There is more time for more possible arrangements of the constituents of the universe. Nevertheless, the present evidence suggests a finite age of some fifteen billion years.

An alternative change in physics might be a discovery that the laws are other than previously supposed, again in such a way that they bring forth intelligent life out of a much wider range of boundary conditions than had hitherto been supposed. ‘Inflation theory’ suggests just that. Inflation theory tells us that regions of the universe with certain features may have been subject, soon after the Big Bang, to a vast faster-than-light expansion, leading to them very quickly becoming cool homogeneous and isotropic regions.²⁷ So features such as homogeneity and isotropy for which a narrow range of initial conditions were thought vital are—according to inflation theory—to be expected, given certain laws, to arise from a wider range of initial conditions. Yet it may well be that inflation theory can be successful in any of its many variants in removing the need for fine-tuning from

²⁷ Barrow and Tipler, *Anthropic Cosmological Principle*, 430–40.

the initial conditions only by putting more fine-tuning into the laws.²⁸

There remains, however, a consensus among physicists that the values of the constants in the laws of standard theory (as opposed to the variables of initial conditions) must lie within very narrow ranges if life is to evolve anywhere in the universe—ranges that include the actual values of the constants and probably a few other small ranges in which the values of several of the constants are different from their actual ones. And there is also a consensus that, given an initial Big Bang, variables such as the initial velocity of recession have (even on inflation theory) to lie within a narrow range. There may, however, be a more fundamental physical theory that explains the standard theory, and a constant probability density for the constants and variables of the boundary conditions of the simplest form of that fundamental theory may have somewhat different consequences for the intrinsic probability of tuning (for example, that the more readily observable variables can take only certain values).²⁹ But, more generally, there are innumerable possible scientific theories differing in their form from each other, and innumerable different kinds of boundary conditions differing in the number of entities that they postulate (big and small universes), each allowing many different sets of constants and boundary condition variables. A constant probability density over the latter (when each theory is expressed in its simplest form) will yield for each theory a different probability that a universe conforming to it will be tuned. The theories (although of equal scope—telling us about everything) themselves will differ in their simplicity, and so in their intrinsic probability. Hence, given a precise way of measuring simplicity, there will be a true value for the intrinsic probability, the probability if there is no God, that any

²⁸ See J. Earman and J. Mostevin, 'A Critical Look at Inflationary Cosmology', *Philosophy of Science*, 66 (1999), 1–49.

²⁹ It is possible that the derivation of the fundamental laws of nature from string theory would greatly reduce the need for fine-tuning. This has been argued in G. L. Kane *et al.*, 'The Beginning and End of the Anthropic Principle', astro-ph/0001197. They suggest that all string theories are equivalent; and that different possible 'vacua' uniquely determine all the constants and initial values of variables of laws of nature. They acknowledge that much work needs to be done before (if ever) string theory is established and their result can be demonstrated. But, even granted all this tentative speculation, they acknowledge that 'there will be a large number of possible vacua'; and that means both having string theory rather than many other alternative fundamental laws and requiring special variables of initial conditions if human bodies are to evolve.

universe will be tuned. It will be (loosely) that proportion of logically possible universes that are conducive to the evolution of human bodies, each weighted by the simplicity of the laws that govern it and the fewness and simplicity of entities in its boundary conditions. And, given the rough way we have of measuring simplicity, we could still give a rough value to this. So it does not matter—for the purposes of an argument from fine-tuning—whether we have the correct theory of our universe, or whether there is a more ultimate physical explanation of the forces that govern it; and whether only a small proportion of the versions of the correct theory lead to a tuned universe. For the prior probability (in a Godless universe), that a universe will be tuned is a function not of the true physical theory and actual kinds of boundary conditions that govern our universe, but of all the possible theories and boundary conditions there could be for any universe at all. It is not, however, within my ability to calculate this value, nor—I suspect—within the ability of any present-day mathematician.

What, however, I suggest, is fairly obvious is that *no* relatively simple universe would be tuned. For consider the seven features required by a human body listed earlier in this chapter. Such a body has parts. But the parts have to form one body distinct from other bodies and from the inanimate world. In our world this is secured by a chemistry whereby only some bits of matter link to other bits of matter—if I put my hand into a sandpit, my hand will not absorb the sand; but, if I eat some bread, it will become part of my body. Sense organs require an enormous variety of stimuli impinging on a place, which vary with their distant source. In our universe the best of all such stimuli are light waves—an enormous variety of different light waves arrive every second at our eyes, which vary with the states of objects many metres away. The sense organs respond differently to each very small range of incoming stimuli. But we humans are interested only in certain aspects of the states of distant objects—whether they are the bodies of predators, or prey, or mates, and so on for a million possible differences. The stimuli have to cause brain states that give us information of moral or prudential importance. Our information processor will utilize states caused by past experiences to turn the states of sense organs into useful brain states. And, if we are not to be just automata but to reason consciously from past experiences, we need a memory bank to file those states in recoverable form. This requires a chemistry of stable

states (so that memories remain the same as time passes) and metastable states so that certain kinds of input will move a brain element from one state to another (as we learn that some previous belief was erroneous). And for output we need again an enormous variety of brain states corresponding to the different purposes we could form, a processor to turn these into the relevant limb movements (for example, if I want to tell you that today is Friday, to produce the twists of tongue and lip that will cause the appropriate sounds of the English language). And we need a stable inorganic world to which we can make a difference that remains; there is no point in trying to build a house if the bricks immediately liquify.

One way in which all this could be achieved would be by bodies composed of only a few particles, each capable of existing in a trillion trillion trillion different states. But a physics that allowed such particles would be of incredible complexity. The other way, the way operative in our universe, is to have extended bodies, each composed of many fundamental particles of a number of different kinds, each particle capable of existing in a few different discrete states; the differences between bodies being a matter of the number and arrangement of the units and the discrete states of each. To do the job in this way you need a universe with a very large number of particles to compose many bodies and an inanimate environment through which people may influence each other. Change has to be affected through a particle (or group of particles) changing their states, causing other particles to change their states. To secure stable bodies that are nevertheless capable of existing in many different states, you need more than one simple force. One simple force of attraction would lead to crushed lumps of matter incapable of sensitive reaction; and one simple force of repulsion would lead to there being no extended bodies at all. Minimally a combination of two different simple forces (possibly both derivable from one more complicated force) is required. A force of attraction between particles inversely proportional to the square of the distance apart of the particles would be required to be balanced, for example, by a force of repulsion inversely proportional to the cube of their distance apart. Forces of these kinds of the right strength would lead to particles coming together but not collapsing on top of each other. But to preserve states (of the brain correlates of belief, for example) intact, we have to rule out small variations. We need metastability—systems that remain unchanged under forces of a certain strength but that change

from one discrete state to another discrete state when the strength of the force exceeds a certain amount. This is ensured in our universe by the laws of Quantum Theory, which guarantee the stability of the atom. And, to have distinct bodies that do not merge with each other, and distinct brain states that are open to change only under certain kinds of input, we need something like a chemistry allowing substances to combine easily with some substances but not with other substances. This is secured in our universe by chemical substances different from each other by the charge on their nucleus and the arrangements of charge-balancing electrons in shells around the nucleus—in other words, protons, neutrons, and the Pauli principle. And so on.

So we need large numbers of particles of a few different kinds and forces of some complexity acting between them. But universes are simpler, the fewer objects (for example, particles) they contain and the fewer kinds of mathematically simple forces that operate between them. No very simple universe could be tuned, whatever its boundary conditions. Clearly more complicated kinds of possible universes (for example, ours) can be tuned, and maybe normally the tuning needs to be fine-tuning. Maybe, too, some very complicated kinds of universe would produce human bodies for most values of constants and variables of boundary conditions. But the considerable *a priori* weight of simplicity suggests that in a Godless universe it is *a priori* improbable that any one universe will be tuned so as to yield human bodies.³⁰ With *e* as the existence of human bodies, *h* as theism, and *k* as the evidence of a universe conforming to natural laws, $P(e | \sim h \& k)$ is very low.

³⁰ In order to show the improbability of tuning, it is not enough to show that tuning is improbable given the standard theory—that is, given the local area of possible worlds. John Leslie has compared this fine-tuning to a dart hitting a cherry on a wall when there are no other cherries in that area of the wall. He claims that (on the assumption that hitting a cherry is something a dart-thrower might wish to do) the fact that the dart hit the cherry is evidence that it was thrown intentionally by a dart-thrower—even though there are many cherries on other areas of the wall. (See p. 143 of his 'Anthropic Principle, World Ensemble, Design', *American Philosophical Quarterly*, 19 (1982), 141–52.) His claim seems to depend on a feature of his analogy, to which there is no parallel in the universe fine-tuning case. A dart-thrower would naturally try to hit a cherry in an area of a wall a long way from other cherries, his aim being to hit a cherry when that would be difficult for the average human dart-thrower. Hence he aims at the isolated cherry, rather than at cherries close to other cherries. A God tuning a universe seeks to produce human bodies; he has no particular concern to produce them in a possible world where all close possible worlds (except the very closest ones) would not allow their existence. He has no concern to show his universe-

Of course, if there was an infinite number of universes, each with different laws and different boundary conditions, one might expect at least one to be tuned. (Recall my earlier—see p. 133—definition of a universe as a collection of physical objects, all spatially related to each other. A universe other than our own would be a collection of physical objects spatially related to each other, but not to our Earth.) I have already in this chapter made the point that it is the height of irrationality to postulate an infinite number of universes never causally connected with each other, merely to avoid the hypothesis of theism. Given that simplicity makes for prior probability, and a theory is simpler the fewer entities it postulates, it is far simpler to postulate one God than an infinite number of universes, each differing from each other in accord with a regular formula, uncaused by anything else.³¹ There might, however, be particular features of our universe (other than its tuning) that are most simply explained by supposing that it ‘budded off’ from another universe in consequence

tuning skills, only to bring about an end product. So, if the fact that there is a tuned universe is to be evidence for God being its creator, what has to be shown improbable *a priori* is not that there be a tuned universe in our local area of possible worlds, but that there be a tuned universe among all possible worlds. I have given some argument for this—from the impossibility of any very simple universe (and so any universe intrinsically probable) being tuned.

³¹ Max Tegmark has, however, claimed that it is simpler to postulate an infinite number of universes than to postulate just one. See Max Tegmark, ‘Is “The Theory of Everything” merely the Ultimate Ensemble Theory?’, *Annals of Physics*, 270 (1998), 1–51, at 38. ‘Our TOE [Theory of Everything] . . . postulates that all structures that exist in the mathematical sense exist in the physical sense as well. The elegance of this theory lies in its extreme simplicity, since it contains neither any free parameters nor any arbitrary assumptions about which of all mathematical equations are assumed to be “the real ones”’. He explicitly (his p. 44) assumes an account of simplicity, according to which a theory is simpler the fewer the number of computational symbols needed to express that theory. This ‘algorithmic’ account has the consequence that, for example (p. 44), the ‘set of all perfect fluid solutions to the Einstein field equations has a smaller algorithmic complexity than a generic particular solution, since the former is specified simply by giving a few equations and the latter requires the specification of vast amounts of initial data on some hypersurface’. So it is simplest of all to postulate that every possible universe exists, since that needs very few computational symbols indeed to state!

Tegmark’s account of simplicity seems to me to yield in this case a bizarre result, totally out of line with all our inductive practice. If we are postulating entities to explain phenomena, we postulate the fewest number of entities needed to do the job. If we did adopt Tegmark’s approach, we would need to amend and amplify his theory in two crucial respects. First, we would need to amend it to deal with the problem that the supposition that all possible entities exist is incoherent. For the existence of some entities rules out the existence of others. Thus, the existence of an omnipotent all-good God rules out

of a law whereby universes produce daughter universes differing from them in boundary conditions and laws; and so our universe is explained as one of a collection of an infinite number of universes (originally causally connected with each other) differing from each other in boundary conditions and laws. But that is tantamount to postulating a multiverse that has laws and boundary conditions such that it will contain at some time or other a tuned universe. But then there are an infinite number of logically possible multiverses that do not have this characteristic, and the shape of the problem has in no way changed. For the problem that concerns us is not really why is there one (in my sense) universe that is tuned for life, but why among all the universes there are (one or many) is there a universe tuned for life. One way in which this could come about is by there being only one such universe. But another way is by there being a universe-generating mechanism that produces universes of various kinds, including a universe tuned for life. But, although the existence of this possibility does not change the shape of the problem, it draws our attention to a way in which a universe tuned for life could have come into existence. And so, in order to assess the intrinsic probability that there be a universe tuned for life, we need to assess the probability that that would come about by one or other route. And taking this into account may lead us to reassess the value of that probability.

It might seem that the value would turn out to be much higher than we originally supposed. Let us individuate universe-generating mechanisms by the multiverse (the collection of universes) that they generate (at some time or another). Then, if we consider all the possible multiverses, each consisting of r universes, chosen from n logically possible kinds of universe only one of which is fine-tuned for life, it follows mathematically that a proportion $\frac{r}{n+r-1}$ of these

the existence of an omnipotent Devil (in any actual universe at all). And, given God, he will certainly not choose to bring about the existence of certain other states—for example, endless suffering unchosen by the sufferer. And, secondly, the account of the simplicity of a theory in terms of the fewness of computational symbols needed to express it, which led to the claim that it is simple to suppose that every possible universe exists, needs considerable amplification. For how many symbols you need to express something depends on the language you use. All theories can be expressed in the form ' $a = b$ ', when a and b represent some highly complicated multidimensional tensors. But, of course, it needs a language far removed from the language of observation to express the theory in that way. Tegmark's account of simplicity is not a clear one and its consequence in the case of current interest to us is bizarre and contains a contradiction.

multiverses will contain a universe fine-tuned for life. For any $r > 1$ ($r = 1$ being the case where there is only universe), this will exceed $\frac{1}{n}$ (the proportion of universes fine-tuned for life). And the more universes in a multiverse (the larger is r), the closer this value will be to 1. So it might seem that, as we consider more and more possible universe-generating mechanisms (generating more and more universes) (r getting larger and larger), the total proportion of universe-generating mechanisms that will generate a universe fine-tuned for life will approach 1. So if it were equally probable that there exist any possible universe-generating mechanism (most of them generating far more universes than the number of logically possible kinds of universe), it would seem to be very probable that there would occur at least one universe fine-tuned for life.

However, we cannot calculate the intrinsic probability (in a Godless world) of a universe-generating mechanism being such as to produce a universe tuned for life merely by counting the proportion of mechanisms that have this characteristic among the total number of possible such mechanisms. To start with, there will be an infinite number of possible mechanisms of which an infinite number will have the required feature. And infinity divided by infinity has no definite value. We have to divide up mechanisms into a finite number of kinds of mechanism, and then weight each kind by the prior probability of a mechanism being of that kind—which will be a function of the simplicity of the laws involved in the mechanism. Now, clearly mechanisms that yield universes varying from each other only in the constants involved in their laws will be much simpler than mechanisms that yield universes differing in the kinds of laws they have. A mechanism that produced universes with laws of totally different kinds from each other would need itself to be governed by some very complicated laws. Yet, if we are confined to mechanisms that yield only laws of one kind, my earlier arguments suggest that very few such mechanisms yielding only laws of relatively simple kinds (that is laws no more complex than are those of our universe) will yield a universe fine-tuned for life. Secondly, mechanisms that produce universes with simple laws are simpler and so intrinsically more probable than mechanisms that produce universes with more and more complex laws as well. And, thirdly, the existence of a multiverse with a universe-generating mechanism is a more complex supposition than the existence of one universe without such a mechanism.

So, even if there is a large range of possible multiverses tuned for life (in the sense of producing a universe tuned for life), and the proportion of the range of possible multiverses tuned for life is vastly greater than the proportion of the range of single universes so tuned, this holds only because the former range includes very complex multiverses that are intrinsically very improbable. So I stick by my point that it is intrinsically very improbable that there be a universe tuned for life (whether it is a sole universe, or a universe produced by a universe-generating mechanism). Yet it may well be that this improbability is less than the improbability that a single universe would be tuned for life.

The Probability of Spatial Order, given Theism

A God, however, I argued earlier, has good reason for bringing about embodied humanly free agents, such as human beings appear to be; and so, on the hypothesis of theism, it is moderately probable that the universe will be tuned—that is, such as to allow and indeed make significantly probable the existence of human bodies. God could achieve this either by creating such bodies entire, or by creating and keeping in existence a universe designed to bring them about by a long evolutionary process, or even a multiverse designed to bring about such a universe.

What reason would God have for taking an evolutionary route? If his only aim in creating a universe was to populate it with human beings, there would indeed be no point in producing them by a long evolutionary process. But there are other good features of the universe that God has good reason to bring about. I have commented already on the beauty of the inanimate universe shown in the development of galaxies, stars, and planets. God has every reason for bringing about this process of development from the Big Bang for its beauty—even if he were the only person to observe it. But, of course, he is not the only person to observe it. We can observe it through our telescopes reaching further and further back into the earliest stages of the universe. And God has the same reason for bringing about plants and animals—their beauty. And animals are good also, I have argued, in virtue not only of their beauty but also of their ability to have pleasant sensations and true beliefs and spontaneously to do good actions (even if not ones freely chosen). In view of all this, it is not too

surprising that God should take the long (by our timescale) evolutionary route to produce human bodies. And similar, though weaker arguments would show it to be unsurprising if God produced human bodies by an even longer route of going through more than universe to achieve this goal.

It may be that, even given the initial conditions of the universe in all their detail, the laws of nature as such do not necessitate the evolution of human bodies, only make it quite probable. As I wrote earlier, it may be that the way by which God ensures that human free choices make differences to the world is by bringing it about that the fundamental laws of nature are probabilistic, not fully deterministic. And clearly God can guide the way in which the probabilistic laws operate so as to ensure that human bodies do evolve, without in any way preventing their operation, simply by ensuring that the most probable outcome does occur. Yet there will be an argument from the existence of human (and animal) bodies to the existence of God of any great strength, via the route of 'fine-tuning', only if it follows that a fine-tuned universe will (not merely possibly but with significant probability) lead to embodied humans and animals. For fine-tuning as such is merely a necessary, and not a sufficient, condition for the evolution of humans and animals. There is, however, a very considerable, but not unanimous, scientific view that the laws and initial conditions of our universe make it very probable that human life will evolve in more than one place in the universe, and animal life will evolve in quite a number of places. And that is enough to make the argument a cogent one.

So it is quite likely that, if there is a God, the laws and boundary conditions of the universe will be such as to make probable the evolution of human bodies. It is otherwise very improbable that they will have this feature. I represent this evidence of the nature of the laws and boundary conditions as e , with h as the hypothesis of theism, and k as the background knowledge that formed the evidence of the two arguments considered previously—that there is a universe governed by simple laws of nature. The probability then, if there is no God, that the laws and boundary conditions will be such as to have this further feature of bringing about human bodies is $P(e|\sim h \& k)$. The probability that this will happen if there is a God is $P(e|h \& k)$. I have argued that $P(e|h \& k) \gg P(e|\sim h \& k)$, and so—by Bayes's theorem— $P(h|e \& k) \gg P(h|k)$. We have here a powerful C-inductive argument for the existence of God.

The Argument from Beauty

The strength of the argument from the universe and its spatial and temporal order to God is increased when we take into account the beauty of that universe. As we have noted, the universe is beautiful in the plants, rocks, and rivers, and animal and human bodies on Earth, and also in the swirl of the galaxies and the birth and death of stars. Mark Wynn comments that nature is 'uniformly beautiful whereas the products of human beings are rarely beautiful in the absence of artistic intent'. I argued in Chapter 6 that, if God creates a universe, as a good workman he will create a beautiful universe. On the other hand, if the universe came into existence without being created by God, there is no reason to suppose that it would be a beautiful universe. The argument has force on the assumption with which I am happy and commend to my readers that beauty is an objective matter, that there are truths about what is beautiful and what is not. If this is denied and beauty is regarded as something that we project onto nature or artefacts, then the argument could be rephrased as an argument from human beings having aesthetic sensibilities that allow them to see the universe as beautiful. In the latter case, there is certainly no particular reason why, if the universe originated uncaused, psycho-physical laws (of the kind that I shall consider in the next chapter) would bring about aesthetic sensibilities in human beings. But, good though it is that humans should have these sensibilities, it would need to be shown that it would be involved in the equal best kind of act that constituted the creation of humanly free agents to endow them with aesthetic sensibilities.

For not to do would not deprive the universe of a kind of sensibility, since God could himself have it whereas the ability to make significant choices between good and evil is not a kind of goodness that God himself could have. Because the argument from beauty needs, I suspect, an objectivist understanding of the aesthetic value of the universe, in order to have significant strength, and the establishment of such an understanding would require very considerable argument, I shall omit further discussion for reasons of space.³²

³² An argument to God from the beauty of the world was presented by F. R. Tennant in his *Philosophical Theory*, vol. 2, *The World, the Soul, and God* (Cambridge University Press, 1930). There is a good short presentation of this argument and response to objections to it in Mark Wynn, *God and Goodness* (Routledge, 1999), ch 1. For the quotation from Wynn, see *ibid.* p. 20.

I should add that this point does not undermine the earlier point that the beauty of the physical universe (whether objective, or subjective in its perception by persons) provides a good reason for God to produce human bodies by the evolutionary route; my point here is simply that it needs much further discussion to show that the beauty of the physical universe provides a positive argument of significant strength for the existence of God.

Arguments from Consciousness and Morality

In considering the arguments to the existence of God I am passing from those with more general premisses to those with less. The cosmological argument appeals to the existence of the universe as its starting point. Teleological arguments appeal to the universe having some very general over all characteristic; it having laws of nature, and these laws and boundary conditions being such as to produce human bodies. The subsequent arguments that I shall consider appeal to more specific features of the universe, and in particular to the nature, experiences, and history of the conscious beings that inhabit the universe. The main argument that I shall consider in this chapter is an argument from the nature of human beings. I argued in Chapter 6 that there is a significant probability that a God would create embodied humanly free agents; and I claimed in Chapter 8 that probably humans are embodied humanly free agents. I also argued in Chapter 8 that, if there is no God, it is improbable that the universe would be ‘tuned’, so as to allow the existence of human bodies. But the value of human bodies lies in their being the vehicles for the acquisition of knowledge and the execution of purposes by human beings. The value of humans lies in their conscious life—in their acquiring beliefs, having thoughts and sensations, having desires, and (through free choice) executing purposes. How probable is it that, if there is no God, human bodies would give rise to the conscious life typical of humans? I shall argue that it is very improbable.

The argument from consciousness to God is not one developed by any classical philosopher at any great length. John Locke gives a very brief statement of the view that matter is such a different

thing from 'thought' that matter could never produce 'thought' by its own power.

Divide matter into as many parts as you will . . . vary the figure and motion of it, as much as you please . . . and you may as rationally expect to produce sense, thought, and knowledge by putting together in a certain figure and motion, gross particles of matter, as by those that are the very minutest, that do anywhere exist. They knock, impel, and resist one another just as the greater do, and that is all they can do.¹

Only a thinking being could produce thought, Locke claims; given that, as he claims to have established in another passage, there must be an everlasting supreme cause of things, it follows that that cause must be a 'cogitative being', namely God. But there cannot be said to be much argument in this passage, merely an appeal to the apparent obviousness of 'matter cannot produce thought'. I think that many others have felt that matter cannot produce thought, and that this has formed one vague reason people have for looking for a different cause of 'thought', namely God. I think that this feeling can be put into the form of a powerful argument to which philosophers have not given nearly enough attention. In this chapter, I propose to set out and defend such an argument.

The Mental Data

But first let us delineate carefully the data that need explanation, which Locke calls 'thought', and for this purpose we require the terminology of substances, properties, and events, that I introduced in Chapter 2. An event consists in the instantiation of a property in a substance at a time; and we need now to distinguish between mental and physical properties, substances, and events. Although this distinction may be made in more than one way, I shall find it best, in order to articulate the argument from consciousness, to make it in the following way. I shall understand by a mental property one to whose instantiation the substance in which it is instantiated necessarily has privileged access on all occasions of its instantiation, and by a physical property one to whose instantiation the substance in which it is instantiated does not necessarily have such privileged access.

¹ John Locke, *An Essay concerning Human Understanding* (first published 1690), ed. A. C. Fraser (Dover Publications, 1959), 4.10.10.

Someone has privileged access to whether a property *P* is instantiated in him in the sense that, whatever ways others have of finding this out, it is logically possible that he can use, but he has a further way (by experiencing it) that it is not logically possible that others can use. A pure mental property may then be defined as one whose instantiation does not entail the instantiation of a physical property. So 'trying to raise my arm' is a pure mental property, whereas 'intentionally raising my arm' is not; for the instantiation of the latter entails that my arm rises. A mental event is one that involves (or entails) the instantiation of a mental property; and a pure mental event is one that involves (or entails) the instantiation only of a pure mental property. (Mental properties will include both conscious properties and continuing mental properties. Conscious properties are ones of whose instantiation in a subject that subject is necessarily aware while they are instantiated—for example, having the thought that today is Tuesday. Continuing mental properties are ones to which the exercise of the subject's privileged access depends on her choice to introspect, but that continue to characterize her while she chooses not to ask herself about them—for example, the beliefs we have while asleep or thinking about other things, and the desires we have that are not currently influencing our behaviour.) A physical event is one that involves (or entails) the instantiation of a physical property (and no mental property). My argument to God in this chapter will be an argument from mental events rather than from the narrower class of conscious events.

I shall understand by a physical substance a substance all of whose essential properties are physical properties (and any properties entailed thereby). A property of a substance is an essential property (see Chapter 5) if necessarily the substance would not exist without that property. Thus occupying space is an essential property of my desk; it could not continue to exist and yet occupy no volume of space. Tables and chairs, brains and planets, houses and nerve cells are all physical substances. The only properties that they need to have in order to exist are physical properties. A mental substance is one that has as an essential property at least one mental property. A pure mental substance is one all of whose essential properties are pure mental properties (and any properties entailed thereby). (Such a substance may have, contingently—that is, non-essentially—also physical properties.) I shall be arguing in due course that we, myself and my readers, are pure mental substances.

Evidently—more evidently than anything else—there really are pure mental events, as we know from our own experience. They include patterns of colour in my visual field, pains and thrills, beliefs, thoughts and feelings. They also include the intentions that I try to realize through my body or in some other way, which I discussed in Chapter 2. My being in pain at midday yesterday, or having a red image in my visual field, or thinking about lunch, or forming the intention of going to London are such that, if others could find out about them by some method, I could find out about them by the same method. Others can learn about my pains and thoughts by studying my behaviour and perhaps also by studying my brain. Yet I, too, could study my behaviour—I could watch a film of myself; I could study my brain—via a system of mirrors and microscopes—just as well as anyone else could. But, of course, I have a way of knowing about pains, thoughts, and suchlike other than those available to the best other student of my behaviour or brain: I actually experience them. Consequently, they must be distinct from brain events, or any other bodily events. For having a brain event (the instantiation in a person of some physico-chemical property) does not entail having a pure mental event (having some sensation, thought, or whatever). A neurophysiologist cannot observe the quality of the colour in my visual field, or the pungency of the smell of roast beef that I smell. A Martian who came to earth, captured a human being, and inspected his brain could discover everything that was happening in that brain but would still wonder ‘Does this human really feel anything when I stamp on his toe?’ It is a further fact beyond the occurrence of brain events that there are pains and after-images, thoughts, and intentions. Likewise, such events are to be distinguished from the behaviour to which they typically give rise. People have sensations to which they give no expression—pains that they conceal or dream sensations that they report to no one—and, if the sensations give rise to behaviour, the subject is aware of the sensation as a separate event from the behaviour to which it gives rise.

I emphasize my definition of the mental as that to which the subject has privileged access. There are many properties that we attribute to people that we might sometimes call ‘mental’ but that are not mental in my sense but are merely properties of public behaviour. When we say that someone is generous or irritable or a useful source of information, it may be that we are just saying

something about the way they behave in public, not anything about the life of thought and feeling that lies behind such behaviour. We may naturally describe being irritable as a 'mental' property, but it is not a mental property in my defined sense. My concern is to make the point that there are mental events in my sense, distinct from brain events.

In making this point, I do not, of course, deny that most of my mental events are caused by my brain events. Clearly most of the passive mental events—the ones that we find ourselves having, sensations, thoughts, beliefs, and desires—are caused at least in part by brain events, themselves often caused by further bodily events; while some mental events are caused, at least in part, by other mental events. My toothache is caused by a brain event caused by tooth decay. A thought that represents the conclusion of a deductive inference is caused (at least in part) by other thoughts encapsulating the premisses of that inference. And, of course, as we considered in Chapter 2, there is also causation in the other direction. Active mental events—our intentions (that is, purposes)—themselves cause brain events that in turn cause further bodily events. Our embodiment consists in there being these mental–physical connections.

A human would not exist unless it had a capacity for a mental life (a capacity to have sensations, thoughts, etc.); and having such a capacity is itself a mental property (one to the instantiation of which in a subject he has privileged access). Hence humans are mental substances. But there is more to humans than just having essentially a capacity for a mental life, connected to a body. That mental life itself, I now argue, is a state of the mental substance that is the embodied human being in virtue of being a state of a pure mental substance, the soul of the human, which is connected to his body. For what makes me me is the continuity of my mental life, not the continuity of a body to which it is connected. Even if normally the latter continuity is physically necessary for the former, there are two different continuities. And by the continuity of my mental life, I mean simply that the mental events are had by me, a notion which is not further analyzable; but whose non-physical nature we can bring out by giving the name of my 'soul' to the essential part of me which has the mental events (and which is connected to the non-essential part, my body). We can begin to see that this account of personal identity is correct by pointing out that if you knew all the

properties, physical and mental associated with bodies, you would still not know one of the most important things of all—whether you or any other human continued over time to live a conscious life.

Let me illustrate this with the example of brain transplants. The brain consists of two hemispheres and a brainstem. There is good evidence that humans can survive and behave as conscious beings if much of one hemisphere is destroyed. Now imagine my brain (hemispheres plus brainstem) divided into two, and each half-brain taken out of my skull and transplanted into the empty skull of a body from which a brain has just been removed; and there to be added to each half-brain from some other brain (for example, the brain of my identical twin) whatever other parts (for example, more brainstem) are necessary in order for the transplant to take and for there to be two living persons with lives of conscious experiences. Now I am very well aware that an operation of this delicacy is not at present practically possible and perhaps never will be possible for mere human scientists with mere human resources; but I cannot see that there are any insuperable theoretical difficulties standing in the way of such an operation. (Indeed that is a mild understatement—I fully expect it to be done one day.) We are, therefore, entitled to ask the further question—if this operation were done and we then had two living persons, both with lives of conscious experiences, which would be me? Probably both would to some extent behave like me and claim to be me and to remember having done what I did; for behaviour and speech depend, in large part, on brain states, and there are very considerable overlaps between the ‘information’ carried by the two hemispheres that gives rise to behaviour and speech. But both persons would not be me. For, if they were both identical with me, they would be the same person as each other (if *a* is the same as *b*, and *b* is the same as *c*, then *a* is the same as *c*) and they are not. They now have different experiences and lead different lives. There remain three other possibilities: that the person with my right half-brain is me, or that the person with my left half-brain is me, or that neither is me. But we cannot be certain which holds. It follows that mere knowledge of what happens to brains or bodies or anything else physical does not tell you what happens to persons.

It is tempting to say that it is a matter of arbitrary definition which of the three possibilities is correct. But this temptation must be

resisted. There is a crucial factual issue here—which can be shown if we imagine that I have been captured by a mad surgeon who is about to perform the split-brain operation on me. He tells me (and I have every reason to believe him) that the person to be formed from my left half-brain is to have an enjoyable life and the person to be formed from my right half-brain is to be subjected to a life of torture. Whether my future life will be happy or painful, or whether I shall survive an operation at all, are clearly factual questions. (Only someone under the grip of some very strong philosophical dogma would deny that.) Yet, as I await the transplant and know exactly what will happen to my brain, I am in no position to know the answer to the question—what will happen to me. Maybe neither future person will be me—it may be that cutting the brainstem will destroy the original person once and for all, and that, although repairing the severed stem will create two new persons, neither of them will be me. Perhaps I will be the left-half-brain person, or maybe it will be the right-half-brain person who will be me. Even if one subsequent person resembles the earlier me more in character and memory claims than does the other, that one may not be me. Maybe I will survive the operation but be changed in character and have lost much of my memory as a result of it, in consequence of which the other subsequent person will resemble the earlier me more in his public behaviour than I will.

Reflection on this thought experiment shows that, however much we know about what has happened to my brain—we may know exactly what has happened to every atom in it—and to every other physical part of me, we do not necessarily know what has happened to me. From that it follows that there must be more to me than the matter of which my body and brain are made, a further essential non-physical part whose continuing in existence makes the brain (and so body) to which it is connected my brain (and body), and to this something I give the traditional name of ‘soul’.

Take a slightly different example. I die of a brain haemorrhage that today’s doctors cannot cure, but my relatives take my corpse and put it straight into a deep freeze in California. Shortly thereafter there is an earthquake, as a result of which my frozen brain is split into many parts, a few of which get lost. However, fifty years later, when medical technology has improved, my descendants take the bits of my broken corpse, warm it up and mend it, replacing the missing parts from elsewhere. The body becomes the body of a living person who

behaves somewhat like me and seems to remember quite a lot of my past life. Have I come to life again, or not? Maybe, maybe not. Again there is a truth here, about whether I have survived the haemorrhage as I wanted to, and yet a truth of which we cannot be sure, however much we know about the story of my brain. Hence, my survival consists in the continuing of something else, which I call my soul, linked to my previous body; and I survive in this new body if and only if that soul is connected with it. And note that the extra truth is not a truth about what kind of mental life is connected to the brain, it is not a truth about mental properties, about what thoughts and feelings and purposes the revived person has. Rather, the extra truth, the truth about whether I have survived, is a truth about *WHO* that is, which substance those properties are instantiated in. Once we realize that a human's continuing existence does not logically entail the continuing existence of any particular part of his body, we can also come to see that it does not entail the existence of that body at all. For we can tell a coherent story of a human coming to acquire a new body (as would normally be admitted both by theists who say that this sometimes happens, and by atheists who deny that it does). And since my continuing to exist does not entail my body continuing to exist and conversely, the full history of the world must include the history of my body and the history of the essential part of me, a pure mental substance, my soul.

Dualisms of the physical and mental are not popular philosophical positions today. In Chapter 2 I defended explanatory dualism (two different ways of explaining events); and in this chapter I have defended two kinds of ontological dualisms—there are both pure mental and physical events, pure mental and physical substances. I find the arguments in favour of the latter dualisms (as also the arguments in favour of the former) inescapable. You have left something all-important out of the history of the world if you tell just the story of which physical events were succeeded by which other physical events. What people did intentionally (as opposed to what merely happened to them), and how they thought and felt, are all-important. And equally important is who had those thoughts and feelings—when did one person cease to exist and another come into being.

Now certainly, as I have written, we normally know the answers to these questions. Our observation of bodies normally tells us when persons are the same and what they are feeling. Of course, if a baby

screams when prodded with a needle, it is in pain. But it is not so obvious, when a human-looking organism made in a factory or a creature from another planet is prodded with a needle and emits some sound, whether that thing is in pain. And, of course, the person with this body today who has not been subject to a brain operation and shares the same patterns of behaviour as the person with this body yesterday is the same person as the latter. But after humans, let alone creatures from some distant planet, have had massive brain operations, it is not at all clear whether we are dealing with the same person as before. What these examples bring out is that someone feeling pain is a different event from their being prodded by a needle, and this person being the same person as that person is different from this body being the same body as that; even if normally an event of the latter kind goes with an event of the former kind. A full history of the world will tell the story of feelings as well as of brain events, and of persons (and so their essential non-physical parts, souls) as well as of bodies.

These arguments that show that humans have two parts—body and soul—will show that any creature that has a mental life will also have two parts. The same issues will arise for a chimpanzee or a cat as for a human. If some cat is to undergo a serious brain operation, the question arises whether the cat has reason to fear the bad experiences and look forward to the good experiences that the post-operation cat will have. That question cannot necessarily be answered merely by knowing what has happened to every molecule in the cat's brain. So we must postulate a cat-soul that is the essential part of the cat, and whose continuation makes for the continuation of the cat. Only when we come to animals without thought or feeling does such a question not arise, and then there is no need to postulate an immaterial part of the animal. Certainly human souls have different capacities from the souls of higher animals (the former can have kinds of thought—thoughts about morality or logic—that the latter cannot have; and form kinds of purpose—for example, to solve an equation—that the latter cannot). But what my arguments show is that animals who have thought and feeling have as their essential part a non-physical soul.

Just as I do not wish to deny that brain events cause mental events (that is, events in the soul, once it exists) and vice versa, so I do not necessarily wish to deny that events in the brain play a role in causing the existence of souls. Maybe, at some stage of animal evolution, an

animal brain became so complex that that caused the existence of a soul connected to it, and the continued development and operation of that brain sustained the existence of the soul; and, as evolution moves on, similar complexity causes similar souls. The connection between one soul and one brain that gets established is a causal one. It is events in this particular brain that cause events in this particular soul, and events in this particular soul that cause events in this particular brain; this is what the connection between this brain and this soul amounts to.

At which stage of the evolutionary process did animals first start to have souls and so a mental life? We do not know. But fairly clearly their behaviour shows that the mammals do have a mental life. My view is that probably all the vertebrates have a mental life, because they all have a brain similar to the human brain, which, we know, causes a mental life in us; and their behaviour, too, is best explained in terms of their having feelings and beliefs. Dogs and birds and (probably) fish all feel pain. But there is no reason at all to attribute a mental life to viruses and bacteria, nor in my view to ants and beetles. They do not have the kind of brain that we do, nor do we need to attribute feelings and beliefs to them in order to explain their behaviour. It follows that at some one particular moment in evolutionary history there appeared something utterly new—consciousness, a mental life, to be analysed in terms of souls having mental properties.² My argument in this chapter is an argument from mental events as the instantiation of mental properties in souls, and in this respect too it differs from Locke's 'argument'.

The Scientific Inexplicability of Souls and their Mental Life

Given the scientific laws as we believe them to be, which operated to govern the inanimate world for the first nine of the first fifteen billion years since the time of the Big Bang, there is not the slightest grounds for supposing that conscious life would evolve. The laws of Relativity Theory and Quantum Theory, integrated perhaps into a 'Grand Unified Theory' or 'Theory of Everything' by which everything

² My arguments in this section to show that mental life consists in the instantiation of mental properties in souls are fairly brief. I give far more extensive arguments in chapters 1–9 of my book *The Evolution of the Soul* (Clarendon Press, 1997), when I also consider a whole variety of objections to them.

physical might be explained (fully or partially, even if not completely), give not the slightest reason to suppose that some brain state would cause a green sensation or a sensed smell of coffee. But maybe there is more to the laws of nature than the relatively simple integrated system of physical laws envisaged in a vast physical theory. Maybe there are also psycho-physical laws connecting brains and their states with souls and their states, which would produce experienced effects only when brains had reached a certain stage of development.

Since brain events often cause mental events, and mental events often cause brain events, scientists could perhaps establish a long list of such causal connections in humans. The list would state that brain events of a certain kind cause blue images, and brain events of a certain kind cause red images; brain events of another kind cause a strong desire to drink tea; and that a purpose to eat cake together with a belief that cake is in the cupboard cause the brain events that cause leg movements in the direction of the cupboard. And so on. Also, just possibly, scientists could list which primitive brains give rise to consciousness—that is, to souls. The reason why I wrote ‘just possibly’ is that our only grounds for believing that any other organism—whether some animal whose body was formed by normal sexual processes on earth, or some creature on another planet, or some machine made in a factory—is conscious is provided by the similarity of its behaviour and brain organization to ourselves. We do not have an independent check on whether it is conscious. And when the similarities are not strong—as between frogs, say, and human beings—it is in no way obvious whether the animal is conscious. But let us waive difficulties about how we could establish such things, and suppose that we have lists of causal connections in humans between brain events and mental events, and lists of which kinds of primitive brain give rise to consciousness—that is, souls—in which subsequent brain events cause subsequent mental events, and mental events cause brain events. These causal connections constitute very detailed generalizations, similar to descriptive generalizations of chemistry about which particular substances combine under what circumstances with which other substances to form which new substances.

So does the true scientific theory of the universe consist of the hoped-for integrated theory of physics plus these trillion or so causal connections. That is immensely improbable. By the criteria set out in Chapter 3, a scientific theory (of given scope) is likely to be true in so

far as it has considerable explanatory power, and in so far as it is simple. Such an imagined psycho-physical theory would have the requisite explanatory power. (It would lead us to expect the correlations that we find—since they would be part of the theory.) But it would be so complicated that it would be immensely improbable that it provided a full explanation of mind–body interaction. For that, we need an explanation of these causal connections in terms of their derivability from a theory consisting of a few relatively simple laws that fit together (in the way in which the low-level laws of chemistry proved derivable from the atomic theory of chemistry). The theory would need to explain why the formation of a brain of a complexity as great as or greater than that of a certain animal (perhaps an early vertebrate) gives rise to consciousness—that is, to a soul with mental states. And the theory would need to explain why brain events give rise to the particular mental events they do—why a brain event of this kind causes a blue image, and one of that kind causes a red image, and not vice versa; why eating chocolate causes the brain events that cause the taste we call chocolatey rather than the taste we call pineappley? It would need to explain why this brain event causes the thought that Russia is a big country, and that one causes the thought that every human has a vocation; and why this mental event causes the brain event that causes my lips to utter this sentence, and that mental event causes the brain event that causes my lips to utter that sentence.

A mere list of causal connections would be like a list of sentences of a foreign language that translate sentences of English, without any grammar or word dictionary to explain why those sentences are correct translations. In the absence of a grammar and dictionary you are in no position to know whether the sentence provides the right translation in less usual circumstances (for example, when you are talking to a child rather than an adult); and you cannot translate any new sentence. Analogously, without a psycho-physical theory, you cannot predict whether the same connections will still hold when some other part of the brain is in a less usual state, let alone predict which brain events of a new kind would give rise to which mental events of a new kind, and which new kinds of machine would have feelings and which would not.

A postulated theory of mechanics with high explanatory power, leading us to expect a diverse set of mechanical phenomena that are otherwise not to be expected, will probably be true in so far as it is

simple having a few laws all dealing with the same sort of thing—material objects, their mass, shape, size, and position, and change of mass, shape, size, and position. Physical objects differ from each other in respect of these properties in measurable ways (one has twice as much mass as another, or is three times as long as another). Because the properties are measurable, we can have general laws that relate two or more measured quantities in all bodies by a mathematical formula. We do not merely have to say that, when an object of this mass and this velocity collides with an object of that mass and that velocity, such and such results; and so on for innumerable different objects. We can have a general formula, a law saying that for every pair of material objects in collision the quantity of the sum of the mass of the first multiplied by its velocity plus the mass of the second multiplied by its velocity is always conserved. But that can hold only if mass can be measured on a scale—for example, of grams or pounds; and likewise with velocity. So a theory of mechanics can easily have sufficient simplicity so as to render it, if it predicts sufficiently well, probably true.

However, a psycho-physical theory would deal with things of very different kinds. The mass and velocity and electrical and other physical properties of material objects are utterly different from the mental properties of thought and feeling that pertain to souls. Physical properties are measurable. So brain events differ from each other in the chemical elements involved in them (which in turn differ from each other in measurable ways) and the speed and direction of the transmission of electric charge. But thoughts do not differ from each other along scales. One thought does not have twice as much of some sort of meaning as another one. So there could not be a general formula showing the effects of variations in the properties of brain events on mental events, for the former differ in measurable respects and the latter do not. And what goes for thoughts, goes for mental events of other kinds. A desire for roast beef is not distinguished from a desire for chocolate by having twice as much of something. (Of course, the underlying causes of the one may have twice as much of something as the underlying causes of the other, but that is not the same.) So there could not be a general formula showing how certain variations in brain events produce changes of desires; only a list of which variations in the brain cause which changes of desire. And since sensations, thoughts, and so on do not differ from other sensations, thoughts, and so on in measurable ways, even more

obviously sensations do not differ from thoughts, or purposes differ from beliefs in measurable ways; and so there cannot be an explanation deriving from some simple general formula of why this brain event was caused by a purpose and that one caused a belief, and another one caused a taste of chocolate. Not merely are the kinds of property possessed from time to time by physical objects and by souls so different, but, even more obviously, physical objects are totally different kinds of things from souls. Souls do not differ from each other or anything else in being made of more or less of some quantity of stuff. So, it would just be a brute fact, not a consequence of some deeper theory, that in humans and known animals this degree of complexity will not, and one just a little greater will, give rise to a soul. Because there could not be an explanation of this, we could not tell whether some robot made in a laboratory was or was not conscious. Above all, there could not be a formula that had the consequence that this brain would give rise to my soul and that one to yours rather than vice versa. For these reasons there could not be an explanation of soul–brain correlation, a soul–brain theory that was sufficiently simple to be probably true; merely a long list of inexplicable causal connections. We could discover at most that there were these connections, not why there were these connections.

But does not science always surprise us with new discoveries? The history of science is punctuated with many ‘reductions’ of one whole branch of science to another apparently totally different, or ‘integration’ of apparently very disparate sciences into a super-science. Thermodynamics dealing with heat was reduced to statistical mechanics dealing with velocities of large groups of particles of matter and collisions between them; the temperature of a gas proved to be the mean kinetic energy of its molecules. The separate sciences of electricity and magnetism came together to form a super-science of electromagnetism. And then optics was reduced to electromagnetism; light proved to be an electromagnetic wave. How is it that such great integrations can be achieved if my argument is correct that there could not be a simple and so probably true super-science that predicts the connections we find between mental events and brain events?

There is a crucial difference between these cases. Every earlier integration into a super-science, of sciences with entities and properties apparently qualitatively very distinct, was achieved by saying that really some of these entities and properties were not as they

appeared to be. A distinction was made between the underlying (not immediately observable) physical entities and physical properties, on the one hand, and the sensory properties to which they gave rise. Thermodynamics was initially concerned with the laws of temperature exchange; and temperature was supposed to be a property inherent in an object that you felt when you touched the object. The felt hotness of a hot body is indeed qualitatively distinct from particle velocities and collisions. The reduction to statistical mechanics was achieved by distinguishing between the underlying cause of the hotness (the motion of molecules) and the sensation that the motion of molecules causes in observers, and saying that really the former was what temperature was, the latter was just the effect of temperature on observers. That done, temperature falls naturally within the scope of statistical mechanics—for molecules are particles; the entities and properties are not now of distinct kinds. Since the two sciences now dealt with entities and properties of the same (measurable) kind, reduction of one to the other became a practical prospect. But the reduction was achieved at the price of separating off the felt hotness from its causes, and only explaining the latter.

All other 'reductions' of one science to another and 'integrations' of separate sciences dealing with apparently very disparate properties have been achieved by this device of denying that the apparent properties (such as the 'secondary qualities' of colour, heat, sound, taste) with which one science dealt belong to the physical world at all. It siphoned them off to the world of the mental. But then, when you come to face the problem of the mental events themselves, you cannot do this. If you are to explain the mental events themselves, you cannot distinguish between them and their underlying causes and only explain the latter. The enormous success of science in producing an integrated physico-chemistry has been achieved at the expense of separating off from the physical world colours, smells, and tastes, and regarding them as purely private sensory phenomena. What the evidence of the history of science shows is that the way to achieve integration of sciences is to ignore the mental. The very success of science in achieving its vast integrations in physics and chemistry is the very thing that has apparently ruled out any final success in integrating the world of the mind and the world of physics.

As we saw in Chapter 8, the Darwinian theory of evolution by natural selection is able to provide the framework of an explanation of the evolution of human and animal bodies, though not,

I suggested, a complete or ultimate explanation. But that Darwinian explanation would explain equally well the evolution of inanimate robots. Could not Darwinism also tell us something about how bodies came to be connected with consciousness—that is, souls? Natural selection is a theory of elimination; it explains why so many of the variants thrown up by evolution were eliminated—they were not fitted for survival. But it does not explain why they were thrown up in the first place. In the case of physical variants (such as the length of the giraffe's neck), there is no doubt an adequate explanation in terms of a mutation (a random chemical change) producing a new gene with properties that cause the new variant to appear in accordance with the basic laws of chemistry. But our problem is to explain why some physical state causes and sustains the existence of souls with such mental properties as beliefs, desires, purposes, thoughts, and sensations, causally connected in a regular way with brain states. Darwinism is of no use in solving this problem.

Darwinian theory might, however, be of use in solving one different problem, and certainly is of use in solving a third problem; but neither of these problems must be confused with the original problem. The first of these additional problems is why, having once appeared in evolutionary history, conscious animals survived. Darwinian theory might be able to show that conscious organisms have some advantage in the struggle for survival over non-conscious organisms programmed to react to their environment in similar ways. It is difficult to see what that could be, but maybe there is an advantage.

The second additional problem is one to which Darwinism can produce a clear, and to my mind fairly obviously correct, answer. That is this problem. Given the existence of mind-brain connections, and given that organisms with a mental life will be favoured in the struggle for survival, why are the brain events that cause and are caused by mental events connected with other bodily events and extra-bodily events in the way in which they are. Take beliefs. A brain event causes the belief that there is a table present. That brain event is caused by a nerve impulse travelling along the optic nerve from the eye when a table image is formed in the eye by light rays arriving from a table. But an animal could have evolved in which the brain event that caused the table belief was caused by quite different events in the outside world. Why these particular connections between the brain and the outside world? The answer is evident:

animals with beliefs are more likely to survive if their beliefs are largely true. False beliefs—for example, about the location of food or predators or obstacles—will lead to rapid elimination in the struggle to get food or avoid predators. If you believe that there is no table present, when there is one, you will fall over it, and so on. Those in whom the brain states that give rise to beliefs are connected by causal chains to the outside world, in such a way that the causal chain is normally activated only by a state of affairs that causes the brain state, which in turn causes the belief that that state of affairs holds, will normally hold true beliefs about the world and in consequence be more likely to survive. And, just as there is an evolutionary advantage if perceptual beliefs are connected to the outside world in the right way, so there will be an evolutionary advantage if the procedures for forming new beliefs on the basis of previous true beliefs usually lead to new true beliefs. It is most improbable that these procedures would have this feature unless creatures use correct criteria of what is evidence for what. We acquire our beliefs about the world by observing features of the world and then devising theories to explain these features that satisfy the criteria of a probably true explanation, which I analysed in Chapter 3. The simplest case of application of these criteria is simple generalization. Observing several people becoming ill after they have eaten a certain berry (where the diet and circumstances of each person differ from that of the others in other respects, and no one who has not eaten that berry becomes ill in that way), we put forward as a probable explanation of the illness that it was caused by eating the berry. That licenses an extrapolation to ‘that berry always poisons’. It is highly probable that natural selection will ensure the survival of those organisms and only those organisms that use correct criteria of inductive inference. But all this requires there to be brains that throw up various different beliefs connected in various ways with other beliefs, and then nature selects those creatures in which correct inferential procedures are incorporated.

Similarly, given that I am going to have desires caused by brain events, there are evolutionary advantages in my having some under some circumstances rather than others under other circumstances—for example, a desire for food when I need to eat rather than when I do not need to eat. The same kind of account can be given of why the brain events produced by intentions give rise to the intended bodily movements. If, when I tried to move my foot, my hand moved instead, predators would soon overtake me. But this

correct explanation of why (given that intentions cause brain events) the brain is connected by nerves to the rest of the body in the way it is does not explain why we have intentions to move our bodies at all and why they cause brain events, which is a quite different problem. I conclude that the existence of the most novel and striking features of animals and above all of humans (their conscious life of feeling, choice, and reason, causing connected to their bodies) seems to lie utterly beyond the range of successful scientific explanation.

The Argument to God

And yet there are these regular causal connections. These are causal connections (in both directions) between kinds of brain event and kinds of mental event, so detailed and specific that it is most improbable that they would occur without an explanation; yet it is immensely improbable that there could be a scientific explanation of the connections. Mind–brain connections are too ‘odd’ for science to explain; they cannot be consequences of a more fundamental scientific theory, and there are simply too many diverse connections to constitute laws. But once again there is available a personal explanation: God being omnipotent, is able to join souls to bodies. He can cause there to be the particular brain-event–mental-event connections that there are. He can do this by causing molecules when formed into brains to have powers to produce mental events in souls to which they are connected, and the liabilities to execute the intentions of such connected souls (new powers and liabilities not deriving from the ordinary ones, which chemistry analyses). And he can make the souls in the first place and choose to which brain (and so body) each soul is to be connected when foetal brain events require a soul to be connected to the brain.

God has good reason to cause the existence of souls and join them to bodies, in the goodness (on which I commented in Chapter 6) of the existence of humanly free agents who would need to have bodies through which to have enjoyable sensations, form largely true beliefs about the world, and form their own purposes in the light of these beliefs, which would make a difference to the world. I argued that there was a significant probability that God would make such creatures. Their existence involves the existence of regular causal

connections between mental events and events in their bodies. Given that humans are humanly free agents, it involves regular causal connections between mental events and events in human brains. We cannot make a difference to the world if, each time we try to move our leg, some different effect is caused in the brain and thereby in the body—one time the arm moves, one time we find ourselves sneezing, and so on. Likewise, if we are to discriminate between one object and another, they have to look (feel, etc.) different, and so there has to be a regular causal connection between the brain events caused by objects of each kind and the mental visual impressions of them. And, if we are to have the awesome power of reproduction, there have to be regular connections between our sexual acts, the foetus to which they give rise, and some soul or other linked to that foetus. God has reason to set up all these connections. He may have a reason to make this brain state cause a red sensation and that one to cause a blue sensation rather than the other way round, but, if there is no particular reason why one connection is better than a rival one, God has a reason by a ‘mental toss-up’ to produce one-or-other connection. He may have a reason to join this soul to this particular body, but again, if there is no reason for joining one soul to one body rather than to a different body, he has reason by a ‘mental toss-up’ to produce one-or-other connection—that is, to make it a chance matter which connection holds.

So then, because we have every reason to believe that there can be no scientific theory and so scientific laws correlating brain states with souls and their states, we have every reason to believe that the causal connections that exist between them do not have a scientific explanation in terms of the properties of brain states; they are additional causal connections independent of the set of scientific laws governing the physical world. Nothing about the physical world makes it in the very least probable that there would be these connections. Let e be the existence of souls with mental states connected to brain states in the ways in which we have been analysing; k be the premisses of the arguments of the previous chapter—that there is a law-governed physical world of the type analysed in Chapter 8 with laws and boundary conditions tuned so as to allow the existence of human bodies; and let h as before be the hypothesis of theism. Then $P(e|\sim h \& k)$ is very low. But, for all the reasons analysed in Chapter 6, a God has very good reason for creating humans (and good reason to create animals); hence $P(e|h \& k)$ has a moderate value. Hence the

argument from consciousness is a good C-inductive argument for the existence of God.

The argument from consciousness that I have now set out, with, I hope, some rigour, captures in a precise form the feeling of amazement that many people have when, aware of the web of physical laws governing such inanimate physical objects as electrons, protons, and photons operating through boundless space and endless time, they are also aware of something quite different, interacting with the physical objects in very limited regions of space (such as human bodies on earth) over a very limited region of time. Here, they have felt, is something outside that web of physical laws—which needs explanation of a different kind—explanation in terms of the action of an agent in certain respects similar to the human agents whose existence needs explanation. Here, they have felt, are evident footprints of the divine. I have been arguing that this feeling is a justified one.³

Humanly free agents need, in order to develop a science and metaphysics, first the ability to be able to form scientific and metaphysical concepts; and then the ability to put forward scientific and metaphysical theories. Given all that, they then have the choice of advancing and testing such theories or of not bothering to do so. It is good to understand the world and to use science to improve the human condition. Scientific and metaphysical discovery is a communal activity, and not everyone needs to have the ability to make the necessary discoveries, so long as others can use the results of these discoveries and so to some small extent test them for themselves. But there do need to be geniuses if there is to be progress. Further, if all humans are to be able to worship God, they will need to be able to develop some fairly sophisticated concepts (omnipotence, omniscience, etc.) I argued in Chapter 6 that it is fairly probable that God would create humanly free agents with these abilities. So, once again, there is good reason for supposing that, if there is a God, there will be such agents, and no very good reason for supposing that they will evolve if there is no God—why should not evolution have stopped at the level of the monkeys? Given that humans began to evolve, such conceptual development might give them a selective advantage that

³ Alvin Plantinga has developed recently an ‘argument against Evolutionary Naturalism’, which makes some of the same points, as does my exposition of the argument from consciousness. I examine Plantinga’s argument in Additional Note 3.

would ensure their survival. But without a particular kind of biochemistry there will not occur the kinds of mutations that will permit further progress. There is no particular reason to expect that any biochemistry will permit mutations of that kind unless there is God, who has reason to give to the biochemical basis of mind this particular causal power.

The Argument from Moral Truth

Moral awareness is confined to the same conscious beings as is metaphysical awareness. Many theists have held that through morality God makes himself known intimately to humans, for the voice of conscience is the voice of God. We find an argument from morality to God classically in Kant—although he would vigorously deny that he is putting forward anything that can properly be called an argument;⁴ and also in writers of the late nineteenth and early twentieth centuries, who, when the arguments from design were driven out of fashion, cherished this as their favourite argument.

It is crucial to distinguish two different arguments from morality. First there is the argument from the fact that there are moral truths and secondly there is the argument from human awareness of moral truths. I begin with the former. It does, of course, in its premiss take for granted moral objectivism—that moral judgements have truth values. If there are no moral truths, then there is no fact about the world for this argument to take off from. The issue then arises, given that moral judgements are propositions with truth values, whether true moral propositions are logically necessary or

⁴ See I. Kant, *Critique of Practical Reason* trans. L. W. Beck (Liberal Arts Press, 1956), book 2, esp. ch. 2, sect. 5. Kant claims that the existence of God is a 'postulate of pure practical reason'—that is, that the existence of God entails that the *summum bonum*, the perfection of the universe that the moral law commands us to seek, is attainable; and that it would not otherwise be attainable. Hence, he claims, the binding obligation upon us to keep the moral law is something that makes sense. Kant would, however, vigorously deny that this constitutes an argument for the existence of God, since he would deny that the binding obligation to keep the moral law can in any way be formulated as a theoretical truth. We simply find ourselves feeling the force of the moral law, and try to make sense of how it can be that we are under the moral law. But I cannot see how it can be rational for us to conform to the moral law (as Kant believes that it is), unless we believe that what the moral law states is true—for example, that murder is wrong, and promise keeping obligatory. Hence in effect Kant seems to be putting forward an argument from the fact that there are binding moral truths.

logically contingent. An argument that claims that the best explanation of the existence of morality is the action of God who created it must claim that many moral truths are (logically) contingent. For the existence of the phenomena described by (logically) necessary truths needs no explanation. It does not need explaining that all bachelors are unmarried, or that, if you add two to two, you get four. These things hold inevitably and necessarily, whether or not there is a God. Now clearly, if there are moral truths, quite a lot of moral truths are contingent. That whatever action was done by Hitler at 10.00 a.m. on 3 December 1940 was morally wicked is, if true, clearly contingently so; for it is coherent to suppose that Hitler might have done a good action at that time. And, if it is true that it is good to give £10 to some beggar, it is clearly a contingent truth, whose truth depends on whether he will spend it on drugs that will kill him or on food that will prevent him from starving. And, plausibly, whether killing a certain human is wrong may also depend on contingent circumstances, such as whether he is trying to kill you or whether he is simply an innocent passer-by. Generally actions may be picked out simply by their spatio-temporal coordinates or by a description that leaves open their moral status. Yet actions cannot be just obligatory or right or wrong. They must be obligatory or right or wrong in virtue of their possessing certain natural properties (that is, properties that could be recognized by someone without moral concepts). And, once one has described an action in terms of all the natural properties that it possesses (in terms of all its circumstances and effects), then—if it is wrong—it will be necessary that it is wrong; and—if it is good—it will be necessary that it is good. For, if one action is right, and another imagined one is wrong, there must be some natural feature that the second action has and the first action lacks that makes the second action wrong. It is not coherent to claim that *a* is wrong, and that *b* differs from *a* in no natural property but that, unlike *a*, *b* is right. There could not be a world that was different from our world solely in the respect that murder was wrong here, but right there. There would have to be some natural features of the other world that made murder right there—for example, that there murdered people quickly come to life again. It follows, given moral objectivism, that contingent claims that some actions are right (or wrong) hold in virtue of contingent truths that the action has certain natural properties and necessary truths that actions with those properties are right (or wrong, as

the case may be). Fundamental moral principles must be (logically) necessary.⁵

Now, if the basic moral principles are necessary, the existence of what they describe cannot provide an argument for the existence of God. An argument could take off only from the truth of some or all contingent moral truths (for example, from the fact that it is wrong to drop atom bombs on Japan, rather than from the fact that it is wrong to kill people who will not certainly come to life again). Now the fact that certain moral truths hold can confirm, add to the probability of, the existence of God only if it is more likely that those moral truths hold if there is a God than if there is not. The contingent truths that actions, *a*, *b*, *c*, *d* are obligatory (or right or wrong, as the case may be) depend on *a*, *b*, *c*, *d* possessing certain natural properties, *Q*, *R*, *S*, *T*, which of logical necessity make them obligatory (or whatever). So, if there is to be an argument to the existence of God from certain actions being obligatory, it will have a structure somewhat as follows: actions *a*, *b*, *c*, *d* are obligatory; they would not be obligatory unless they were *Q*, *R*, *S*, *T*. It is more probable that they are *Q*, *R*, *S*, *T* if there is a God than if there is not; therefore the obligatoriness of *a*, *b*, *c*, *d* confirms the existence of God.⁶

Now the most plausible candidates for actions that would not be obligatory unless there were a God are actions like promise keeping and truth telling, for the obligatoriness of which there is no easy utilitarian justification. Starting from the obligatoriness of such actions, we could construct a Kantian argument along the following lines. (Although this argument is very much in the spirit of Kant's *Critique of Practical Reason*, Kant himself denies that he is putting forward *an argument* to the existence of God—see my p. 212 n. 4). 'Promise keeping is always obligatory. But an action is obligatory if and only if it conduces to the perfection of the universe—what Kant calls the *summum bonum*. It is more probable that promise keeping will conduce to the *summum bonum* if there is a God than if there is not. (This, it may be urged, is because keeping secret promises to dying people would be pointless if there is no life after death, in

⁵ For further argument on these points, see *The Coherence of Theism* (Clarendon Press, 1993), ch. 11.

⁶ The argument could be varied in various ways such as making the step from *a*, *b*, *c*, *d* being obligatory to their being *Q*, *R*, *S*, *T* a mere probabilistic one, and making the next step a deductive one.

which the promisee can perceive that the promise has been kept; and it is more probable that there is life after death if there is a God than if there is not.) Therefore the obligatoriness of promise keeping confirms the existence of God.' This argument is valid, but its first and third premisses are highly questionable. Some will deny the first premiss—that promise keeping is always obligatory. Others will deny the third premiss—that it is more probable that promise keeping will conduce to the *summum bonum* if there is a God than if there is not. For they would say that the very act of promise keeping as such contributes to the *summum bonum*, and any further consequences are irrelevant. A moralist of a teleological viewpoint will tend to deny the first premiss, and ones similar to it; and a moralist of a deontological viewpoint will tend to deny the third premiss and ones similar to it. What goes for the particular argument that I stated is liable to go for similar arguments (for example, ones that appeal in their first premiss to the invariable wrongness of lying). Now, of course, both the first and third premisses could be true even though most people have an initial inclination to deny one or the other. But to get the moral argument off the ground you would need arguments to show the first and third premisses to be true. As it stands, the argument is not a good argument (for the reason that the premisses are not accepted by disputing parties). I am too pessimistic about the prospects to devote more time to attempting to supplement the argument by producing good arguments to support its premisses. One reason for this is that I cannot see how anyone who holds one of the first and third premisses but not the other is going to be persuaded by a process of rational argument to hold the other, unless he is *first* persuaded by some other argument that there is a God. For this reason I cannot see any force in an argument to the existence of God from the existence of morality.

The Argument from Moral Awareness

Very different from the argument from the fact that there are moral truths is the argument from human awareness of significant moral truths.

If humans are to make significant choices at all they must have the concepts of moral goodness and badness (in my sense of overall goodness and badness). They must be able to see some actions as

good to do, and of these to see some as obligatory; and to see some actions as bad, and of these, some as wrong. I argued earlier that, in order to do any action at all, an agent must see it as in some way a good action to do; and so it follows merely from the agent having intentions that he will have the concept of an action being good (and so of the concept of an action being bad). But his understanding of goodness might be extremely limited. He might see the goodness of an action merely in terms of the action being good to do, simply for its own sake (he just wants to do it; he cannot think why), or because it produces pleasurable sensations. He might have no conception of the distinction between what he desires to do and what is (overall) desirable (good) to do; and so no conception of an action being good because it makes other people happy, extends their lives, or fulfils a promise, even though he does not desire (want) to do the action. If God is to give us significant choices, he will ensure that we develop this kind of moral awareness.

But if there is no God, how likely is it that embodied creatures with a mental life will progress to this stage? If genetic mutations produce creatures naturally inclined to behave altruistically towards others of their community (though there does not seem to be any particular reason why animal biochemistry should be of such a kind that these mutations would occur), then there may well be a good Darwinian explanation for their survival.⁷ For a community of creatures who help each other (defend others of the community against predators, divide tasks so that some care for the young while others forage for food and so on) may well be more likely to survive than a group of creatures who do not help each other. Such creatures may help each other spontaneously and naturally, as many groups of animals do. But having the understanding of these actions as morally good (even when we do not desire to do them) is something beyond mere altruistic behaviour.⁸ And there seems no particular reason why any mechanisms of mind–body interaction that cause creatures to have beliefs should produce moral beliefs. Creatures would need brain states with the power to produce moral beliefs or the ability to

⁷ As argued in Elliott Sober and David Sloan Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behaviour* (Harvard University Press, 1998). And see the corrections to their model suggested in Peter Gildenhuys, 'The Evolution of Altruism: The Sober/Wilson Model', *Philosophy of Science*, 70 (2003), 27–48.

⁸ 'Just as an individual can be an altruist without being moved by moral principles, the converse is also possible' (Sober and Wilson, *Unto Others*, 239).

acquire them; and so their genes would need to produce mutations that would cause such brain states, and for this purpose they would need special kinds of gene. So again, while God will give some creatures moral beliefs as features essential to their being humanly free agents, there is otherwise no particular reason why whatever processes give creatures beliefs should give them moral beliefs. This is shown by the fact that, as far as we can tell, there are many species of animals that are naturally inclined to help others of their species, and yet do not have moral beliefs—there is no reason to suppose that lions and tigers have moral beliefs, or could ever develop them. A tiger may desire and so believe it a good thing to help one particular fellow tiger in distress without believing it to be of overriding importance to help any other tigers in distress, when he did not wish to do so.⁹ However, if some mutation gave to creatures of some community such moral beliefs or the ability to acquire them, such beliefs would reinforce any selective advantage possessed by a natural inclination to altruistic behaviour. For the moral beliefs would move those creatures to altruistic behaviour when the natural inclination was lacking (for example, it would move them to care for offspring for whom they no longer had much natural affection).

While having a desire to do an action that is in fact morally good does not require the belief that it is morally good, the belief that some action is morally good does require a desire, however weak, to do that action—as I argued in Chapter 5. You could not really believe that it is morally good to comfort someone in distress unless you had some minimal desire to help them, even if you have stronger desires to do other things instead. And, as I also argued in Chapter 5, having

⁹ For this reason I wish to retract my endorsement in previous editions of this book, and elsewhere, of Darwin's claim that 'any animal whatever endowed with well-marked social instincts . . . would inevitably acquire a moral sense or conscience, as soon as its intellectual powers had become as well, or nearly as well, developed as is man' (Charles Darwin, *The Descent of Man* (2nd edn. John Murray, 1875), 98). For moral awareness, the creature needs to be able to contrast what he desires to do and what is desirable; and he needs also, I suggest, to have an understanding of the desirable (the overall good) as including the good of other creatures. Once he has that understanding, he will have the concept of moral goodness and be equipped to extend his understanding of what things are morally good—e.g. in virtue of his other 'intellectual powers', he will be able to recognize the goodness of helping members of other communities as an action sufficiently similar to an action of helping members of his own community, for the former as well as the latter to be a morally good action. On the naturalness of this process, see my *The Evolution of the Soul*, chs. 11 and 12. But I argue now, as I did not then, that the initial stage of this process does require the acquisition of a new concept, perhaps as a result of a genetic mutation.

a free choice between doing what you believe good and what you believe bad requires having a stronger desire to do the latter. Without temptation, a strong desire to do what is bad, there would be no free choice between good and bad. Moral choice requires moral awareness and desires both good and bad. And, of course, our world has plenty of both good and bad desires. There is a natural love of parent for child, man for woman, fellow for fellow, which includes desires to remove the suffering and promote the well-being of others. And there are also desires for fame and fortune greater than others have, and desires to hurt and maim others.

My claim is that, like each of the other phenomena discussed in Chapters 7 and 8 and earlier in this chapter, there is no great probability that moral awareness will occur in a Godless universe, and an increasingly large improbability, as we consider more and more phenomena, that they will *all* occur—for example, not merely will there be a universe, but it will be governed by simple laws, etc., etc., and contain conscious beings with moral awareness. Yet, I claimed in Chapter 6, a God has significant reason to bring about conscious beings with moral awareness (and with further features yet to be discussed), and so to bring about all the phenomena discussed previously that are necessary for their existence. The reason for God to give to humans moral awareness is to give them a free choice between good and evil. I argued briefly in the previous chapter that there is reason to suppose that we do have free will. This ‘significant reason’ made it significantly probable that God would produce a universe in which there are humanly free agents (as well as the conditions necessary for their existence); and I gave to this probability the somewhat arbitrary value of $\frac{1}{2}$. This probability entails that there will be a probability of much greater than $\frac{1}{2}$ that God will create each particular one of the necessary conditions (a physical universe, governed by natural laws, which are such as together with boundary conditions are conducive to the existence of human bodies, these bodies being the bodies of conscious beings who have moral awareness) for this. For the probability that God will create humanly free agents is equal to the product of the probabilities that he will create each of these necessary conditions. Whereas, as we consider more and more phenomena, the probability that they will all occur if there is no God gets less and less. Hence the fact of moral awareness provides one more good C-inductive argument for the existence of God.

The Argument from Providence

If we had choices between forming intentions to do good actions and intentions to do bad actions, but our intentions never made any difference to what happened, we would be living in a deceptive world. A good God would not subject us to such a radical deception. And a God who wished to give us significant free choices would give us efficacious choices, choices that made important differences to ourselves, each other, and the physical world. I argued in Chapter 6 that God had very good reason to create humanly free agents with significant responsibilities for themselves, each other, and the world. In this chapter I describe the enormous scope of the responsibilities possessed by humans.

This argument from the opportunities we have for making significant differences I call the argument from Providence. The world in which we are placed is in this all-important respect providential. The argument echoes points made by many thinkers over the past two or three millennia; but I do not know of anyone who has put them together in the form of precise argument for the existence of God. However, these opportunities to make significant differences require the occurrence of actual evils,¹ and the possible occurrence of many more evils; and the question inevitably arises of whether a good God would be right to give us these opportunities in view of the evils they bring with them. I shall briefly draw attention in this chapter to many of the evils necessary for us to have these opportunities; and then consider in the next chapter whether allowing them and some other

¹ All that I mean by calling an event or state of affairs an 'evil' is that it is a state such that in itself, apart from its circumstances, causes, and consequences, it is bad that it should occur. I do not imply that bringing it about or allowing it to occur would be (what would ordinarily be called) an evil act, or even a bad act at all. I argue in Chapter 11 that it is sometimes a good act to allow or even bring about a bad state. I call such bad states 'evils' simply to conform with much philosophical usage.

evils also necessary for this and other purposes is compatible with the goodness of God.

Man's Opportunity to Provide for Himself

It is a great good for humans to be able to affect themselves—to choose how they are to live, whether to acquire knowledge of the world, and to have the opportunity to form their own character; to have, that is, what I may call a choice of destiny.

Merely to have a body (as I defined it in Chapter 6) involves having a machine room for the maintenance of which we are responsible. We have the choice of continuing to exist (by giving ourselves food and drink), giving ourselves pleasures and pains by what we do with our bodies, damaging or increasing our bodily powers (by rest, exercise, and sleep). But the range of choice is greatly extended by our living in an environment where geography provides dangers, food is limited, there are predators, and other humans compete with us for the things that will satisfy our desires.

Geography is dangerous—there are rivers and seas where we may drown, cliffs over which we may fall, forests in which we may get lost, weather in which we may freeze. Food is limited—edible plants grow in some places and not others; edible animals need to be caught. There are predators—the first humans had to avoid tigers and snakes. And other humans had desires also for food, drink, shelter, and mates; they were in competition with each other. In this environment humans had to learn how to survive and flourish; it became a complicated business with a range of choices of short- and long-term actions. Humans had to learn how to hunt land animals, catch fish, and grow plants, and to choose to use the method of food acquisition that would most probably produce the best results in their particular environment. We had to learn to build shelters, where other humans would find it difficult to tear them down in order to use the material to build their own shelters. And so on. God who seeks to entrust our well-being to us, and wishes the best for us, has reason to give us a range of choices where more worthwhile goals can be obtained by more effort and longer-term policies. But the mere operation of some laws of biochemistry that produced human bodies could have placed them in a lazier environment. Rivers might all have been shallow, cliffs non-existent, food plentiful, no predators

preying on humans, plenty of shelter for all humans, all of whom had relatively few children so that there was little competition for shelter and other good things.

If humans were to have the significant choice of learning (or not bothering to learn) various things necessary for their survival and flourishing in a dangerous environment, they had to begin in ignorance of them. If they were to survive at all until they acquired knowledge, they needed natural desires and aversions to guide them. And, of course, nature provides these. There are natural desires for food, sleep, drink, and sexual intercourse that ensure our own survival and that of our race until we learn the consequences of eating or not eating, of sleeping regularly or doing without sleep, and so on. And there are natural aversions—to darkness and heights. And there are the biologically useful painful sensations that we feel when we touch hot or sharp objects, which lead us instinctively not to do such things again. But, as we learn how the world works, so we have the choice of taking the risk of walking in darkness or along cliff paths, or avoiding the danger.

I have already drawn attention in Chapters 8 and 9 to the two crucial features of the world that make it possible for us to discover the effects of our actions—regularities in the behaviour of things; and our acquiring beliefs about these regularities and using those patterns of inference from them that prove successful and so (given natural selection) survive in the human race. It is highly improbable that our inferences would be successful unless we used correct criteria of what is evidence for what. I noted in Chapter 9 that we need particular kinds of mental ability in order to do science. But there are two particular features of our universe that respectively make this task challenging, and also possible. The first feature, to which I drew attention in Chapter 8, is that the phenomenal laws on which we naturally rely and our knowledge of which we naturally extend by simple generalization depend on fundamental laws. It is this feature that makes the pursuit of science no mere matter of discovering correlations between observable phenomena, but a matter of proposing and testing deep theories. This makes the choice of whether to do fundamental science a very significant choice for individuals and societies about how to use time, energy, and money. The second feature is that human intelligence would have been unable to discover the fundamental laws unless those laws had been such as to allow the construction of instruments that would allow us to detect what was

happening on the small scale and at a large distance away. To advance science we needed microscopes and telescopes; unless light (and other electromagnetic radiation) could be refracted and reflected by various materials (initially, glass), nature would not have revealed its secrets. As we discover the fundamental laws, we can choose whether to build atomic bombs, or send rockets to the moon, or cure cancer, or not to bother. A generous God would seek to give humans a range of choices. But, barring God, there is no particular reason to expect that we should have this ever-widening range of choice.

It is a further greatly significant choice to be able gradually to change one's character—to make it the case that some heroic actions that we cannot now choose to do eventually become natural to us; or, alternatively, to be able gradually to allow ourselves to opt out of morality, to choose to be uninfluenced by moral considerations. Humans are able to form their characters as a result of a crucial contingent feature of their nature—that doing a good action when it is difficult makes it easier to do a good action next time; and doing a bad action when it is not too difficult to resist doing it makes it more natural to do it next time. Each choice for good or ill marginally shifts the range of actions open to us—frequent good choices make heroic actions serious possibilities for us, when previously they were not live options; frequent bad actions bring really wicked actions within the range of psychological possibility. Humans can in these ways form their characters. In a Godless world there is no reason to expect that, even given that creatures make moral choices, these choices would affect character in this way. It might be just as hard to show courage after you have shown courage on innumerable past occasions as on an occasion when you had never shown courage before. Yet a God who wants to give us really significant choices would give us this choice of forming our characters for good or ill. But he would, I suggest, ensure that our character became formed only through a series of choices over time, expressing a fixed determination and not mere sudden impulse to become or allow ourselves to become a certain sort of person. All this, of course, holds in our universe.

Man's Opportunity to Provide for Others

So far I have been talking about individual agents, and the desirability of their having control of their own destiny. I have been

considering them as though each lived like Robinson Crusoe on his own island. But in our world, of course, things are not like that. Humans are interdependent to an enormous degree. The very birth of a human being requires the cooperation of a man and a woman. And when society is organized so that there is specialization of labour (one person is a carpenter, one is a mason, one is a potter, one is a farmer, and so on), a decent standard of living is reached for all that would be difficult if not impossible to attain if each person did all the jobs for himself. And, of course, building aeroplanes, cyclotrons, and radio-telescopes, and sending people to the moon, are quite impossible without cooperation. A world in which good things can be attained only by cooperation is one that a God has reason to make—for cooperation in a worthwhile task is a good thing; and it is good that we should have the opportunity to cooperate when cooperation matters. Our world is like this.

All that I have suggested so far is that it is good that the world should be one in which *A* can benefit *B*, and *B* can also benefit *A*, if they agree to work together. In such a world, however, benefiting others will always have its reward for the benefactor. Dependence can go a lot further than that. We can, for example, have a situation where *A* can benefit *B* but *B* can confer no benefit on *A*, although he can benefit *C*. Is a world with the opportunity for unrequited benefit good thing? Surely yes, for it is a great good for me to benefit you, to be able to give things to you and do things for you. Think how awful it would be if we could never be of any use to anyone. And if God creates a world in which we can benefit others without reward, he creates a world in which we can share in his creative work on the same (unrewarded) terms as himself. Our world has plenty of opportunity for such benefit. One obvious case is provided by the parent-child relationship. In his infancy and childhood a human being depends very largely for his existence and health, for his knowledge of the world, and for encouragement to develop his character, on parents and others, including doctors and teachers. As parents and children get older, parents may become dependent on their children. But parents may die before they need their children to care for them. Yet our children in their turn can give good things to their own children. It is especially good that people should have the opportunity to help others and show their concern for others when the others are at the lowest. It is a privilege for someone to be able to care for the infirm, help the sick, and talk to the lonely.

It is a good thing not merely that agents should need to cooperate with each other and depend on others, but that they should enjoy doing so, and more generally that they should get pleasure from collaboration in work with others, in helping others, and in simple companionship with them. Once again, our world is one in which there is plenty of such pleasure. Indeed it is one in which creatures have a variety of needs for collaboration and company and service to others, and a variety of kinds of pleasure from the satisfaction of these needs. There is the need for a parent, for a child, for a permanent mate, for a friend, for a casual acquaintance, for a colleague with whom to collaborate in working hours on a common project. The world is such that collaboration, companionship, and service of such varied kinds is useful, gives pleasure, and is wanted for its own sake. It is good too that the world should be such that there is the opportunity for a kind of cooperation between many people over many generations to build up human knowledge and extend human power. This cooperation in the acquisition of knowledge has sometimes been going on for many generations—sometimes generations of investigators (especially scientists) have helped each other forward, and worked consciously in order to add their bit to the corpus of knowledge. Yet people seem only to be beginning to take the opportunities that exist for cooperation for long-term practical ends. Planning of towns and schools and populations so as to provide for the welfare of people many decades in the future is something to which politicians have only recently given much attention. But, as knowledge and power increase, so the opportunity for such planning increases radically.

As well as being good that people should have the opportunity to benefit each other, is it good that they should have the opportunity to harm each other? I suggested in Chapter 6 that it is. A God has the power to benefit or to harm. If other agents are to be given a share in his creative work, they must have that power too (although perhaps to a lesser degree). A world in which agents can benefit each other but not do each other harm is one where they have only very limited responsibility for each other. If my responsibility for you is limited to whether or not to give you a camcorder, but I cannot make you unhappy, stunt your growth, or limit your education, then I do not have a great deal of responsibility for you. Your well-being will not then depend greatly on me. God has reason for going beyond that. A God who gave agents only such limited responsibility for their

fellows would not have given much. He would be like a father asking his elder son to look after the younger son, and adding that he would be watching the elder son's every move and would intervene the moment the elder son did a thing wrong. The elder son might justly retort that, while he would be happy to share in his father's work, he could really do so only if he was left to make his own judgements as to what to do within a significant range of the options available to the father. A good God, like a good father, will delegate responsibility. In order to allow creatures a share in creation, he will allow them the choice of hurting and maiming, of frustrating the divine plan. Our world is, of course, one where creatures have just such deep responsibility for each other. I can not only benefit, but harm my children. One way in which I can harm them is that I can inflict physical pain on them. But there are much more damaging things that I can do to them. I can deprive them of adequate food, of play, and, above all, of affection. They are made so as to need certain things for their well-being. If God has made me, he has given me the opportunity to deprive them of what they badly need.

One crucial way in which one agent *B* may depend on another *A* is that whether and how *B* grows in freedom, power, and knowledge depends on *A*; and, if *A* is to have a very deep responsibility for *B*, he will have that sort of responsibility. That allows *A* to stunt *B*'s growth in freedom and power, give him false belief instead of knowledge, and generally make him retarded. I commented earlier on a human's ability to improve his own character or allow it to deteriorate, by shifting the range of actions that come naturally to him towards the good or towards the bad end of the spectrum of possible actions. But we can influence each other for good or ill in this process of character formation; in particular we can influence our friends and neighbours and family and above all our children. We can teach our children which actions are good, and we can encourage them to do these actions by word and example, reinforced by an evident loving concern for them. Or we can show no concern for their well-being, lie to them, and be violent towards them, which will influence them to behave in a similar way towards us and others. The influence we can have on the character of our children is frightening in its enormity but—thank God—not total. They are subject to other influences as well as those of their parents, and have some freedom to resist influence.

An obvious way in which a God can give to agents the opportunity to harm or benefit each other in the long run is to produce a world of decay. Embodied humans might have been such as to flourish, needing to do nothing in order to live a reasonably enjoyable life either forever or until they suddenly died. But in fact our bodies are subject to accident and disease; we need to act continually to prevent bad things happening to our bodies. And we need to act to clear away the weeds to allow our plants to produce food, and to repair our houses as their bricks and beams collapse. Our machines tend to break down, and so we need to act to repair them. All this gives us three choices of how to use our knowledge of those of the world's regularities that produce bad effects. We can use them deliberately to produce bad effects; we can actively inhibit their operation; or we can simply not bother to do anything, which will lead to some bad effects (but not as many or as bad effects as when we deliberately use the regularities to cause bad effects). If we learn that it is dangerous to swim in some river, then either we can try to stop people swimming there or we can encourage our enemies to swim there, or we can not bother to do anything, which will leave it quite probable that some people will drown. This greater range of choice available in a decaying world allows us to do wrong either maliciously or through negligence. The temptation to be negligent always exists, since we are all subject in the end to a strong desire to do nothing—sloth. To do wrong by negligence is not as bad as to do wrong through malice, and so we have the choice of doing wrong in ways of different seriousness. The choice of using or not using natural processes for good or ill exists once we know what these processes are and how so many of them produce bad consequences. Others of these processes we do not yet understand, and again we have a greatly significant choice—of investing time and money in attempting to understand these processes, or of not bothering. Our world is a world in which *A*'s pain gives reason for *B*'s research, which, in cooperation with *C*, using money provided by *D*, may lead to discovery of the cause of the pain, which *E* then produces a drug to relieve, with money provided by *F*.

Our actions of helping others are reinforced by the good desires for their well-being, to which I have drawn attention earlier. But there is a point in such desires not always occurring in sufficient strength in the relevant circumstances, so that people sometimes have the heroic choice of forcing themselves to do what is good

when they have little natural inclination to do so. They can then by such a choice manifest in a particularly full way their dedication to the good. Further, if desires are not automatic, they can be cultivated or repressed; and that gives agents an additional way of controlling their character and also thereby influencing their fellows. Our world is, of course, one in which such control is possible. Love that comes naturally in some circumstances can be encouraged or inhibited. We can cultivate love for a child by doing things with it, or inhibit affection for a woman by avoiding her and consorting with other women instead. (Such methods are not, of course, infallible, but they often work. We do have limited control over such desires.) Desires for specific ends are helped by more general emotions of affection and compassion. And, as we can cultivate good emotions, so we can allow bad ones, such as envy and jealousy, which come to us passively, to develop. They are just another aspect of our situation of temptation, having bad desires, necessary if we are to be moral beings.

Because of the opportunities that others have for harming us (as well as for other reasons to which I will come in the next chapter), inevitably things may go badly wrong with us. Our plans may fail, our beloved may be taken from us by circumstances or by some other agent. How will a God arrange for us to react to such misfortune? He could, of course, make us turn our minds straight away to other things, make us have no feeling about things going wrong. But surely the world is better if agents pay proper tribute to losses and failures, if they are sad at the failure of their endeavours, mourn for the death of a child, are angry at the seduction of a wife, and so on. Such emotions involve suffering and anguish, but in having such proper feelings a person shows his respect for himself and others. A man who feels no grief at the death of his child or the seduction of his wife is rightly branded as insensitive, for he has failed to pay the proper tribute of feeling to others, to show in his feeling how much he values them, and thereby failed to value them properly—for valuing them properly involves having proper reactions of feeling to their loss. And only a world in which people feel sympathy for losses experienced by their friends is a world in which love has full meaning. Once again a God would not necessarily make the expression of such emotions inevitable. It would be good if we have the opportunity to cultivate or repress them. Other emotions that have an appropriate place are emotions of regret and penitence at our wrong actions; emotions

also that it would be good if agents have the opportunity to cultivate or repress.

I have been arguing for the last few pages that a God would have reason to make a world in which agents have the opportunity to benefit or harm each other. There are a variety of different possible worlds, according to the time scale and nature of the mutual dependence involved. First you could have what I shall call a World-I. Here there would be an unchanging set of immortal humanly free agents. The world and its inhabitants would have their evils and imperfections, but the world would be perfectible by the cooperation of agents within a finite time. By the world being perfectible I mean that all evils would be removable and the world and society could be brought to a very happy and beautiful state such that no effort of agents could make it more beautiful. For the rest of eternity all they would need to do would be to keep it ticking over. There would be reason for God to make such a world—the happiness of agents is a good thing, and each agent (after enduring a little) would be able to attain it in such a world. But the trouble with such a world is that, after a finite time, agents would have nothing demanding to do. Although they could always tinker with the world, it would (once their initial labours were completed) as a result of the tinkering be no better a world than it was before; and, having attained much knowledge, agents would realize this. For this reason God seems to me to have more reason to make what I shall call a World-II. Here once again there is an unchanging set of immortal humanly free agents, but in this case there is an infinite number of such agents, and there would be no limit to the extent to which they could go on improving each other and their world. There would be an infinite amount of world and of agents to improve, infinite knowledge and freedom to acquire; and all this in World-II would need an infinite time. However, by the very description of World-II, one good thing would be barred to agents—giving birth to new agents and forming them from the beginning of their existence. Clearly it is good that agents should have such power. In what I call a World-III the number of immortal agents can be increased through the activity of existing agents. If a God has reason to make a World-II, he has, *a fortiori*, reason to make a World-III.

Birth is fine, but what about death? Does a God have reason to make a world in which, either by natural causes or by the action of agents, there is death? I believe that he does have a number of reasons

to make mortal agents. The first is that, if all agents are immortal, there is a certain kind of bad action (of a qualitatively different kind to other bad actions) that agents cannot do either to themselves or to others—they cannot deprive of existence. However much I may hate you or myself, I am stuck with you and me. And in this vital respect humanly free agents would not share the creative power of God. In refusing them this power, a God would refuse to trust his creatures in a crucial respect. To let a person have a gun is always a mark of profound trust. Secondly, a world without death is a world without the possibility of supreme self-sacrifice and courage in the face of absolute disaster. The ultimate sacrifice is the sacrifice of oneself, and that would not be possible in a world without death. ('Greater love hath no man than this, that a man lay down his life for his friends.') Supreme generosity would be impossible. So too would cheerfulness and patience in the face of absolute disaster. For, in a world without death, the alternatives would always involve continuance of life and so too the possibility that others would rescue one from one's misfortunes. There would be no absolute disaster to be faced with cheerfulness and patience.

Thirdly, a world with natural death would be a world in which an agent's own contribution would have a seriousness about it because it would be irreversible by the agent. If I spend all my seventy years doing harm, there is no time left for me to undo it. But, if I live for ever, then, whatever harm I do, I can always undo it. It is good that what people do should matter, and their actions matter more if they have only a limited time in which to reverse them. Fourthly, a world with birth but without natural death would be a world in which the young would never have a free hand. They would always be inhibited by the experience and influence of the aged.

The greatest value of death, however, seems to me to lie in a fifth consideration, which is in a way opposite to my second one. I wrote earlier of the great value that lies in agents having the power to harm each other. Only agents who can do this have significant power. Yet, for the sake of the potential sufferer, there must be a limit to the suffering that one agent can inflict on another. It would, I believe that we would all judge, be morally wrong for a very powerful being to give *limitless* power to one agent to hurt another. Giving to agents the power to kill is giving vast power of a qualitatively different kind from other power. It is very different from a power to produce endless suffering. Clearly the parent analogy suggests that it would

be morally wrong to give limitless power to cause suffering. A parent, believing that an elder son ought to have responsibility, may give him power for good or ill over the younger son. But a good parent will intervene eventually if the younger son suffers too much—for the sake of the younger son. A God who did not put a limit to the amount of suffering that a creature can suffer (for any good cause, including that of the responsibility of agents) would not be a good God. There need to be limits to the intensity of suffering and to the period of suffering. A natural death after a certain small finite number of years provides the limit to the period of suffering. It is a boundary to the power of an agent over another agent. For death removes agents from that society of interdependent agents in which it is good that they should play their part. True, a God could make a temporal limit to the harm that agents could do to each other without removing them from each other's society. But that would involve agents being in mutual relation with each other while being immunized from responsibility for each other, as well as depriving them of the possibility to make the choices of great seriousness described above.

I conclude that God would have reason to make what I shall call a World-IV. In a World-IV agents are born and die and during their life give birth, partly through their own choice, to other agents. They can make a difference to the world; but there is endless scope for improvement to it, and each generation can only forward or retard its well-being a little. Agents can make each other happy or unhappy, and can increase or decrease each other's power, knowledge, and freedom. Thereby they can affect the happiness and morality of generations distant in time. Our world is clearly a World-IV. A God has reason for making such a world.

In it there is the possibility of agents damaging each other over a number of generations until they fall badly down the ladder of ascent to divinity. Many, perhaps all, of the large-scale moral evils of recent centuries—the genocides, hatreds, large-scale enslavements, such as the Gulag Archipelago, the Slave Trade, and the Holocaust—are not the results solely of a wicked decision of some modern leader. They are the product of innumerable acts of individual members of one group against members of another group; which creates a climate in which some leader can motivate some others to institute the larger-scale evil, and yet others can ignore its commission.

But also in a World-IV there is the possibility of man's gradual ascent up the evolutionary scale, of man gradually developing his moral and religious awareness, and of each generation handing on to the next some new facet of that awareness. Man may grow in understanding moral truths and in applying them to the care of the less fortunate; he may grow in sensitivity to aesthetic beauty and in the creation and appreciation of works of art; in the acquisition of scientific knowledge and in its application to the betterment of the human condition and to the exploration and comprehension of the universe. Given that, as I argued in Chapter 6, we may expect God to create humanly free agents with a large degree of free choice and responsibility, subject to a limit of harm (that is, positive evil) that they can do to each other, it is moderately probable that God will make a World-IV, including natural death for all and free agents having the power to cause death. Yet such a world is obviously a very unsatisfactory one in the crucial respect that lives capable of flourishing happily for years to come, if not for ever, are cut short, deprived of future experiences and choices. God would have reason to intervene in the process to preserve in existence in some other part of this world agents who cease to exist in our part (and of course Christian and other forms of theism claim that he does so intervene). But, if the advantages of a world with death are to remain, the evident mutual interdependence in this world must cease after a finite period (to give a limit to the suffering allowed herein) and the future existence must in no way be foreknown for certain by agents (else there would be no opportunity in our part of the world for certain choices of great seriousness). If God did intervene in this way, our part of the world would still be, as far as appears to its inhabitants, much like a World-IV.

The Place of Animals

So far, this chapter has been concerned with the nature and circumstances that a God would have reason to give to humanly free agents. However, I suggested in Chapter 6 that God would have reason also to create conscious beings unburdened with the mixed blessing of moral awareness and free will—that is, the higher animals. I suggested that it was good that there should be beings who learn what is to their good and harm and seek the good and avoid the harm, and

through doing so enjoy the world and have pleasurable sensations. The results of this chapter about the nature that God would have reason to give to humans and the circumstances in which he would have reason to place them have limited application to the higher animals. As they lack free will and moral awareness, God would have reason to give to them less responsibility, and hence less power and knowledge than he would give to humans. For a substantial reason for giving great responsibility to humans, and with it the power and knowledge of how to produce great harm, was that they were not predetermined in their choices but were independent sources of how the world was to be, aware to some extent of the goodness or badness of alternative choices. Since animals lack free will, they must not be given such power to do harm; and in our world we find that they do not have such power. Animals can hurt and kill other animals, but they cannot indulge in genocide or atomic warfare, in giving perjured evidence in court or breaking solemn undertakings, in unjust imprisonment of their fellows for long periods, or in the subtler forms of torturing or humiliating their fellows. There are, however, reasons that I shall discuss in the next chapter for giving animals limited power to hurt. One is the higher-order virtues that it allows them to exercise. Another is in terms of the knowledge that it gives to other animals and to humans of how to avoid harm. I shall argue in the next chapter that one can acquire knowledge only through rational inference and inquiry of how harm is to be avoided by experience of the circumstances in which it occurs. With that knowledge, and with that alone, animals can do many of the things that we saw in this chapter it is good that they should be able to do—preserve themselves and their offspring, and perhaps others of their species also, from harm. When the gazelle sees other gazelles killed by tigers, that gives it knowledge of how to use its power for good ends—for its own benefit and that of its offspring—by escaping from other tigers and helping its offspring to escape from them.

In general, God has reason for giving to animals some power and knowledge, with opportunity to use it for good ends—for their own benefit and for that of others, and knowledge of how to do this. Clearly animals are in this position. Like humans, they have the opportunity to continue their existence by escaping from prey, walking in safe places, eating, drinking, and caring for health. Lacking free

will, they cannot choose whether to pursue these ends. Hence, if they are to do so, they must be given the desires to do so, including desires to get out of circumstances fatal to their well-being—and that means, as we saw earlier, biologically useful pain. The higher animals, like humans, also have the opportunity to care for others, notably their offspring, their mates, and to a lesser extent members of their own species. The care of offspring is a very central activity in the lives of animals. Again, since they lack free will, they cannot choose whether to pursue these ends, but they must be given the desire to do so. And, of course, they are in general abundantly blessed with the parental instinct. There are exceptions—animals who are bad parents—and we will come to those in the next chapter.

God could have put animals in a separate world from ourselves, or he could have put them in our world and given them as an additional responsibility for humans. For hundreds of millions of years there were on earth animals and no humans. But latterly there have been both. There is a clear risk in having both belong to the same world—that humans will abuse animals, causing them much suffering or sorrow. But there are advantages. For both humans and animals there is opportunity for responsibility and a new kind of cooperation and friendship. The horse, for example, has the opportunity to do interesting work together with man that he could not do apart from him, and friendship with a different kind of being. In return man has a helper and a friend. In seeing how much the welfare of animals now depends on him, man has come to realize that animals are his responsibility. And the responsibility for animals is not merely a short-term responsibility for individual animals whom we own or come across. It is now recognized as a responsibility for the preservation of animal species, and ensuring that there are environments suitable for animals of different species where they can be happy. (The doctrine of Genesis 1: 28 is that God gave to man this responsibility.)

As well as having this responsibility for animal welfare, humans do, of course, have responsibility for the physical world—for the plants, rivers, and rocks of this earth and of the moon; and in due course no doubt for the plants, rivers, and rocks of other planets and their moons. It lies within the power of humans to preserve species and large regions of natural beauty; and we have the choice of whether to do so or not to bother.

Conclusion

I argued in Chapter 6 that there was a significant probability to which I gave the somewhat artificial value of $\frac{1}{2}$ that a God would create humanly free agents—that is, beings who could choose how to make important differences to themselves, each other, and the world. I commented, then and subsequently, that such free choice brings with it significant evils; and, of course, the greater the range of choice, the greater the extent of actual and possible evils. Clearly there is a limit to the temporal extent and the intensity of the suffering that a good God would allow conscious beings to undergo (except through their own choice). And there is abundant room for dispute about whether this world contains more and more intense evils than a good God would allow, and we will come to that issue in the next chapter. But, if we suppose for the moment that there are not too many or too intense such states, it would seem good that the world contains humanly free agents, who have great power over themselves and each other and the physical world and great power to increase their power, knowledge, and freedom. These agents are ourselves, human beings. If God had given us more power, the resulting evil would plausibly have been too great. On the assumption (which, I have argued, is probable) that we do have free will, God could hardly have been more generous. If there is a God, we would expect a world containing such creatures with significant probability. And it is not unexpected that God would create unfree creatures with lesser powers and knowledge, the higher animals.

There are, however, many other worlds, which, if there were no God, would be as likely to come into existence as this one, characterized by very different general features. To take crucial examples—there could be worlds in which agents with the nature analysed in Chapter 9 cannot improve their characters, are not at great risk from their environment, cannot greatly improve their knowledge of how the world works, do not depend on or love each other (above all do not depend for their life on each other). We might all live in shells, unable to talk to each other or to hurt each other, merely making contact at the moment of reproduction (if then). Or, although we might be able to hurt each other physically, our characters might be unalterable, dependent entirely on our genes. Or we might be unable to make any long-term changes to our race and the environment in which we live. Or we might live for ever and have the power to cause

each other to suffer endlessly. And so on. In most of these worlds humans would not have much responsibility, and in the latter world they would have too much. That the laws and boundary conditions of the world are such as to endow humans with such natures and put them in such circumstances as to give them the sort of responsibility described in this chapter is something ‘too big’ for science to explain. Once again it depends on features from which science starts. And, in view of the diversity of possible worlds into which agents with the nature analysed in Chapter 9 could have been put, there does not seem to be any very great intrinsic probability that the laws and boundary conditions should be such as to give them the nature and circumstances of life described in this chapter. But, if I am right, a God may be expected with significant probability to produce a world of our kind and so to make the boundary conditions and natural laws such that they led to a world of this kind. The features of human beings and the world described in this chapter thus constitute further evidence for the existence of God. With k as the existence of human bodies connected with a mental life of the kind described in the last chapter, h as ‘there is a God’, and e as laws of nature and boundary conditions of the universe and the mind–body connections being such as to bring about the features described in this chapter, I claim that $P(e|h \& k) > P(e|k)$.

The Problem of Evil

I argued in Chapter 6 that God has substantial reason to create humanly free agents. They will have bodies of the kind analysed in Chapter 8, a mental life as analysed in Chapter 9, and something like the nature and circumstances of life considered in Chapter 10. To all appearances, humans are humanly free agents. In a world of the kind described in Chapter 10, there will inevitably be evils in the sense of bad states or actions.¹ I divide the world's evils in the traditional way into moral evils (those brought about by human intentional choice, or knowingly allowed to occur by humans, together with the evils of their intentional bad actions or negligence) and natural evils (all other evils, such as bad desires that we cannot help, disease and accidents). If there is a God, he permits moral evils to occur, and apparently himself brings about natural evils (through creating the natural processes that cause bad desires, disease, and accidents.) The moral and natural evils include animal pain, either caused by humans or by other animals or natural processes. However, since brain complexity and sophistication of behaviour decrease as we move further away from humans down the evolutionary scale, it seems reasonable to suppose that animal pain is less intense than human pain, and that animals feel pain less, as we go down the evolutionary scale from the primates to the least developed vertebrates. And, since the brains of invertebrates are of a different kind from those of vertebrates, I see little reason for supposing that the latter suffer pain at all.

¹ Chapters 10 and 11 summarize my arguments in *Providence and the Problem of Evil* (Clarendon Press, 1998), where I discuss the problem of evil at the length of a whole book. I have, however, here once again altered and, I hope, improved my account of 'the argument from the need for knowledge'.

I have already given reason for supposing that, in a world providential in the ways described in Chapter 10, there will be evils of certain kinds. There will inevitably be biologically useful unpleasant sensations, such as the pain that someone suffers until he escapes from a fire or the feeling of suffocation that one gets in a room full of poisonous gas, and the emotion of fear in dangerous circumstances. Also, since humans have the power to do each other significant hurt and they are not causally determined to do what they do, it is vastly probable that in such a world there will be a lot of further suffering, inflicted by humans on each other. And there will also be the moral evil of people choosing to do what they believe to be wrong, in inflicting such suffering; an evil that will exist even if they are not successful in inflicting the suffering. There will be the evil of bad desires, temptations to do wrong, whether or not we yield to them, that make possible the choice between good and evil. And when bad things happen to us or are done by us, or good things come to an end, there will be feelings of grief and sympathy and regret. But it might appear that most of the world's natural evils are in no way necessary to secure the good purposes described so far. And it might seem to some that, although humans having a free choice between hurting or benefiting each other logically requires the possibility of humans being caused by others to suffer, a good God would not be justified in permitting these moral evils for the sake of the good that the possibility of their occurrence unprevented by God makes possible.

Clearly, however, sometimes perfectly good people will allow evils to occur when they could easily prevent them. For sometimes some greater good can be achieved only by a route that involves suffering; and it is right to try to achieve it despite the suffering. Parents rightly allow children to suffer pain in the dentist's chair for the sake of the resulting healthy teeth. But God, unlike human parents, could produce healthy teeth without the need for the pain of dental surgery. Yet, as we saw in Chapter 5, even God cannot do the logically impossible. And that makes it plausible to suppose that a perfectly good God may allow an evil *E* to occur or bring it about if it is not logically possible or morally permissible to bring about some good *G* except by allowing *E* (or an evil equally bad) to occur or by bringing it about. I suggest that there are three further conditions that must be satisfied if, compatibly with his perfect goodness, God is to allow an evil *E* to occur. The second condition is that God also in fact brings

about the good *G*. Thirdly, God must not wrong the sufferer by causing or permitting the evil. He must have the right to make or permit that individual to suffer. And, finally, some sort of comparative condition must be satisfied. It cannot be as strong as the condition that *G* is a greater good state than *E* is a bad state. For obviously we are often justified in order to ensure the occurrence of a substantial good in risking the very unlikely occurrence of a greater evil. A plausible formal way of capturing this condition is to say that the expected value of allowing *E* to occur—given that God does bring about *G*—must be positive; that it is probable that the good will outweigh any evil necessary to attain it. I shall summarize the claim, with respect to some evil *E* that, if there is a God, he could, compatibly with his perfect goodness, allow it to occur in order to promote a good *G*, as the claim that *E* serves a greater good. I shall illustrate what these conditions amount to by examples, and defend the thesis that all the world's evils do probably serve a greater good—at least if we add to theism one or two further hypotheses. My treatment of evils, until the final section of this chapter, is concerned only with those evils that would be evils whether or not there is a God. I discuss certain states that would be evils only if there is a God in the final section entitled 'The Argument from Hiddenness'.

How Evils Serve Greater Goods

I begin with the first condition. This is evidently satisfied in the case of moral evil, as I have pointed out earlier. If humans are to have the free choice of bringing about good or evil, and the free choice thereby of gradually forming their characters, then it is logically necessary that there be the possibility of the occurrence of moral evil unprevented by God. If God normally intervened to stop our bad choices having their intended effects, we would not have significant responsibility for the world. And, as pointed out earlier, for us to have a free choice between good and evil we must (of logical necessity) have some temptation to do the evil. Hence the natural evil of bad desires. But what of the major natural evils of disease and accident?

I must begin by commenting briefly on three well-known but imperfect defences that theists have offered to the problem of natural evil. First there is the defence that much of the evil suffered by a human being is God's punishment for his sins; such punishment is a

good, and suffering is necessary to achieve it. Although this might account for some natural evil, it is clearly quite unable to account for the suffering of babies or animals. Secondly, there is the claim that God ties to the choices of some humans the well-being of humans (and animals) of later generations by routes other than normal causal processes (such as that of the influence we have over our children). God gives to our ancestors a great responsibility for our well-being. If they behave well, we flourish. If they sin, we suffer for their sins. The good of their responsibility, it may be claimed, requires the possibility of our suffering when they abuse it. Again, although this second defence might account for some natural evil, it clearly does not provide a satisfactory account of all such evil. There is the major difficulty that the good of agents having a choice between good and evil depends on their knowing the good and evil consequences that follow from their different actions. It is implausible to suppose that our early ancestors had any conception that their actions might cause the range of natural evils that their descendants suffer. And anyway this defence cannot explain the suffering of animals long before humans arrived on earth.²

More substantial is the third defence, used by many theistic writers down the centuries,³ that natural evils have been brought about by free agents other than humans—namely, fallen angels. If there is reason, as I argued that there is, for allowing humanly free agents to hurt other agents, then there is reason for allowing free agents other than humans to inflict such hurt—if, as may be postulated for angels, but cannot be believed with plausibility about our ancestors, they are agents with significant freedom and moral awareness⁴ fully aware of the consequences of their actions. This defence, unlike the first two, is adequate to cope with natural evils of all kinds. But it does have the major problem that it saves theism from refutation by adding to it an extra hypothesis, a hypothesis for which there does not seem to me to be much independent evidence—the hypothesis that angels of this

² Both of these defences were rejected by Jesus, according to St John's Gospel (John 9: 3), as accounts of why one particular person was born blind.

³ This defence has been used recently by, among others, Alvin Plantinga. See his *The Nature of Necessity* (Clarendon Press, 1974), 191–3.

⁴ Given the traditional view (see Chapter 6 n. 8) that angels have a fixed character (good or bad) as a result of one free original choice, we would have to suppose that they made that choice in the knowledge that God had promised them limited power over his creation, and that their free choice of character involved a choice of how they would use that power.

kind exist created by God and have limited power over the rest of God's creation. This hypothesis is not entailed by theism, nor does theism make it especially probable; any need for God to create creatures with free choices between good and evil that make great differences to other creatures is satisfied by the creation of human beings. A hypothesis added to a theory complicates the theory and for that reason decreases its prior probability and so its posterior probability. I shall, however, argue that we do not need to add this additional hypothesis to theism, since there are two other substantial reasons why, without natural evils, our ability to make the significant free choices that 'the free-will defence' rightly sees as such a good thing would be gravely diminished. That is, God has himself two substantial reasons for bringing about natural evils; and so we do not need to postulate that fallen angels are responsible for them. But I shall be arguing later that we do need to add one or two different further hypotheses to theism in order to justify the claim that God has the right to impose the degree of suffering that some individuals suffer.

The first of the substantial reasons why our ability to make significant free choices would be gravely diminished in the absence of natural evil is what is known as the 'higher-order' defence. This claims that natural evil provides opportunities for especially valuable kinds of emotional response and free choice. It begins by pointing out that the great good of compassion (the natural emotional response to the sufferings of others) can be felt only if others are suffering. It is good that we should be involved with others emotionally, when they are at their worst as well as at their best. But of course, the objector will say, even if pain is better for the response of compassion, better still that there be no pain at all. Now obviously it would be crazy for God to multiply pains in order to multiply compassion. But I suggest that a world with some pain and some compassion is at least as good as a world with no pain and so no compassion. For it is good to have a deep concern for others; and the concern can be a deep and serious one only if things are bad with the sufferer. One cannot worry about someone's condition unless there is something bad or likely to be bad about it. If things always went well with someone, there would be no scope for anyone's deep concern. It is good that the range of our compassion should be wide, extending far in time and space. The sorrow of one in a distant land who really cares for the starving in Ethiopia or the blinded in

India or the victims of carnivorous dinosaurs millions of years ago is compassion for a fellow creature, even though the latter does not feel it; and the world is better for there being such concern.

The 'higher-order-goods' defence goes on to provide the first reason why natural evil makes possible significant free choices. It points out that certain kinds of especially valuable free choice are possible only as responses to evil. I can (logically) show courage in bearing my suffering only if I am suffering (an evil state). I can 'show' sympathy for you (a phrase that I shall use to designate performing an action as opposed to having the passive feeling of compassion), and help you in various ways, only if you are suffering and need help. If I make the wrong choices and I ignore or laugh at your suffering, or indulge in self-pity at my own suffering, further possibilities of choice become available—you can resist the impulse to resent my lack of sympathy, or try gently to encourage me to show courage rather than self-pity. More generally, each bad action or state of affairs gives to the victims, perpetrators, and observers a free choice of how to react by actions, good or bad (as also does each good action and state of affairs). It is good that we should have the opportunity (occasionally) to do such actions as showing courage or sympathy, actions that often involve resisting great temptation, because thereby we manifest our total commitment to the good. (A commitment that we do not make when the temptation to do otherwise is not strong is not a total commitment.) Help is most significant when it is most needed, and it is most needed when its recipient is suffering and deprived. But I can (logically) help others who are suffering only if there is the evil of their suffering. In these cases, if there is a God, he makes possible the good of free choices of particular kinds, between good and evil, which—logically—he could not give us without allowing the evils (or evils equally bad) to occur. Or rather, it is the only morally permissible way in which he could give us the free choice. For God could make a basically deceptive world in which other people appeared to be in great pain when really they were not. Then we could have the choice between helping them or not helping them (or at any rate the choice of trying to help them or refusing to do so). But it would not be morally permissible—in my view—for God to make a world where people are moved to help others at great cost when the others do not really need help at all. God, if he is not to deceive us and yet give us a real free choice between helping and not helping others, must make a world where others really do suffer.

But could not the absence of a good (of an ability to walk, say, or the ability to talk French) give to the victim equal opportunity: whether to endure it with patience, or whether to bemoan his lot; and to friends, whether to show sympathy or whether to be callous? To answer this question, it is important to consider why pain is a bad state and so, if uncaused by humans (and not negligently allowed to occur by them), a natural evil. Pain is a sensation of a kind that we do not dislike if we have it in a very weak degree; indeed, we may often like it—we may like the sensation of warmth, which we dislike if it gets a lot stronger and becomes a sensation of great heat. And there are a few abnormal people who appear not to dislike the sensations that we call ‘pains’ at all. A sensation is a pain and so a natural evil only in so far as it is strongly disliked. Any state of affairs not caused (or negligently allowed to occur) by humans, disliked as strongly, would be just as bad. Some people dislike their disabilities just as much as they dislike pain; they so dislike their inability to walk that they will undertake a programme to conquer it that involves their ‘overcoming the pain barrier’. True, it would be unusual for anyone to dislike anything quite as much as some of the pains caused by disease or accident (and to call those pains ‘intense’ just is to say how much they are disliked). And, for that reason, pain normally provides more opportunity for evincing patience rather than self-pity than does anything else. But any state of affairs disliked as much would be equally bad and so provide as much opportunity. And the choice between being sympathetic rather than callous matters more than it does if the suffering is less. If the absence of the good is not disliked nearly as much as the sensations caused by disease and accident, then, of course, it is still very good to show courage in bearing that absence, but the courage is not in the face of such strong dislike for the existing state of affairs.

It may, however, be suggested, secondly, that adequate opportunity for the higher-level good emotions and especially valuable kinds of free choice would be provided by the occurrence of moral evil without any need for suffering to be caused by natural processes. You can show courage when threatened by a gunman, as well as when threatened by cancer; and show sympathy to those likely to be killed by gunmen as well as to those likely to die of cancer. Yet just imagine all the suffering of mind and body caused by disease, earthquake, and death not immediately preventable by humans removed at a stroke from our society. No sickness, no unavoidable diminution of powers

in the aged, no birth deformities, no madness, no accidents. Then, not merely would none of us have the opportunity to respond with sympathy or courage or reforming zeal, etc. directly, but so much of the oppression of one group by another stimulated by such suffering would also be removed. Starvation and disease in one group have so often served as the triggers leading to their oppression of another group whose good things they seek to possess. (That is, those natural evils have so strengthened the desires of the former for food and easier living that they have yielded to them, despite their dim awareness that they had no right to oppress the other group.) Then so many opportunities for coping with difficult circumstances would have been removed that many of us would have such an easy life that we simply would not have much opportunity to show courage or indeed manifest much in the way of goodness at all. It needs those insidious processes of (currently) unavoidable accident and dissolution that money and strength cannot ward off for long to give us the opportunities, so easy otherwise to avoid, to become heroes. True, God could compensate for the absence of natural evil by subjecting humans to such temptation deliberately (or at any rate knowingly) to cause suffering to each other that there was again plenty of opportunity for courage. He could make us so naturally evil that we lacked much natural affection and had inbuilt urges to torture each other (or at any rate allow each other to suffer), in face of which we could show courage and sympathy. But it is, I hope, in no way obvious that it would be better for God to replace disease by such an increase of inbuilt depravity (that is, a system of strong desires for what is known to be bad or to cause what is bad). Rather, I would have thought, the reverse. A world in which humans (and animals) lacked much natural affection for parents, children, neighbours, etc. would be a horrible place.

So, by bringing about the natural evil of pain and other suffering, God provides an evil such that allowing it, or an equally bad evil, to occur makes possible, and is the only morally permissible way in which he can make possible, many good states. It is good that the intentional actions of serious response to natural evil that I have been describing should be available also to simple creatures lacking free will. As we saw earlier, good actions may be good without being freely chosen. It is good that there be animals who show courage in the face of pain, to secure food and to find and rescue their mates and their young, and who show sympathetic concern for other animals. An

animal life is of so much greater value for the heroism it shows. And, if the animal does not freely choose the good action, it will do the action only because on balance it desires to do so; and, when its desire to act is uncomplicated by conflicting desires, the good action will be spontaneous. And (even if complicated by conflicting desires), animal actions of sympathy, affection, courage, and patience are great goods.

Yet an animal cannot go on looking for a mate despite failure to find it unless the mate is lost and the animal longs for it; nor decoy predators or explore the vicinity despite risk of loss of life unless there are predators, and unless there is a risk of loss of life. There will not be predators unless sometimes animals get caught. A hunt would be only a game unless it was likely to end in an animal getting caught and killed; and animals would not then be involved in a serious endeavour. And there will not be a risk of loss of life unless sometimes life is lost. Nor can an animal intentionally avoid the danger of a forest fire or guide its offspring away from one unless the danger exists objectively. And that cannot be unless some animals, such as fawns,⁵ sometimes get caught in forest fires. For you cannot intentionally avoid forest fires, or take trouble to rescue your offspring from forest fires, unless there exists a serious danger of getting caught in fires. The intentional action of rescuing despite danger simply cannot be done unless the danger exists and is believed to exist. The danger will not exist unless there is a significant natural probability of being caught in the fire; and to the extent that the world is deterministic, that involves creatures actually being caught in the fire; and to the extent that the world is indeterministic, that involves an inclination in nature to produce that effect unprevented by God.

True, the deterministic forces that lead to animals performing good actions sometimes lead to animals doing bad intentional actions—they may reject their offspring or wound their kin—and in this case the bad action cannot be attributed to free will. Nevertheless, such bad actions, like physical pain, provide opportunities for good actions to be done in response to them; for example, the

⁵ Those familiar with recent philosophical writing on the problem of evil will realize that I choose the example of a fawn caught in a forest fire because of its prevalence in that literature. This example was put forward by William Rowe ('The Problem of Evil and Some Varieties of Atheism', *American Philosophical Quarterly*, 16 (1979), 335–41) as an example of apparently pointless evil. I shall be pointing out in several places the good purposes that the fawn's suffering subserves.

persistence, despite rejection, of the offspring in seeking the mother's love or the love of another animal; the courage of the wounded animal in seeking food, especially for its young, despite the wound. And so on. The world would be much the poorer without the courage of a wounded lion continuing to struggle despite its wound, the courage of the deer in escaping from the lion, the courage of the deer in decoying the lion to chase her instead of her offspring, the mourning of the bird for the lost mate. God could have made a world in which animals got nothing but thrills out of life; but their life is richer for the complexity and difficulty of the tasks they face and the hardships to which they react appropriately.

The Argument from the Need for Knowledge

The second substantial reason why without natural evils, such as disease and accident, our ability to make significant free choices would be greatly diminished is that natural evils provide us with the knowledge required to make such choices. Natural evils are necessary if agents are to have the knowledge of how to bring about evil or prevent its occurrence, knowledge that they must have if they are to have a genuine choice between bringing about evil and bringing about good. Or rather, they are necessary if agents are to have this knowledge without being deprived of the good of rational response to evidence, and rational inquiry.

We saw in Chapters 6 and 8 that there need to be regular connections between an agent's bodily states and events beyond his body if he is to be able intentionally to perform mediated actions—that is, by his basic actions intentionally to produce effects beyond his body. But, if he is to acquire knowledge of how to perform these mediated actions by rational inference from observations of regularities in the world, and if he is to have the choice of whether to try to acquire this knowledge by rational inquiry (that is, by looking for such regularities), these regularities must be simple and observable, and the agent will need to extrapolate from what he observes by the criteria for a theory and so its predictions being probably true, as described in Chapter 3—what I shall call normal inductive inference. The simplest case of normal inductive inference is where I infer that a present state of affairs *C* will be followed by a future state *E*, from the generalization that, in the past, states of affairs like *C* on all occasions of which

I have knowledge have been followed by states like *E*. Because on the many occasions of which I have knowledge a piece of chalk being liberated from the hand has fallen to the floor, I can infer that the next time chalk is liberated it will fall. However, as we saw in Chapter 3, normal induction may take a more complicated form. From a vast collection of data about the positions of the sun, moon, and planets a scientist may infer a consequence of a different kind—for example, that there will be a very high tide on earth when the moon is in such-and-such a position. Here the data render probable a scientific theory of which the prediction about the high tide is a somewhat remote consequence: the similarities between the data and the prediction are more remote than in the simplest cases. (But the similarities exist and are the basis of the prediction. In both the data and the prediction there are material bodies attracting each other.)

Whether the normal inductive inference is simple or complicated,⁶ certain general points can be made about the claim to knowledge of the future that results from it. The first is that the more past data there are, the better established is a claim to such knowledge. This is because the data confirm a claim about the future by confirming a theory or a simple universal (or probabilistic) generalization (for example, ‘states like *C* are always followed by states like *E*’), which in turn licenses the claim about the future. The more data there are, the more they show that the theory or generalization holds in many different circumstances and so is more likely to hold in the future instance in question. (However similar in many respects are the circumstances under which the past data have been observed, they are almost bound to differ from each other in some observable or unobservable respects; if the generalization has worked so far despite many such differences, that gives it greater probability.) Secondly, the surer my knowledge that the past data occurred as stated, the better grounded is my claim to knowledge of the future. If the data are mental experiences of mine or events that I myself have seen, then my knowledge of their occurrence is sure. If they are experiences that others report or events that others claim to have seen, then my knowledge of their occurrence is less sure. My knowledge of their occurrence will be still less sure, if I need to make a complicated

⁶ Note that the complexity of an inferential process does not as such make its conclusion any less likely to be true; whereas the complexity of a postulated hypothesis does make it less likely to be true.

inference from other data to prove their occurrence. Clearly, in so far as an inference is licensed by certain data, then to the extent to which it is doubtful whether the data are correct, it is doubtful whether the inference is justified. Thirdly, in so far as the data are qualitatively rather dissimilar from what is predicted, and a complicated scientific theory is needed to generate the prediction, the claim to knowledge will be less surely based. Thus, suppose that by a process of complex extrapolation from a number n of astronomical data I reach a very complex theory of mechanics, from which I conclude that in a very unusual set of circumstances (when the planets are in just such-and-such configurations) if I let go of a bit of chalk it will rise into the air. And suppose that these circumstances are to be manifested on earth uniquely in my study during this hour. Do I *know* that when shortly I let go of the chalk it will rise? Doubtfully so. Clearly I do know it and know it a lot better if I have already actually let go of the chalk n times during the hour, and it has risen. Fourthly, if a complicated inference is needed in order to reach a prediction, then, in so far as the inference is of a type that has proved successful in the past or the inference is done by persons with known predictive success from this kind of work in the past, that is grounds for believing the prediction. These four points about the strength of knowledge obtained by normal induction may be summarized by saying that our claims to knowledge are better justified, the closer they are to our experience.

Now, if agents are knowingly to bring about states of affairs, or to allow states of affairs to come about through neglecting to prevent them, they must know what consequences will follow from their actions. Normal inductive knowledge of consequences, it follows from what has just been said, is to be obtained as follows. Consider an action A , which I am contemplating doing in circumstances X . Suppose that A consists in bringing about a state of affairs C , the result of A (see p. 35 for a definition of 'result'). How am I to know what its effects will be, what will follow from it? Most certainly, by having done such an action myself many times before in similar circumstances, and having observed the effects of its result. I could come to know most surely what will result from my drinking eight double whiskies—that I shall be unable to drive my car safely, by having done such an action often before. I know the effect less surely by having seen the effects of others doing the action, or by having seen the effects of the result of the action when this was brought

about unintentionally, all in similar circumstances to those in which I am considering doing the action or by others telling me what happened on different occasions when they drank eight double whiskies. I know that this will lead to inability to drive less surely, because I suspect that I am different from the others (have more will power, am more conscious of the dangers, am a better driver than the others anyway).

Less sure knowledge still is obtained by observing the result occur in somewhat different circumstances (for example, when drinkers drink the whiskies much more quickly, or when tired). Still less sure knowledge is obtained by having observed goings-on only somewhat similar, and having to make allowance for the difference—for example, I may only have seen the effects of people drinking different quantities of beer or gin. Or my knowledge may depend on reports given by others; then it will be still less certain. The witnesses may have exaggerated, not noticed differences in circumstances, etc. The least certain knowledge of all is that which is reached by a process of more complicated inference from goings-on only remotely similar to A. However, it is difficult to see how a theory that predicted the occurrence of an evil such as pain could have any justification unless the data on which the theory was built were cases of pain. If you had no knowledge of anything causing pain, how could other kinds of data substantiate predictions about pain? For pain is so different from other kinds of goings-on and has no natural connection with particular brain or nerve conditions rather than with others. (As I argued in Chapter 9, there is no reason for supposing that stimulation of this nerve will cause pain and of that one will cause pleasure, other than knowledge that that is what has happened in the past.)

So proximity to experience gives more certain knowledge. It is notorious that people are much more inclined to take precautions against some disaster if a similar disaster has happened to them previously or to those close to them than if they are warned of the need for precaution by some impersonal distant authority. Someone is far more inclined to take precautions against fire and burglary if she or her neighbours have previously been victims of fire or burglary than if the police warn her that these things have happened in the next village. My point is that this is not just irrational perversity. It is the height of rationality to be influenced more by what is known better. People know better that it can happen to them if they know that it has happened to them or to others like them. With a mere

police warning, they always have some reason for suspecting that police exaggerate or that things are different in the next village. What is irrational is not being influenced at all by the police warning; what is not irrational is being influenced more by goings-on closer at hand of which we have more intimate experience.

Now, for any evil that people knowingly inflict on each other, there must have been a first time in human history at which this was done. There must have been a first murder, a first murder by cyanide poisoning, a first deliberate humiliation, and so on. The malevolent agent in each case knows the consequence of the result of his action (for example, that causing someone to imbibe cyanide will lead to their death). *Ex hypothesi*, he cannot know this through having seen an agent give another person cyanide for this purpose. His knowledge that cyanide poisoning causes death must come from his having seen or others' having told him that on other occasions taking cyanide accidentally led to death. (If, in my example, you think that knowledge of the effects of imbibing cyanide might be gained by seeing the effects of taking similar chemicals, the argument can be put more generally. Some person must have taken previously a similar poison by accident.) What applies to the malevolent agent also applies to the person who knowingly refrains from inflicting evil on another or stops evil occurring to another. There must be naturally occurring evils (that is, evils not knowingly caused by humans) if humans are to know how to cause evils themselves or are to prevent evils occurring. And there have to be *many* such evils, if humans are to have sure knowledge, for, as we saw, sure knowledge of what will happen in the future comes only by induction from many past instances. A solitary instance of a person dying after taking cyanide will not give to others very sure knowledge that in general cyanide causes death—maybe the death on the occasion studied had a different cause, and the cyanide poisoning had nothing to do with it. And, unless people have knowingly been bringing about evils of a certain kind recently, there have to have been many recent naturally occurring evils if people are currently to have sure knowledge of how to bring about or prevent such evils.

To take another example—we know that rabies causes a terrible death. With this knowledge we have the possibility of preventing such death (for example, by controlling the entry of pet animals into Britain), or of negligently allowing it to occur, or even of deliberately causing it. Only with the knowledge of the effects of rabies are such

possibilities ours. But for us to gain knowledge by normal induction of the effect of rabies it is necessary that others die of rabies (when the rabies was not preventable by humans), and be seen to have done so. Generally, we can have the opportunity to prevent disease affecting ourselves or others or to neglect to do so, or the opportunity to spread disease deliberately (for example, by indulging in biological warfare), only if there are naturally occurring diseases. And we can have the opportunity to prevent incurable diseases or to allow them to occur, only if there are naturally occurring incurable diseases.

What applies to individuals in the short term applies in the longer term and to races. If humans are to have the opportunity by their actions or negligence to bring about evil consequences in the distant future, or to avoid doing so, they must know the long-term consequences of their actions, and the most sure inductive knowledge of those consequences can only come from past human history. How are we to have the opportunity to stop future generations catching asbestosis, except through knowledge of what causes asbestosis, and how is that to be obtained except through records that show that persons in contact with blue asbestos many years ago have died from asbestosis thirty years later? Or—to take a slightly different kind of example—suppose that people are to have the choice of building cities along earthquake belts, and so risking the destruction of whole cities and their populations hundreds of years later, or of avoiding doing so. How can such a choice be available to them unless they know where earthquakes are likely to occur and what their probable consequences are? And how are they to come to know this, unless (unpredicted) earthquakes have happened in the past, and the circumstances of their occurrence are studied and provide evidence making probable a theory of when earthquakes are likely to occur? And if humans in the past could not predict the earthquakes, it is highly probable that sometimes earthquakes would occur where humans have built cities. And so again (though in a slightly different kind of way), natural evil provides us with a wider range of choice of actions by which we can affect ourselves, each other, and the physical world.

What has happened to sentient creatures other than humans also gives knowledge of what will happen to us, though very much less sure knowledge. Indeed, a great deal of our knowledge of the disasters for man that would follow some action comes from study of the actual disasters that have befallen animals. For a long time it has been

normal to discover the effects of drugs or surgery or unusual circumstances on man by deliberately subjecting animals to those drugs or surgery or circumstances. Before putting humans into space, humans put animals into space and saw what happened to them. Such experiments do not give very sure knowledge of what would happen to humans—because from the nature of the case there are very considerable differences between animals and humans—but they do give considerable knowledge. The evils that have naturally befallen animals provide a huge reservoir of information for humans to acquire knowledge of the choices open to them, a reservoir that they have often tapped—seeing the fate of sheep, humans have learnt of the presence of dangerous tigers; seeing cows sink into a bog, they have learnt not to cross that bog, and so on. And the evils that provide information need not just be physical ones, and the ways in which they are produced may be by the actions of other animals not blessed with free will. The effects of bad parenting by gorillas may help us to see some of the effects of bad parenting by humans.

The scope for long-term choice available to future generations must not be underestimated. They may have the choice not merely of whether to build cities so as to avoid earthquakes, but of whether to drive the earth nearer to the sun or further from it, to take air and water to Mars and live there instead, to extend the life span, to produce new manlike organisms in laboratories, and so on. But rational choices on these matters can be made only in the light of knowledge of the consequences of alternative actions. While knowledge of the disasters and benefits that have befallen past human generations can be of great use in providing such knowledge, when we are considering the *very* long-term consequences of changes of circumstances, environment, or climate, the story of animal evolution provides our main information. Human history so far is too short to provide much useful knowledge. For example, the fossil record indicates that the earth's magnetic field periodically undergoes a sudden reversal every few hundred thousand years. (Magnetized parts of metal that previously pointed north thereafter point south and vice versa). We need to know how this will affect humans when it next occurs; and if its effects will be harmful, whether there are precautionary measures that we can take to minimize the harm. Our main evidence on the long-term effects of such a reversal can only come from the fossil record of how it has affected animals. But in any case the story of pre-human nature 'red in tooth and claw'

already provides some very general information crucially relevant to our possible choices. For suppose that animals had come into existence at the same time as human beings (for example, 4004 BC) always in situations where humans could save them from any suffering. Naturally it would then seem a well-confirmed theory that (either through act of God or nature) suffering never happens to animals except such as humans can prevent at that time. So humans would seem not to need to bother to take action now to prevent later animal suffering. But the story of evolution tells us that that is not so—animals may well suffer in circumstances when humans are not available or able to prevent this, because this has happened in the past. This knowledge gives us a choice of taking action now to prevent animal suffering in future, or of not bothering to do so. The story of pre-human evolution reveals to man just how much the subsequent fate of animals is in his hands—for it will depend on the environment that he causes for them and their genes, which he may cause to mutate.

Like earlier examples, the above argument illustrates the more general point that, if agents are to have knowledge of the evil that will result from their actions or negligence, laws of nature must operate regularly. Thus I mentioned in the last chapter that among the advantages of the pain caused by fire is that it leads the sufferer to escape from the fire. But the pain still occurs when the sufferer is too weak or paralysed to escape from the fire. Would it not be better if only those able to escape suffered the pain? But, if that were the case, then others would know that it mattered much less that they should help people to escape from fire and that they should prevent fire. And so the opportunities for humans to choose whether to help others and guard against their future sufferings will correspondingly diminish. And in general, if God normally helps those who cannot help themselves when others do not help, others will not take the trouble to help the helpless next time, and they will be rational not to take that trouble. For they will know that more powerful help is always available.

My argument so far has been that, if humans are to have the opportunity to bring about serious evils for themselves or others by actions or negligence, or to prevent their occurrence, and if all knowledge of the future is obtained by normal induction, that is by rational response to evidence—then there must be serious natural evils occurring to man or animals. I have argued earlier that it is good

that we should have the former opportunity. What of the possibility that God should give us the necessary knowledge by a different route?

Normal inductive inference from the past is not the only possible route to knowledge of the future. Why do we need to acquire this knowledge by rational response to evidence? Why could not God ensure that we simply found ourselves having true basic beliefs that this action would cause pain and that action would cause pleasure, for actions of various kinds and pleasures and pains of various kinds?⁷ A basic belief is one that we find ourselves having, not on the basis of inference from other beliefs, and from which we may infer to other things that we then come to believe. For example, for most of us, the immediate deliverances of perception—that I am looking at a tree, or listening to a lecture—come to us as basic beliefs. By the Principle of Credulity, which I defend in Chapter 13, all basic beliefs with which agents find themselves are—in the absence of counter evidence—probably true; the mere fact that you have a belief is grounds for believing it. This route to knowledge of the future would be inductive, but not use induction of the normal kind. Given that (for the good reason adduced on p. 226) our world is a world of decay, our basic beliefs would need to include beliefs about what will happen if we do nothing—for example, about when a disease epidemic will strike unless we begin a programme of inoculation. It would not, however, be possible for any of us to know with any reasonable certainty *all* the long-term consequences of our actions, since those long-term consequences depend on whether other free agents help or hinder our actions attaining the consequences that we intend. So the most that would be possible is for us to know those consequences that are independent of the actions of others, and also conditional consequences (for example, ‘if no one else interferes, action *A* will have consequence *C*’). But, if God gave us true basic beliefs about the consequences of all our actions subject to those restrictions, we would know what would be the whole future of the world if humans did not interfere with it, and what would happen if

⁷ A third way, beside that of normal inductive inference and providing such knowledge in the form of basic beliefs, by which God could provide us with knowledge of the consequences of our actions is by telling us himself what these consequences would be, and making it evident that he is telling us. But that would make his presence so evident to us that it would be open to all the disadvantages of that which I shall discuss shortly when I come to the ‘argument from hiddenness’, as well as depriving us of the opportunity for rational inference and inquiry.

they did interfere with it in various ways. And so, among the other things that we would know would be the outcomes of all the experiments we might do to attempt to confirm any scientific or metaphysical theory. We could still decide between competing theories on the basis of the *a priori* criteria of simplicity and scope. But the decision would be limited to a decision between theories that had exactly the same observable consequences as each other (even in the distant future); and in consequence the interest and importance of such a decision would be extremely low. For a major reason why some conclusion that a certain theory is more probable than some other is of great interest and importance is that the former makes predictions that the latter does not. But in the postulated situation we would not need to do science in order to know the future.

As things are in the actual world, most moral decisions are decisions taken in uncertainty about the consequences of our actions, even if we discount the possibility of interference by other agents. I do not know for certain that, if I smoke, I will get cancer; or that, if I refuse to give money to Oxfam, another person will starve to death. Maybe I will be one of the ones who does not get cancer, and maybe my failing to make my small gift to Oxfam will make no difference to the number of people who starve to death. For suppose that the only difference made by the absence of my gift is that each starving person gets an allocation of food a tiny bit smaller than what they would have got anyway; and I know that this will be the immediate effect of my action. But what I may not know is whether that difference of allocation is so small as to make no difference to the future condition of the starving. So we have to make our moral decisions on the basis of how probable it is that our actions will have various outcomes—how probable it is that I will get cancer if I continue to smoke (when I would not otherwise get cancer) or that someone will starve if I do not give (when they would not starve otherwise). These decisions under uncertainty are not merely the normal moral decisions; they are also the hard ones. Since probabilities are so hard to assess, it is all too easy to persuade yourself that it is worth taking the chance that no harm will result from the less demanding decision (that is, the decision that you have a strong desire to make). And, even if you face up to a correct assessment of the probabilities, true dedication to the good is shown by doing the act that, although it is probably the best action, may have no good consequences at all.

But if we are often in this situation (and for the above reasons it is good that we should be), then it is good (because we rightly seek to do good actions) that we should have the opportunity to obtain more certain knowledge of the consequences of our actions—that will involve getting more data about the consequences of events, for example, data from the past about what has happened to people who have smoked in ignorance of the possibility that smoking causes cancer. Seeking more certain knowledge, in other words, involves once again relying on normal induction.

Above all, if our knowledge of the consequences of our actions is limited, we have the all-important choice of whether or not to pursue scientific inquiry to extend our knowledge, and of teaching or not teaching others the results of such inquiry. The rationality that is necessary if we are to make serious moral choices is, quite apart from its value for this purpose, a great good in itself. One of the very greatest glories of humans is their ability to be responsive to evidence and reach probable conclusions about the effects of their actions, about how the world works, and about what is our origin and destiny. Rationality is a quality for which it is worth paying a considerable price. We rightly value greatly the scientist who investigates the causes and effects of things and who opens himself to applying objectively correct criteria to discovering how nature operates, and about which events cause pain and which cause pleasure. And it is a further glory of humans that they cooperate in the activity of reaching probable conclusions; some humans teach others, and the others build on those foundations. And humans have the choice whether or not to investigate, to cooperate in investigation, and to teach the results of investigation. To have these various serious moral choices, we need initially to be (more or less) ignorant of the consequences of our actions, for good or evil. The occurrence of natural evil gives us the choice of improving our knowledge of these consequences, which we cannot obtain in any other way without a serious loss of good.

The higher animals themselves also acquire knowledge by normal induction, knowledge of where to obtain food, drink, and fellowship; and also knowledge of the causes of pain, loss of health, and loss of life. While animals do not have the free choice of active investigation into causes and effects, the higher animals do learn (that is, acquire knowledge) by the use of objective criteria of what is evidence for what—by generalizing, the simplest kind of normal inductive

inference. Seeing the suffering, disease, and death of others in certain circumstances, they learn to avoid those circumstances. Seeing a fawn caught by a fire in a thicket, other animals learn to avoid the thicket. (Animals, and especially the lower animals, do of course avoid many situations and do many actions instinctively; but in those cases they do not do the action or avoid the situation through rationally acquired knowledge of its consequences.) And some of the higher animals (albeit unfreely) do investigate the consequences of actions before doing them. A cat often tests the strength of a tree branch before putting his full weight on it. Their rationality in this simple respect is one of their glories, another limited respect in which they have some similarity to humans. It is good that they too should save their lives and those of their offspring through learning the consequences of things by observation and inquiry. Yet other animals must suffer if some animals are to learn to avoid suffering for themselves and their offspring.

We may not know exactly when and where the past natural evils occurred, but the mere knowledge that suffering of a certain type occurred to certain kinds of creatures under certain conditions provides us with very good reason to avoid actions that may produce those conditions. Indeed, since all natural evils occur as a result of largely deterministic natural processes (there are no kinds of natural evil that occur in a totally random way), all such knowledge helps to build up knowledge of natural processes that we can utilize to produce or prevent future evils. All past and present human and animal natural evils of which we know thus contribute to the widening of human choice when we learn about them. And (except at a heavy cost) we could not learn, and especially choose to learn, without them. The great good of choice between good and bad acts cannot be had without knowledge of the consequences of our actions. If we were given this knowledge in the form of basic beliefs, we would be deprived of the great good of rational response to evidence, rational investigation, and the choice of whether to pursue it. To have all this knowledge provided for us would radically reduce the range of the choices we have. Natural evil makes it possible for us to make many more serious moral choices, including the choice of whether or not to pursue rational investigation.

In developing the higher-order-goods defence and the argument from the need for knowledge, I have been arguing that by bringing about natural evils God makes possible various goods and that doing

the former is the only logically possible and morally permissible way in which he can secure the latter. I have argued that theism can justifiably claim that this first condition for God to allow evils to occur is satisfied without the need to invoke any extra hypotheses such as the fallen-angel hypothesis discussed earlier. All the evils of our orderly world of which we know are ones to which we can choose to react in the right way—minimally by compassion and by learning from them, and often in many other ways as well. But, if God is to be justified in bringing about or allowing others to bring about these evils, he must also bring about the good that they make possible; humans who choose between good and evil must do so of their own free will. The bad nature and bad effects of human free choices being so much worse than the bad nature and effects of instinctive animal reactions, the free nature of their choices is, I suggest, needed to justify a good God allowing them to cause such evil as they can. I have argued briefly that there is no reason to deny that things are as they seem to be in this respect, and that humans do have the requisite free will. The second condition for God to be justified in bringing about natural evils is also satisfied.

God's Right to Inflict Harm

I have been arguing that, by permitting moral evil and bringing about natural evil, God gives us (and animals) a good that he could not give us in any other morally permissible way. But does God have the right to impose harm on us for the sake of a resulting good; and, in particular, does he have the right to allow some to suffer for the benefit of a good to others?

God as the author of our being would have rights over us that we do not have over our fellow humans. To allow someone to suffer for his own good or the good of someone else, one has to stand in some kind of parental relationship towards him. I do not have the right to force some stranger, Joe Bloggs, to suffer for the good of his soul or of the soul of Bill Snoggs, but I have *some* right of this kind in respect of my own children. I may let the younger son suffer *somewhat* for the good of his and his brother's soul. I have this right because in small part I am responsible for the younger son's existence, its beginning, and continuance; I feed him and educate him. I have the right to demand something in return, that he is open to the possibility of his

elder brother inflicting (limited) harm on him. If this is correct, then, *a fortiori*, a God who is, *ex hypothesi*, so much more the author of our being than are our parents, has rights so much greater in this respect.

It is for this reason that, it is not a good objection to my argument that, if pain serves various good purposes (such as providing opportunities for sympathetic and courageous action, and knowledge of when pain is caused), it would be good for us humans to cause more pain to each other and to animals. We are not the primary cause of the existence of other humans and animals, and so we do not in general have that right. God has that right, and in very specific circumstances and to a very limited degree humans in parental situations do also.

I suggest that we can generalize these intuitions by the following principle. A benefactor has the right to take back, or to tie bad aspects to, some of the benefits that he gives to some one, so long as he remains on balance a benefactor. God who gives so much has the right to take back something in return; though he who is perfectly good will never seek it unless to take it is logically necessary in order to provide some good to this beneficiary or to someone else.

One might feel that people ought to be asked whether they want to receive a benefit, especially one with bad aspects (for example, pain) attached. Surely no one has this right to inflict harm on some person for his greater good, let alone for the greater good of another, without that person's consent. We judge that doctors who use people as involuntary guinea pigs for medical experiments are doing something wrong. The all-important difference, however, is that the doctors *could* have asked the patients for permission; and the patients being free agents of some power and knowledge could have made an informed choice of whether or not to allow themselves to be used. God's choice is not about how to use already existing people, but about the sort of people to make and the sort of world into which to put them. In God's situation there are no people to be asked. Thus in the previous chapter I argued that it would be good that one person, *A*, should have deep responsibility for another one, *B*. Ought not God to have asked *B* if he wanted things thus? But this is not possible, for, if *A* is to be responsible for, *B*'s growth in freedom, knowledge, and power, there will not be a *B* with enough freedom and knowledge to make any choice before God has to choose whether or not to give *A* responsibility for him. The creator has to make the choice independently of his creatures, and he has a reason for choosing to

make them deeply interdependent. Again, God has reason, we saw, to create a world in which some suffer, to give others knowledge through rational inquiry. But humans cannot choose in what sort of a world by what route they are to acquire knowledge, for until they have acquired knowledge they cannot choose anything. God has to make the choice for them. And, if there is a God, he often pays us the compliment of assuming that, if we had the choice, we would want our lives to be of use to others.

It may seem that my condition that God may not impose on us evils that (equal or) outweigh the good things he gives us is not satisfied in many cases. Does it not seem that many humans live lives that are such that it was better for that human never to have lived? I believe that this is the case far less often than it seems initially for two reasons. The first is that it is a great good simply to be alive, even if life involves quite a lot of suffering, perhaps many periods in which the bad outweighs the good. I hope that a reader will not think me callous in making the comment that, if many people really thought that it were better that they had never lived, there would be many more suicides than there are. The objector may say that the reason why there are not more suicides than there are is that people feel obligations to their parents, children, spouses, etc. to stay alive. That is true, and they are surely often right to feel these obligations. But that brings me to the second reason why a judgement that it were better for some humans never to have lived is so often mistaken. It ignores the great good of being of use.

That helping is an immense good for the helper has always been difficult for humans to see, but it is especially hard for twenty-first-century secularized Western man to see. But just think how awful life would be if we were of no use, if we could not ever help others in some way or another. And most of us can see that sometimes: when, for example, we seek to help prisoners, not by giving them more comfortable quarters, but by letting them help the handicapped; or when we pity rather than envy the 'poor little rich girl' who has everything and does nothing for anyone else. And one phenomenon prevalent in modern Europe draws this especially to our attention—the evil of unemployment. Because of the systems of social security common in Western Europe, the unemployed on the whole have enough money to live without too much discomfort; certainly they are a lot better off than are many employed in Africa, Asia, or nineteenth-century Britain. What is evil about unemployment in

Western Europe is not so much any resulting poverty but the uselessness of the unemployed. They often report feeling unvalued by society, of no use, 'on the scrap heap'. They rightly think it would be a good for them to contribute; but they cannot.

It is not only intentional actions freely chosen, but also ones performed involuntarily, that have good consequences for others, that constitute a good for those who do them. If the unemployed were compelled to work for some useful purpose, they would still—most of them—regard that as a good for them in comparison with being useless. Or, if they would not so regard it, I suggest that most of us who are employed, and not directly involved in their plight, can see it as a good for them. Or consider the conscript killed in a just and ultimately successful war in defence of his country against a tyrannous aggressor. Almost all peoples, apart from those of the Western world in our generation, have recognized that dying for one's country is a great good for him who dies, even if he was conscripted.

And it is not only intentional actions but experiences undergone involuntarily (or involuntary curtailment of good experiences, as by death) that have good consequences—so long as those experiences are closely connected with their consequences—which constitute a good for him who has them (even if a lesser good than that of a free intentional action causing those consequences, and a good often outweighed by the evil of the experience in question). Consider someone hurt or killed in an accident that leads to some reform that prevents the occurrence of similar accidents in future (for example, someone killed in a rail crash, which leads to the installation of a new system of railway signalling that prevents similar accidents in future). His relatives often comment in such a situation that at any rate the victim did not suffer or die in vain. They would have regarded it as a greater misfortune for the victim if his suffering or death served no useful purpose. It is a good for us if our experiences are not wasted but are used for the good of others, if they are the means of a benefit that would not have come to others without them, which will at least in part compensate for those experiences. It follows from this insight that it is a blessing for a human (or animal) if the possibility of his suffering makes possible the good for others of having the free choice of hurting or harming him; and if his actual suffering provides knowledge for others, and allows others to feel compassion for him and gives to them the choice of showing or not

showing sympathy to him. Thus it is a good for the fawn caught in the thicket in the forest fire that his suffering provides knowledge for the deer and other animals who see it to avoid the fire and deter their other offspring from being caught in it. (I should make it clear here again that—for the reasons given earlier—I am *not* saying that humans have any right to cause train crashes or forest fires of the kind just described. God alone has that right.) It is much better if the being-of-use is chosen voluntarily, but it is good even if it is not.

Someone may object that the good is not (for example) dying for one's country, but knowingly dying for one's country when one believes it good to do so—having the experience of 'feeling good' that one is sacrificing oneself. But that cannot be right. It could be good to have the experience in question only if one's beliefs were correct. There would be nothing good about believing one was dying for a good cause, when in fact it was a lousy cause. To take an analogy, it is a good thing to rejoice that you have passed your exams only if it is a good thing (independently of whether or not you believe it) that you have passed your exams. Believing that it is good that something has happened cannot make it good that it has happened; it will be good to have the belief only if it is good anyway that that thing has happened. So, while believing truly that it is good that one is dying for one's country is a good, it will be a good only if dying for one's country is a good anyway (whether or not one believes it).

One consequence of all this is that, if someone refrains from committing suicide because he knows that someone who loves him will be greatly hurt thereby, he is very fortunate. To be loved, and for your life to be valuable to someone else, are an enormous benefit for you. And, if the would-be suicide does not see that, he is simply mistaken.

If there is a God, the greatest good of all in this respect must be being-of-use to God himself; and all human suffering that is of use to others will also be of use to God, who has designed the world so that suffering does benefit those others. Just as when a doctor helps a small child to use injured limbs, he benefits both the child and the child's parents, one of whose major goals in life is that the child shall flourish, so all human suffering that helps others is of use to God in forwarding his purposes. And one who is of use to the perfectly good source of all being is indeed fortunate.

If, however, when even this great good of being of use is taken into account, there are humans whose lives on earth are such that on

balance it would have been better for them never to have lived, then God has an obligation to provide them with enough (in quantity and quality) of a good life after death, so that their total life (on earth and hereafter) is on balance a good life. Being omnipotent, he can do this, and we have no reason to suppose that (if there is a God) he does not. There is nothing wrong in a benefactor providing a life that (within limits) in its early stages is on balance bad, so long as over its whole period it is on balance good. We can see that from the fact that doctors and parents are surely right to allow a foetus or very young child to suffer in consequence of some serious surgery or other medical intervention, so long as it is the means to a good life overall.

In this respect also what goes for humans goes for animals too. God must give to each of them on balance a good life; the goodness of their lives must outweigh evils in it. But for them too being of use to others—either to humans or to other animals or to God himself—is a great good, whether or not they recognize this. And I am inclined to think (though I may be mistaken) that, because the pains of animals are less than ours, when the great good of being of use is taken into account, God's obligation to provide for each animal a good life would be satisfied without any need for life after death.

However, if the evils of this world are such that, in order to save the hypothesis of theism from refutation, I need to add to it the hypothesis that God provides a period of good life after death for anyone whose life on earth is on balance a bad one (one such that it would be better if they had never lived), I am complicating theism and therefore decreasing its probability—in the same way as if I were to have added to it the fallen-angel hypothesis. I shall return to this point shortly.

Still, given this crucial assumption of a compensatory life after death for any whose lives on earth are on balance bad, I claim that God has the right to allow humans (and animals) to suffer for a limited period and to a limited extent. But, in view of the fact that he is so much more a benefactor than are human benefactors, his rights in this regard are so much greater than are theirs. And, of course, there are limits—of time (roughly eighty years) and of intensity (suffering beyond a certain point leads to death). (Again, I hope the reader will not think me callous in making these points. They are not points that I would put to a victim in his moment of suffering, where the need is for comfort and not theodicy; but the points are correct, and need to be taken into account by anyone who wishes to

investigate this subject with proper logical rigour.) The crucial issue, however, is whether the actual limits are too wide.

The Quantity of Evil

It may be urged that, despite the good ends that its actual or possible occurrence serves, there is too much evil in the world. My fourth condition for a perfectly good God to allow or bring about some evil is that it is probable that the good will outweigh any evil necessary for attaining it. And, even if it does outweigh it, there are—we have noted—limits to God's right to impose evil. So—is there in the world too much evil for a perfectly good God to have imposed it? An objector may agree that one does need a substantial amount of various kinds of evil in order to provide the opportunity for greater goods, and in particular a choice of destiny for human beings. But he may feel that there is just too much evil in the world, and that less evil would produce adequate benefit. It might be said that a God could give to man choice enough by allowing him to inflict quite a bit of pain on his fellows, and could deter humans from harmful actions by some nasty headaches. In our world, the objection goes, things are too serious. There is too much evil that humans can do to their fellows, and too many and too unpleasant natural evils to subserve the good of the opportunities for sympathetic and courageous response and for rational inference and inquiry that they give to humans. The suffering of children and animals is something that rightly often appals us. This is, I believe, the crux of the problem of evil. It is not the fact of evil or the kinds of evil that are the real threat to theism: it is the quantity of evil—both the number of people (and animals) who suffer and the amount that they suffer. If there is a God, the objector says in effect he has given humans too much choice. He has inflicted too much suffering on too many people (and animals) for the purpose of making it possible for them to have a free choice and to make greatly significant differences to themselves, each other, and the world on the basis of knowledge obtained by rational inquiry. No God ought to have allowed Hiroshima, the Holocaust, the Lisbon Earthquake, or the Black Death, claims the objector. With the objection that, if there is a God, he has overdone it, I feel *considerable initial* sympathy. The objection seems to count against the claim that there is a God.

But then I reflect that each bad state or possible bad state eliminated eliminates *one* actual good. Each small addition to the number of actual or possible bad states makes a small addition to the number of actual or possible good states. Suppose that one less person had been burnt by the Hiroshima atomic bomb. Then there would have been less opportunity for courage and sympathy; one less piece of information about the effects of atomic radiation, less people (relatives of the person burnt) who would have had a strong desire to campaign for nuclear disarmament and against imperialist expansion. And so on. Of course removal of one bad state or the possibility of one bad state will not remove much good, any more than the removal of one grain of sand will make much difference to the fact that you still have a heap of sand. But the removal of one grain of sand will make a bit of difference, and so will the removal of one bad state.

What, in effect, the objector is asking is that God shall very greatly diminish the number of sufferers and intensity of the suffering produced by natural processes, and the harm that humans can do to each other. What this means is that, yes, there should be diseases, but not ones that maim or kill; accidents that incapacitate people for a year or two but not for life; we could cause each other pain or not help each other to acquire knowledge, but not damage our own or each other's characters. And our influence would be limited to those with whom we come into contact; there would be no possibility of influencing for good or ill distant generations. And most of our beliefs about how to cause effects, good or evil, would be beliefs with which we would be born. Such a world would be a toy-world; a world where things matter, but not very much; where we can choose and our choices can make a small difference, but the real choices remain God's. The objector is asking that God should not be willing to be generous and trust us with his world, and give us occasional opportunities to show ourselves at our heroic best.

I have already suggested that God would not have the right to give anyone an earthly life that is on balance bad unless he provided for them a compensatory period of good life after death. To add to theism the hypothesis that he does so is to complicate theism. I am also inclined to suggest that, if God makes humans (and animals) suffer to the extent to which he does, albeit for good purposes, he would in virtue of his perfect goodness share our suffering himself. (He would recognize it as a best act to do so.) We think that good

parents who make their children eat a plain diet because of some disease that they have will often share that diet (although they do not themselves suffer from the disease); or, if they make their children play with difficult neighbouring children who are badly in need of friendship, they will show special friendship to the neighbouring parents (even if the parents are less in need of friendship). Good kings and queens share the suffering that they demand of their subjects for good purposes (for example, to win a war against an oppressor), even if the suffering of the king or queen itself would not help to forward that particular good purpose. If he makes us suffer as much as we do, God must become incarnate and share our suffering. But to add to theism the hypothesis that he does so is further to complicate theism. For, while his allowing the *kinds* of evil that he does is as such compatible with his perfect goodness, and not unexpected in view of the good states that it makes possible, my concern (as that of most people who are concerned with the problem of evil) is with the degree of that evil (the amount particular individuals have to suffer). That, my claim is, God would be justified in allowing only if he provides a compensatory period of good life after death (where necessary) and perhaps also shares the suffering of humans and animals by becoming incarnate. While I am not myself confident that there are any humans such that it would be better for them not to have lived, let me nevertheless allow the objector his claim that there are such. In that case theism needs one or may be two additional complicating hypotheses. Given them, and so the additional good that the additional evil makes possible, the degree of evil is not unexpected. For God might well be expected to ask a lot from us in order to give a lot to us.

So, given both of these additional hypotheses, and conscious of the *very* short temporal span of human and animal life (and to a lesser extent of the limits to the intensity of pain and suffering within that life that can be experienced), my own final verdict is that a God would not be less than perfectly good if he were to bring about or allow to occur that amount of suffering that exists for the sake of the greater good that results. Still, the need for additional hypotheses in order to save theism makes the resulting theistic theory more complicated than theism on its own (bare theism), and so reduces the probability of bare theism. Put another way, bare theism makes it less probable that we would find evil of as great a degree as we do than it would be on background evidence alone, because theism is

compatible with this evidence only if we add to theism a further hypothesis or hypotheses. Hence evil provides a good C-inductive argument against the existence of God. But it does not provide a very strong one, for the reason that providing life after death for many humans (not merely those who need compensation) and becoming incarnate to share their suffering are the kinds of act that a good God might well do anyway—for they are good acts (and perhaps good acts of different kinds from the other acts of God that we have been discussing, and maybe even acts of best kinds), whether or not required in order for God justifiably to allow the amount of evil that occurs. (See p. 231 for the goodness of an act of the former kind, and pp. 288–90 for additional reasons that God might have for becoming incarnate.) So, with e as the occurrence of the moral and natural evils known to us, h as the hypothesis of theism, and k as the evidence considered in previous chapters, $P(h|e \& k) < P(h|k)$, but the former is not less than the latter by very much.

Note further that, while evil may provide a good C-inductive argument against the existence of God (bare theism); it does not provide a good C-inductive argument against Christian theism (theism plus the central Christian doctrines incorporated in creeds), for life after death⁸ and God becoming incarnate are already part of the more detailed hypothesis of Christian theism; which, because of its more detailed character (its greater scope), is always as such less probable than bare theism. So any further evidence in favour of these two detailed Christian claims⁹ will diminish further the force of the C-inductive argument. (And if the only extra hypothesis required were life after death, then, since that is part of many more specific forms of theism (for example, Islam), evil would not provide a good C-inductive argument against these forms of theism.)

⁸ The more central function of a good life after death in the Christian system is to provide a reward for the virtuous, rather than compensation for the suffering. But there is quite a bit in later Church tradition that teaches that those who die as babies have a basically good life after death (in the medieval Catholic tradition, Heaven was for baptized babies and Limbo for the unbaptized). If there are any lives that are on balance bad lives, the lives of suffering babies are perhaps the most obvious example. And Jesus's parable of Dives and Lazarus (Luke 16: 19–31) sees Lazarus's good life after death as straightforward compensation for his adult life of suffering on earth.

⁹ Apart from a brief mention in Chapter 12 of the possibility of evidence that God became incarnate in Jesus Christ, the present book is not concerned with evidence for specifically Christian claims. For this, see my *The Resurrection of God Incarnate* (Clarendon Press, 2003).

The arguments of opponents who claim that the occurrence of evil provides a stronger argument against the existence of God than I am willing to allow stem, I claim, from a failure to appreciate the deepest needs of human beings and other conscious beings, and a failure to appreciate the strength of the logical constraints on the kinds of world that a God can make. The reader will sympathize with my verdict in so far as he believes that it is more important what an agent does (the choices he makes, the changes he produces in the world and the effects of his life on others) than what happens to him (the sensations and disappointments he experiences).

The Argument from Hiddenness

The physical and moral evils that I have considered so far would be evils, whether or not there is a God. But there are certain states of affairs that would be bad only if there is a God. Of these the state that has seemed most evidently to constitute evidence against the existence of God is honest agnosticism (which I understand in such a way that it includes atheism). If there is a God who is our loving creator, surely he ought to make himself known to all creatures with the capacity to understand what it is for there to be a God, a capacity that humans evidently have. Fathers who absent themselves too much from their children are rightly judged less than adequately loving. God's failure to make himself known is surely, an objector will say,¹⁰ in view of God's supposed perfect goodness, evidence against his existence. This argument is a variant on the normal argument from evil against the existence of God. Some agnosticism may be due to people not taking the trouble to investigate whether or not there is a God, or hiding from themselves the force of arguments for the existence of God; or be due to those who do believe that there is a God failing to announce the 'good news' to others. Agnosticism arising from such causes would be a moral evil for which there

¹⁰ This objection has recently been presented very thoroughly in a book devoted solely to the objection: John Schellenberg, *Divine Hiddenness and Human Reason* (Cornell University Press, 1993). Schellenberg's positive claim is that a perfectly good God would provide 'probabilifying evidence' (p. 35) of his existence (whether by experience of God or the availability of public objective argument) for all human beings capable of being aware of God, 'at all times' (p. 25); and that, since we do not all always have such evidence, there is no God.

would be justification of the same kind as for other moral evils. All-important free choices (of whether to investigate, or evangelize) would be unavailable to humans if God had already provided all humans with knowledge of his existence.

But clearly there is much honest agnosticism. Many people who have devoted much time to considering whether or not there is a God are unconvinced by any arguments known to them, and have themselves had no relevant religious experiences. Their failure to be convinced by arguments may be due to their inadequate acquaintance with cogent arguments. But, whatever the cause, they are honestly unconvinced. Why should a good God hide himself in this way from many human beings capable of reacting to him in worship and service?

My answer is twofold. Agnosticism makes possible a good for the agnostic, and it makes possible a good for the religious believer. To start with the former—a deep conviction of the existence of God inhibits someone's ability to choose freely between good and evil. It makes it too easy to choose the good for anyone who has either a strong desire to be liked by good persons (and especially any on whom he depends for his existence), stronger than any contrary bad desire; or a strong desire for his own future well-being combined with a strong belief that it is quite likely that a God would not provide a good afterlife for bad people. Why it makes it too easy to choose the good is because, as we saw earlier, in order for someone to have a free choice between good and evil, he needs temptation—a (balance of) desire to do what is evil, which he can then resist, if he so chooses. Our good desires have to be outweighed in their causal influence on us by our evil desires if we are to make a free choice in favour of the good.

It is good for us to have a desire to be liked, to like to be thought well of by others, and to have them seek our company and respect our opinions and achievements. To like and to like to be liked are essential elements of friendship (even though sometimes we may need to risk temporary unpopularity in order to secure some very important goal). And friendship with the good, and above all with the perfectly good perfectly wise all-powerful source of our existence, would be an enormous good. There would be something deeply wrong with someone who did not mind whether such a God liked him much or not. And yet how could such a God like him very much if he did and was inclined to do wrong actions? For God loves the

good and hates the evil; and, while he may still love us because he has made us and we still have some potential for doing good, he will not love us for our wrong actions.

Now, if I acquire a deep awareness of the presence of God, I will then become deeply aware that, if I do bad, and especially wrong, actions, the all-good creator will strongly disapprove. Hence, if I have the proper desire to be liked, I will have a strong inclination not to do wrong; and, unless that is overborne by some even stronger desire to do wrong, there will be a balance of desire against choosing wrong and so no overall temptation to do wrong. I will inevitably do the good. The desire to be liked may be of various strengths, as may the desire to do what is bad, and the belief that there is a God. But, if the good desire is stronger than the bad one and I have a deep awareness of the presence of God (that is, such that God's existence is not open to question), then the balance of inclination will be to the good and there will be no free choice between good and bad. We will be in the situation of the child in the nursery who knows that mother is looking in at the door, and for whom, in view of the child's desire for mother's approval, the temptation to wrongdoing is simply overborne. We need 'epistemic distance' from God in order to have a free choice between good and evil.

The only way in which a strong awareness of the presence of God will leave open the possibility of free choice between good and evil will be if the desire for divine approval is weaker than the desire to do wrong. If God makes us naturally malicious enough, keen to hurt and deceive others with no natural affection for them, then the choice may remain open. But, of course, just as it is good that we desire the love of God, so it is bad if we are naturally malicious and lack natural affection. Yet it is not logically possible that God give us both a strong awareness of his presence and a free choice between good and evil at the same time as giving us a strong desire for his love and some natural affection for our fellows. These latter are great goods; God can give us certain goods only if he does not give us others.

God could give us much more moderate awareness of his presence, ambiguous experiences but ones best interpreted as experiences of God, or arguments to show that on a slight balance of probability there is a God. Yet, even given only a moderate belief that there is a God, the possibility of a free choice between right and wrong will exist only given a certain ratio of strength between the desire to please

God and the desire to do wrong. Even if the influence of the former desire is modified by doubt about the existence of God, it cannot be too much stronger than the latter desire if there is to be that balance of desire for wrong action over right that alone makes possible serious free choice between right and wrong. Even God cannot give us that choice if he gives us fairly strong natural desires for good (including the desire to be liked by the good) and shows us, even only on a balance of probability, that he exists. But the more uncertainty there is about the existence of God, the more it is possible for us to be naturally good people who still have a free choice between right and wrong.

This point is strengthened when we take into account the other relevant desire—the desire for our own future well-being, another desire that it is very good that we should have. It is good that we should seek to make our lives good lives and want them to continue as good lives (even if sometimes we may need to sacrifice our lives for some very important goal). But, if we believe that there is a God, we will believe that whether we will have a life after death and of what kind that life will be will depend on his decision. If our present lives are bad ones in which we get pleasure out of wrongdoing, we may reasonably suspect that God may see no point in allowing us to continue to exist; or, we might suspect, God would punish us for the abuse of our present lives. Yet for God to force us to live good lives in the next world (which we would enjoy only if we wanted to live such lives) would be forcing on us desires and a destiny contrary to our present free choices. That, we may reasonably suspect, God is unlikely to do. So, if we believe there is a God, we are likely to believe that, if we want to have a good afterlife, it is to our selfish advantage to do good now. And also for the reasons given on pp. 228–31, to the extent to which we are confident that there is life after death (and we are likely to have such confidence only if we believe that there is a God), we are deprived of the possibility of making certain kinds of very serious choice.

I conclude that, if we believe that there is a God and desire the approval of any good God there may be for our actions and desire to have a good life after death, we shall be less open to temptation to do what is bad and have less opportunity for serious good actions. That will make our commitment to the good a less serious one. Only a significant balance of desire in favour of evil (with many opportunities for serious good actions) gives us the possibility of showing total

commitment to the good. When religious believers do what is wrong, either they do not have the right desires of the kind discussed above in sufficient strength, or they are subject to other strong desires to do evil, or their wrongdoing is worse than that of the agnostic—for they have yielded to a very weak temptation.

The agnosticism of the agnostic also makes possible a great good for the religious believer. It allows the believer to have the awesome choice of helping or not helping the agnostic to understand who is the source of his existence and of his ultimate well-being (helping the agnostic not merely by verbal preaching but by an example of what living a religious life is like). The existence of honest agnosticism may, if there is a God, be due to the failures of believers to help agnostics in these ways.

But while, if there is a God, there are these good states that the evil of agnosticism makes possible, the goodness of these states (as of some of the other good states discussed in this chapter) depends on their being temporary. Agnosticism allows the agnostic to make a more serious commitment to the good than he would be able to make if the presence of God were more obvious. As his earthly life progresses, so he begins to form his character for good or ill. Once he has become committed to the good, the advantage of agnosticism in helping him to do it with great seriousness disappears. If he makes himself a good person, he makes himself a person ready to worship his creator if he learns that he exists, whether in this life or another one. And the goodness for the religious believer of the existence of agnosticism is for him to have the opportunity to abolish it. It loses its point if the believer makes himself so hard-hearted as to be indifferent to it. So, of course, if God has made us, it is a great good that he should show us his presence, and I shall be arguing in the next two chapters that he does show his presence to many humans. And there would be no good in the existence of agnosticism for the religious believer if there were no religious believers. So some must be aware of the presence of God (either through religious experience or seeing the force of arguments) if the existence of agnosticism is to provide an opportunity for them.

So I suggest the conditions for a perfectly good God to permit or bring about the evil of agnosticism are satisfied. Doing so makes it possible for the agnostic to have the great good of a serious free choice between good and evil and also at the same time to have the great goods of the desires to be liked and to have a good future,

without having strong malicious desires. It also makes it possible for others to have the great good of the choice of whether or not to help the agnostic to reach true religious belief. As far as we can tell (that is, given that we have free will), we do have these great goods. God has the right to allow us to suffer the evil of agnosticism for the period of our earthly life for the same reason as he has the right to allow us to suffer other evils—so long as there remains a balance of good in our lives, which normally there will be. Given such a balance, I cannot see that God has any obligation to reveal his presence to the agnostic after his death, good though it is that he should do so. Agnosticism, though (if there is a God) an evil, is not an evil as awful as some of the suffering that does require God to provide a compensatory afterlife or to become incarnate. And the great good of significant free choice (while having the two good desires in great strength) is a great good indeed, sufficient to outweigh—I suggest—the evil of agnosticism. I conclude, therefore, that the argument from hiddenness does not constitute a good *C*-inductive argument against the existence of God. With *e* as the existence of honest agnosticism (and atheism), *h* as theism, and *k* as the evidence that we have considered, previously $P(h|e \& k) = P(h|k)$.

Arguments from History and Miracles

Arguments from History

The arguments for and against the existence of God that I have discussed so far have been arguments from very general and evident features of experience. The premiss of the cosmological argument, for example, that there is a universe is both a very general feature of experience and one that is evidently true. The same applies to the argument against the existence of God from the existence of evil. The existence of evil in the world is evident to everyone, and it is a general feature of the world manifest at many places and many times. I am now reaching arguments for the existence of God from more particular goings-on. If there is a God, one might well expect him to make his presence known to human beings not merely through the overall pattern of the universe in which he placed them, but by dealing more intimately and personally with them. In this chapter and the next I shall consider claims that he has done so. In this chapter I consider the claim that God has made himself known in ancient and contemporary public events within human history, the occurrence of some of these events being very evident and the occurrence of others being much disputed.

Many people have claimed to see the hand of God in the sudden appearance of wise men, prophets, and great leaders on the historical scene, bringing great moral, political, or religious truths to their generation and inspiring them to act in accord with them. All theistic religions have in them central figures of this kind like Muhammad or Jesus Christ or Abraham, but they have less central figures as well—people who have carried good news to new countries, reformed

religious institutions, revived religious life, or applied religion to social and political affairs. From the work of such leaders, whole religions, and traditions within religions, have flowed. The life, death, and purported Resurrection of Jesus Christ gave rise to the Christian Church, a group of Galilean fishermen who suddenly announced the story of Jesus to the world. The Church grew rapidly, yet for proclaiming its news its members were persecuted ruthlessly by the civil authorities, and very many were executed. They never offered physical resistance to the authorities, and hundreds of thousands joined them—so many in fact that the persecutions ceased and the illicit religion became the official religion of the Roman Empire. For many Christians these remarkable events have been evidence not merely that a God in whom they already believed was at work in a particular way, proclaiming a particular message, but that there is a God.

Then we have particular events not central to a religious tradition, but ones that in some way occur within a religious tradition, and that have been seen by those few who were close to them as especially compelling evidence of the hand of God in history and so of God's existence. For many people the 'answering' of their prayers (sometimes by the literal fulfilment of those prayers, sometimes by an unexpected outcome that the person praying feels to be equally satisfactory in removing the concern that led him to prayer) has seemed just such compelling evidence. Very often these are prayers for someone's recovery from a serious illness, physical (such as cancer), or mental (such as deep depression).

As well as events whose occurrence is not disputed, there are other events, a few of them central to a religion, but most of them, while occurring within a religious tradition in no way central to it, whose occurrence is normally disputed. The reason why the occurrence of these events is disputed is that their occurrence would seem (rightly or wrongly) to be (in a sense that I shall make more precise shortly) a 'violation' of the laws of nature and so an intervention in those laws by a power outside nature, such as God. This class of (purported) events includes the bodily Resurrection of Jesus, levitations (people who are praying rising in the air contrary to the law of gravity), and medical miracles, such as a spontaneous recovery from cancer. There is no space in this book to discuss detailed evidence for or against particular historical claims. All that I can do is to list the kind of evidence that is relevant to establishing disputed claims about whether certain historical events occurred; and to discuss generally

what evidential force the occurrence of such disputed events would have (if they occurred), and the occurrence of other striking but non-disputed events does have, in pointing to the existence of God.

A historical event *E* may be an event of either of two very different kinds. It may be an event brought about by deterministic laws or allowed to occur by probabilistic laws, or it may be (in the sense yet to be clarified) really a 'violation' of natural laws. The undisputed public events are plausible candidates for events of the former kind, while the apparent 'violations' such as those mentioned are plausible candidates for being really violations. But of course it may be that some of the undisputed events are really violations, and some of the apparent violations are not really violations at all.

All the undisputed public events of the kinds that I have described are ones that God has reason to bring about. I argued in earlier chapters that God has reason for giving human beings freedom and leaving them to work out their own destiny. But suppose they start to abuse that freedom; suppose they make wrong choices and fail to grow in moral and religious knowledge (in understanding the nature of God and their duties towards their fellows). What will God do? He has reason of concern and compassion to raise up prophets and leaders to announce moral and religious truths and to encourage societies to pursue right paths. God has reason for not giving the leaders manifest 'supernatural' powers to overawe or dragoon societies into doing the right thing—for, if he did, people would see that it was evidently in their self-interest to behave in a morally right way. That would severely inhibit their freedom to choose between good and evil in a way similar to the way that I analysed at the end of the last chapter. Still the occasional appearance of wise men and great leaders is perhaps to be expected if there is a God—for God wishes to encourage, though not to compel, pursuit of the good. Parental analogy suggests that God might act as I have described. When teenage children start going wrong, the good parent will not overrule their freedom (they have the right to make a mess of things); but he will not leave them without some instruction and encouragement. Similarly God has reason for bringing about a sudden recovery of the sick—for suffering is always something that God has reason to remove, and quickly (even though he also has reasons of the kinds discussed in the previous chapter for allowing it to occur in the first place). And likewise God has reason for bringing about 'converting' events in a person's life.

But although God has reason for bringing these things about, he also has reason for not bringing them about or not bringing them about too automatically in response to human needs. We have seen at some length in the last three chapters that God has reason not automatically to remove suffering. But he also has reason for not inevitably bringing about guidance for communities who go astray or converting events in lives, which are too obvious and compelling. The major such reason is that it is good that humans should decide for themselves whether to warn or convert others, that humans should have some responsibility for the (immediate or long-term) spiritual destiny of their fellows. Still, the parental analogy suggests that a good God will come to the help of the human race if communities go too far astray and people forget about God. But the argument so far provides no reason for supposing that God needs to intervene in the natural order to achieve these goals. He needs simply to make the laws of nature such that sometimes people recover spontaneously from illness, and such that there are occasional people sensitive to the deepest needs of the culture of their time who believe that they have a call to influence others to deal with these needs, and that others should be able to respond to their message.

In so many cases where people have wished to attribute some particular event to divine agency, we simply do not know what are the natural laws operative in the relevant field and so whether the event occurred in conformity with those laws or not. This is particularly the case with psychological healings, individual and society-wide conversions, acts of unusual heroism, and cases of deep prophetic insight. The way the early Christians behaved is contrary to the way most of us behave most of the time. But perhaps there are no laws in this field, and the unusual sometimes occurs. Or perhaps there are laws, and, given the circumstances of Palestine in the first century AD, Christians behaved as the laws of sociology (as yet unknown to us) predict. Or again, a man prays for the recovery of his friend from depression and the friend spontaneously and permanently recovers. Do we have a violation here—who knows? Maybe human mental states are not subject to any more than the vaguest of probabilistic laws; or, if they are subject to strict laws, we do not know what these laws are. But certainly we have good reason to suppose that most quick recoveries from physical illness without medical aid occur as a result of the normal operation of the laws of biochemistry (even though we may not know what these laws are and how they operate

in the particular case). It may well be that God often intervenes in the natural order and sets aside the natural laws in order to effect psychological cures, or to call individuals to do various jobs. I shall, however, be arguing shortly that God has reason to bring about some answers to prayer, and so some recoveries from physical illness (since so much prayer is for God to heal someone) by violating natural laws rather than by making the laws such as to produce spontaneous healings. But in general in cases of these kinds, our ignorance of what are the natural laws means that there is no public evidence that he has set aside their normal operation. Still, the (in these cases) undoubted fact that such world-changing, religious-reforming, psychological-healing and converting events occur; and that (if there is a God) he has reason to bring these events about (whether via natural laws, or by setting them aside) provides a bit more evidence for the existence of God. For the laws of nature need a special character to permit quick recoveries from physical and psychological illnesses, and life-altering conversions. And, if there are to be leaders who revive religion and apply it to social and political affairs, the probabilistic laws that govern human behaviour will have to make it probable that such figures will emerge from time to time when they are needed and that people will be open to their influence. If there is no God, there is no particular reason to suppose that the laws of human behaviour would have these features. That the laws governing human behaviour do have these features derives from the fundamental laws of nature and so is something 'too big' for science to explain.

The Nature of Miracles

There are, however, certain events, the occurrence of which is normally disputed, which are such that if they occurred they would constitute a 'violation' (or 'quasi-violation') of laws of nature. On my preferred account of laws of nature, the substances-powers-and-liabilities account, 'laws of nature' are simply summaries of the powers and liabilities of substances that have the same powers and liabilities as other substances of the same kind. I now define a 'violation' of a law of nature as the occurrence of an event that is impossible, given the operation of the actual laws of nature; that is, it involves something happening that could not happen if physical objects have the powers and liabilities that are summarized in laws

of nature. A 'violation' will be an exception to a universal (or deterministic) law—that is, a law of the form 'all *A*'s are *B*' or 'all *A*'s do *B*'. It will be an *A* that is not *B* or does not do *B*. If laws of nature provide complete explanations of all events within their scope, the only changes in respect of which substances exist and in the powers and liabilities of substances, are those caused or permitted by other substances in virtue of their powers and liabilities. In that case there can be no violations of laws of nature. There can only be violations either if there occurs some uncaused event of a kind that will not be repeated in similar circumstances; or if some cause from outside the system of law-governed objects intervenes to bring about an event not permitted by laws. That cause must be a unique physical object that does not behave in the same way as other objects of its kind, or a non-embodied person—God or some lesser spirit.

The evidence that some event *E* is a violation is that its occurrence is incompatible with what are probably (on the evidence we have) the true laws of nature. As I argued in Chapter 3, the evidence that a purported law is a true law comes from its explanatory power and its prior probability. To take one example that I used there, suppose that you have observed many positions of planets, and propose a law that 'all planets move in ellipses'. From the proposed law and some of the observed positions, you can predict all the other positions, positions that you would have no other reason for expecting—which gives the proposed law high explanatory power. The proposed law is a simple one, which (unless dissonant with background knowledge of laws in analogous fields) suffices to give it high prior probability. All this justifies the claim that the proposed law is a law of nature and hence justifies the expectation that it will hold in future without exception.

Suppose now that one day Mars moves out of its elliptical path for a brief period and then returns to the path. There are two possibilities. This wandering of Mars may occur because of some current condition of the universe (for example, the proximity of Jupiter drawing Mars out of its elliptical path), such that, if that condition were to be repeated, the event would happen again. In this case the phenomenon is an entirely regular phenomenon. So what might have appeared originally to be a fundamental law of nature proves now not to be one. It proves to be a consequence of a more fundamental law that the original purported law normally holds, but that, under circumstances describable in general terms (for example, 'when other

planets are close to Mars'), there are exceptions to it. Such repeatable exceptions to a purported law merely show that the purported law is not a fundamental law of nature. The other possibility is that the exception to the law was not caused by some current condition, in such a way that if the condition were to recur the event would happen again. In this case we have a non-repeatable exception to a law of nature. We would have grounds for believing that the exception is non-repeatable in so far as any attempt to amend the purported law of nature so that it predicted the wandering of Mars as well as all the other observed positions of Mars would make it so complicated internally and so dissonant with the rest of scientific knowledge that constitutes our background evidence that we would have no grounds for trusting its future predictions. It is no good, for example, amending the law so that it reads 'all planets move in ellipses except in years when there is a competition for the World Chess Championship between two players both of whose surnames begin with K'. Why not? Because this proposed law mentions properties that have no other place in physics (no other physical law invokes this sort of property) and it mentions them in the form of an exceptive clause, 'so-and-so holds except under such-and-such circumstances' (and thus involves two unconnected variables, 'so-and-so' and 'such-and-such circumstances'), where the clause does not follow naturally from any theory. What we need if we are to have a more adequate law is a simple formula, of which it is a consequence that the exception to the original law occurs when it does.

In these ways we could have grounds for believing that an exception to a purported law was non-repeatable. This would show that its occurrence was incompatible with the laws of nature that operated in that field, and so was a violation of a natural law. Any cause of this event would lie outside the system of natural laws. Claims of this sort are, of course, corrigible—we could be wrong; what seemed inexplicable by natural causes might be explicable after all. (The true natural laws, although they appeared to be universal laws, might be such as to permit some quite random event to occur at very rare intervals.) But then we could be wrong about most things, including claims of the opposite kind. When I drop a piece of chalk and it falls to the ground, every one supposes that here is an event perfectly explicable by natural laws. But we could be wrong. Maybe the laws of nature are much more complicated than we justifiably suppose, and the laws of Relativity Theory and Quantum Theory are mere

approximations to the true laws of mechanics. Maybe the true laws of mechanics predict that almost always when released from the hand chalk will fall to the ground, but not today because of a slightly abnormal distribution of distant galaxies. However, although the true laws of nature predict that the chalk will rise, in fact it falls. Here is a stark violation of natural laws, but one that no one detects because of ignorance of natural laws. 'You could be wrong' is a knife that cuts both ways. What seem to be perfectly explicable events might prove, when we come to know the laws of nature much better, to be violations. But of course this is not very likely. The reasonable investigator goes by the available evidence here, and also in the converse case. He supposes that what is, on all the evidence, a violation of natural laws really is one. There is good reason to suppose that events such as the following if they occurred would be violations of laws of nature: resurrection from the dead of a man whose heart has not been beating for twenty-four hours and who counts as dead by other currently used criteria; water turning into wine without the assistance of chemical apparatus or catalysts; a person growing a new arm from the stump of an old one; levitation (some one praying rising up in the air).

Suppose now that the laws of nature in some field are probabilistic. They will not then rule out any occurrence in the field and so the strict notion of a 'violation' of a law of nature will have no application. If a purported law says that '99.99 per cent *A*'s are *B*' (that is, 'the physical probability of an *A* being *B* is 0.9999'), then, however many *A*'s that are not-*B* occur, their occurrence will not be ruled out by the law; it is merely very very improbable (logically) that a large proportion of some finite class of *A*'s will be not-*B*. Such an occurrence might be far more probable if some other law was the true law. For example, if 1,000 out of 2,000 observed *A*'s were found to be *B*, that would be very, very improbable if it were the law that '99.99 per cent *A*'s are *B*' and far, far more probable if the true law were '50 per cent *A*'s are *B*'. Among events compatible with some law *L*, one (*E*) might occur that was so improbable, given *L*, that its occurrence counts very strongly against the claim that *L* holds without exception. Yet there may be no other simple formula that predicts more accurately what happens. Any attempt to amend or replace *L* so as to have a law that predicts more accurately might make it so complicated internally and dissonant with the rest of scientific knowledge that we would have no grounds for trusting its subsequent

predictions. In such a case, it is highly probable that either *E* is an uncaused event of a non-repeatable kind, or that *E* is caused and its cause lies outside the system of natural laws, and so that *E* is what I shall call a quasi-violation of natural laws.

The evidence that some event *E* is a quasi-violation is any evidence that a probabilistic law *L* is the true fundamental law in the field, and that *E* is very, very improbable given *L*, in comparison with other events described in equal detail that could have occurred instead of *E*. It is compatible with Quantum Theory that all the atoms in a large block of Carbon-14 should decay simultaneously in the next minute, but it would be so vastly improbable that this should happen to one of the very few blocks we have studied in the history of the universe that the occurrence of such an event would rightly be regarded as casting very grave doubt on that theory. Nevertheless there might be so much other evidence in favour of Quantum Theory that to amend it solely to deal with such an apparent counter-instance would make it so complex as to make it unlikely that its other predictions would be correct. In those circumstances we may reasonably conclude that the apparent counter-instance is a quasi-violation of a statistical law.

If the laws of nature are universal, it follows merely from a sufficiently full description of *E* whether or not (if it occurred) it would be a violation of them. Probability will enter into the issue of whether *E* is a violation only in so far as it is merely probable on our other evidence what are the laws of nature. But, if the laws of nature are probabilistic, then it can only be very very probable, not certain, given a description of *E*, that *E* is a quasi-violation of laws; and probability will, as before, also come in so far as it is merely probable what the laws of nature are. The example of the last paragraph, however, shows the kind of event that we would be right beyond reasonable doubt to count as a quasi-violation. And so, even if all the fundamental laws are the probabilistic laws of Quantum Theory, levitations and all the other events listed four paragraphs ago would still be rightly considered to be beyond question quasi-violations of natural laws—so improbable is it that the small indeterminacies allowed by Quantum Theory would permit their occurrence. Violations and quasi-violations of natural laws are—if they occur—‘too odd’ to be explained scientifically.

These events might be uncaused random events, but the very fact that there are laws of nature (universal or probabilistic) operative in the relevant field and all other fields of which we know makes this very

improbable. Likewise, in the absence of any other evidence for its existence, it is very improbable that there is some unique physical object that upsets laws of nature from time to time. That leaves us to look for a personal explanation of the occurrence of violations or quasi-violations—by the agency either of God or of some lesser spirit (ghost, poltergeist, demon, or whatever). We have seen in earlier chapters that there is quite a bit of evidence that counts far more strongly for the existence of God than for the existence of lesser spirits. If there is a God, natural laws can be set aside only by the action or with the permission of God who sustains them in operation. And, indeed, it is simpler to suppose that, if natural laws can be set aside by the action or permission of some agent, that is the agent who keeps them in normal operation. Hence, in the absence of positive evidence for the existence of lesser spirits whose action does not depend on the permission of God, the most probable explanation of any violation or quasi-violation is that it was brought about by or with the permission of God. Further reason to suppose that this is the case would be provided by showing that the relevant event was an event that God would have wished to bring about. I shall come shortly to consider what sorts of violations (or quasi-violations) God would have reason to bring about. I shall now follow a normal usage of the word ‘miracle’ and call a violation (or quasi-violation) of a natural law by the action or permission of God a miracle.¹ (Henceforward I shall omit the ‘or quasi-violation’ and ask the reader to assume that what I have to say about violations applies also to quasi-violations.)

Evidence of the Occurrence of Miracles

But how can we learn that an event *E* of a kind that, if it occurred, would be a violation of natural laws and so probably a miracle in fact

¹ Aquinas wrote that a ‘miracle’ in a wide sense is any event brought about by a rational agent in virtue of powers greater than normal human powers; and so many events brought about by demons or angels would count as miracles. But in a strict sense, he claims, a miracle is that which occurs outside the whole system of created nature; it is that which no other agent except God has the power to bring about. See *Summa Theologiae*, Ia.114. Hume defined a miracle as ‘a transgression of a law of nature by a particular volition of the deity, or by the interposition of some invisible agent’ (*An Enquiry concerning Human Understanding* (1777 edn.), ed. L. A. Selby-Bigge, 2nd edn. (Clarendon Press, 1902), 10.1.90 n.). Note the move from defining miracle in terms of an exception to the powers of substances, to defining it in terms of an exception to laws of nature now thought of as more fundamental than the former.

occurred? For the very fact that, if *E* occurred, it would have been violation of a natural law is in itself of course evidence against its occurrence, as Hume classically argued. This is because, if *L* is a law of nature, then it is vastly probable that what occurs will occur in accord with *L*; the past phenomena that make it (logically) probable that *L* is a law of nature make it probable that, on the occasion in question, things conformed to *L*. But there may be much other evidence that the event *E* in fact occurred. There are four kinds of evidence about what happened on a particular past occasion. First, each person has his own apparent memory of what happened—I seem to remember seeing John yesterday. Secondly, we have the testimony of others as to what they seem to remember—several people may claim to have seen John dead the day before yesterday. Thirdly, we have traces of the past, physical remains such as footprints, fingerprints, cigarette ash, Carbon-14, which allow us, given knowledge of laws of nature (other than *L*), to retrodict what happened in the past. Knowing that cigarette ash is caused by smoking (or otherwise burning) cigarettes, and very rarely in any other way, we can retrodict from the presence of the ash that previously a cigarette was smoked (or otherwise burnt). Evidence of these three kinds I shall call the detailed historical evidence about what happened. For most of us and for most purported violations the only relevant detailed historical evidence is that of testimony. The fourth kind of evidence about what happened is our background knowledge about how things are or behave on other occasions—and this acts as a corrective to discount some of the claims made on the basis of the first three kinds of evidence. If I report that I met a man ten feet tall, you will be suspicious of my report on the ground that men do not normally reach a height of ten feet. I argue in detail in the next chapter that, other things being equal, we should believe our apparent perceptions (that what we seem to see is really there), our apparent memories of them, and the testimony of others about what they claim to have seen. I analyse there what reasons we can have for not believing our apparent memories of what we perceived, and the testimony of others. Crucial among those reasons are that we have very strong other evidence that things are not as we or others seem to perceive them; and such reasons may include general background evidence that the sort of thing we seem to perceive does not happen (to take the sort of examples considered there—that the person we seem to perceive does not exist, or would not say the sort of thing we think we heard him say).

The most weighty kind of background evidence for discounting apparent observations or testimony would be that the kind of event reported would, if it happened, be a violation of a law of nature—claimed Hume.² Even given Hume's assumption that in such cases the main relevant background evidence just is our evidence about what are the laws of nature, I see no adequate reason to suppose that this evidence always counts decisively against the report. Maybe so many careful witnesses report very clearly what happened that their evidence can outweigh the evidence from the normal operation of laws of nature; and so the weight of evidence might show that a violation of law of nature occurred. But Hume's main mistake was his assumption that in such cases our knowledge of what are the laws of nature is our main relevant background evidence. Yet all other evidence (of the kind discussed in other chapters of his book) about whether there is or is not a God is also relevant, since, if there is a God, there exists a being with the power to set aside the laws of nature that he normally sustains; whereas, if there is no God, there is far less reason to suppose that violations might sometimes occur. Hume might be right in his view that, if evidence about possible violations of laws of nature were the only evidence for or against the existence of God, this evidence would never be sufficient to 'make it a just foundation for any . . . system of religion.'³ But, of course, such evidence is not the only evidence. Any evidence that there is no God is evidence against the occurrence of an event such as *E* (which, if it occurred, would be a violation); and any evidence that there is a God is evidence that laws of nature can be violated, and then testimony to the occurrence of *E* might be good enough to outweigh the evidence of what, given the laws of nature, would probably occur on this occasion.

This will be the case in particular if *E* is an event of a kind that a God would have reason to bring about as a violation of natural laws. What reasons could God have to bring about some event that he has reason to bring about, by violating the natural laws that he has made rather than by making the laws such as to bring about that event in the normal course of their operation? There are, I suggest, reasons of two kinds. The first is to respond to free human actions, either actions of praying to God for some good thing to happen, or other

² See the whole argument of section 10 of his *Enquiry concerning Human Understanding*.

³ *Ibid.* 10.2.98.

good or bad actions that they do. An entirely regular world in which everything (apart from human choices) occurred in accordance with natural laws would not be a world in which God had any living interaction with human beings. It would be a world in which God had planned in advance what would happen if some human did this or if some human did that, and he would have so arranged natural laws that his responses to human actions and situations was built into the laws. He would never respond to the sins of humans as they committed them, their requests and acts of worship as they made them. And that would give our dealings with him a very impersonal quality. He could, of course, still have made the world so that different things happened to people according to the different free choices that they made. Thus he could have arranged it so that, if Augustine found it very difficult to break free from his sinful life, he would hear someone saying ‘take and read’, which led to his opening his Bible at a crucial passage, which urged the reader to turn away from a sinful life.⁴ But it would have been built into the world in advance that this would happen; God would not have been responding to that person’s choice of a selfish path as he made it. God has the reason of friendship to seek living interaction with people whom he has made who are conscious beings and free agents like himself. Hence one would expect him to intervene in the natural order occasionally in response to the human situation, especially in answer to request (that is, petitionary prayer) for good things. Just as a good parent wants to give the good things to children that they choose to request (so that it is up to the children whether or not they get these good things), so God wants it to depend on freely made petitionary prayers whether we and those for whom we pray get many good things. But in these cases one would not expect it to be too obvious that natural laws had been violated, at any rate very often. No parent advertises that she can be relied on to get her children out of difficulties—at least no parent who wants her children to grow up advertises that. So, although—if there is a God—God might well intervene in the world to answer prayers quite often, one would not expect there to be too much evidence of God’s existence from miracles done in answer to prayers for particular needs.

⁴ See Augustine’s description of the incident in his *Confessions*, 8.12. See the full analysis by Nicholas Wolterstorff in his *Divine Discourse* ((Cambridge University Press, 1995), chs. 1–2 and subsequently) of how the words that Augustine heard someone saying can be regarded as God’s command to Augustine.

The second kind of reason why God might intervene in the natural order is just occasionally to put his signature on the work or teaching of some prophet in order to show that that work or teaching was God's work or teaching. In these cases God has reason to provide some (not necessarily enormously strong) evidence that he has violated natural laws. He could do this by causing the prophet's life's work to achieve its goals or his teaching to be propagated on a large scale, as the result of an apparent violation of natural laws. Humans need to know which reforming movements are good and which are not, and sometimes this is not obvious. It helps us to see which movements are good if (the evidence suggests) God makes them flourish by intervening in nature rather than by allowing natural laws to have their normal effects—for only God who conserves natural laws in operation can set them aside. And humans need to know truths that are beyond their capacity to discover with reasonable certainty for themselves. Although our natural reason shows us some clear cases of good and especially obligatory actions, and some clear cases of bad and especially wrong actions, our moral intuitions are unclear in respect of so many kinds of action. Are abortion, euthanasia, homosexual relationships, war, capital punishment, corporal punishment, etc. always wrong, or only sometimes? And humans may need to know also metaphysical and historical truths beyond their capacity to discover with reasonable certainty for themselves. They may need to know more than natural reason can show about what God is like and what he has done for humans, in order that they may worship him properly. In all these ways humans need information, and so God has reason to reveal the truth on these matters in a revelation authenticated by miracle. Judaism, Islam, and Christianity have all claimed divine approval for their movements by means of miraculous foundation events. The plagues of Egypt that forced the Pharaoh to free the people of Israel, the burning bush through which God revealed his identity, the manna by which the Israelites were fed daily in the wilderness, and the fall of the walls of Jericho have in the past—rightly or wrongly—been claimed as events by which God showed that Israel were a chosen people. Islam has claimed that the production of the Koran by an uneducated Mohammad showed God's miraculous witness to its truth. And, more than all other major religions, Christianity has claimed miraculous authentication, as I shall spell out more fully shortly.

So evidence that some event was an event that God would have reason to bring about miraculously could be in the form of evidence that it would have been (if it occurred) an event of one of the above kinds. It could be evidence that it was an event of a kind that God might be expected to bring about as a response to particular human actions—for example, an answer to prayer for some good cause. Or it could be evidence that, if God brought it about, he would thereby have vindicated some religious movement or the teaching of some prophet in favour of some plausible moral or other doctrinal claim. I emphasize that the religious movement or moral or doctrinal teaching must already have some modest degree of probability; the detailed historical evidence for the miracle will then increase that probability. But, if some movement is obviously wicked, or some teaching obviously mistaken, that must count against the occurrence of any event that, if it occurred, would be a violation of natural laws and so probably to be attributed to God giving his authentication to that movement or teaching. It is only in respect of movements or teachings that might well be good or true that we can attribute any apparently miraculous support to divine agency. A purported miracle in support of a movement of racial hatred, or confirming the goodness of child sacrifice, could not have been a real miracle—either it did not occur or it was not miraculous.

My conclusion to this section is that, in so far as we have historical evidence (normally in the form of testimony) to the occurrence of an event E that is such that, if it occurred, it would probably be a violation of natural laws and that is of a kind that there is some probability that a God would have reason to bring about, that makes it more probable than it would otherwise be that there is a God. For a God would be expected occasionally to bring about such events. Yet, if there is no God, there is no significant probability that such events will occur. Evidence that an event is of this kind is evidence that it is an event 'too odd' for science to explain. Hence, with k as the evidence discussed in previous chapters, and e as evidence of testimony of the above kind, and h the hypothesis of theism, $P(h|e \& k) > P(h|k)$. It will depend on the strength of the testimony, by how much $P(h|e \& k)$ exceeds $P(h|k)$. In so far as other evidence suggests that there is no God, testimony to the occurrence of such violations is probably misleading—but how probably will depend on how improbable it is to start with that there is a God and how strong

the testimony is (that is, how many witnesses of established honesty and accuracy in reporting testify to the event).

Hume commented that ‘every miracle . . . pretended to have been wrought in any of these religions [“of ancient Rome, of Turkey, of Siam, and of China”], as its direct scope is to establish the particular system to which it is attributed; so has it the same force . . . to overthrow every other system’.⁵ That is, put in my terminology and using a particular example, evidence of a violation that, if it occurred, would show God’s approval of an Islamic doctrine would thereby be evidence against the occurrence of any violation that, if it occurred, would show God’s approval of some Christian doctrine incompatible with the Islamic doctrine; and conversely. If the production of the Koran really was a miracle, the Resurrection could not have happened, and conversely. But, while this point is formally correct, note that very few purported miracles have this character. Most are simply answers to prayers for the needs of particular individuals; and it is compatible with both Christian and Islamic doctrines and the doctrines of most other religions that God may answer the prayers of members of all religions. And many doctrines of one religion are compatible with doctrines of another religion. Christianity incorporates most of Judaism, and is certainly happy to recognize the occurrence of its foundation miracles. But there are cases of conflict, and for those cases Hume’s point is correct. It follows that that religion (if any) which has the best authenticated miracles has the best evidence from this source in its support.

Incarnation

It would be appropriate before concluding this chapter to illustrate in more detail the role of the founding miracle of Christianity, that religion among the major religions in which miracles have had the most important role. Christianity has claimed that there is one special event that God had reason to bring about, a particular intervention of himself into the world that he made, an incarnation. Suppose that the human race gets into a really bad mess. Suppose that people so abuse their freedom that they teach others evil and not good. They do not altogether know which actions are right and

⁵ *Enquiry Concerning Human Understanding*, 10.2.95.

which are wrong, and they conceal from themselves even what they do know. They show little interest in where they came from (for example, whether they have a creator to whom thanks and service are appropriate), or in whether their existence has any point and their race any destiny. They do not care for their fellows, but live for self. Now the Christian view is roughly that such was the human condition at the outset of civilization; and that, but for various, especially Christian, influences from without, it still is. Now whether this Christian view is correct here is of course a matter for argument, which will turn both on issues of history and psychology and on moral issues; and once again there is no space to pursue these issues. However, few at the beginning of the twenty-first century would deny that this view has a certain plausibility. Suppose that this Christian view of the human condition is correct. What does God have reason for doing about it?

There may be more than one thing that God has reason for doing about such a human condition; but one kind of response that God has reason to make is the following. He might conclude that things had gone so wrong that an atonement was needed; that the human race ought by sacrificial action to show its contrition to its creator. Yet he might also conclude that it was not within the capacity of a fallen race to make this kind of atonement; and that, if atonement was to be made, it would have to be made on behalf of the race by a human being preserved from the worst influences to which humanity was normally subject. But it would not be right of God to single out any ordinary human being to make such a sacrifice. God could insist on the sacrifice of none other but himself. So God has a reason to bring about an incarnation of some kind by himself becoming human in order to make an atonement. Once again, there are big Christian assumptions here which there is no space to discuss—for example, whether atonement of this kind is morally good, or whether it is better for people just to forget wrongdoing; and there is also a big philosophical assumption—that it is coherent to suppose that a God can become incarnate (perhaps there is some self-contradiction in supposing a God to become a human being), and there is not space to discuss that issue either. But, in order to continue the argument further, let us suppose that the Christian moral view of the propriety of atonement is correct, and that the concept of an incarnation is coherent. As well as concluding that an incarnation to make atonement would be a good thing, God might also conclude that the

human race needed a new start with a supreme leader and inspirer to found a society in which his work would be continued. The leader would need to teach the race moral truths that it had only dimly perceived; perhaps, in virtue of his status, also to give it new moral laws, and to show it by example how to behave. Yet again, to preserve human freedom, the powers of the leader and of the society must not be too evident or too 'supernatural'. And thirdly and most importantly, God might decide to impose on the human race a very considerable burden of evil for the sake of the very considerable good that it makes possible. In that case, as I argued in the last chapter, it would be not merely very good but obligatory for God to become incarnate to share the burden with us.

For these reasons, given the stated assumptions, if there is a God and if the human condition falls low, we may well expect there to appear on Earth a human being who lived a humble and sacrificial life and suffered the evil that humans do to other people (for example, by suffering an unjust death at their hands), who taught great moral and religious truths, who even suggested that he was God, and who founded a society to continue his work. He might manifest the divine compassion by healing, and the divine power by apparently violating natural laws in order to do so. He might show to people that his atonement availed and that it was possible for them in his new society to reform the world, by natural laws being violated in a supreme way by his resurrection from the dead. All of this, however, would be none too obvious in order that it might remain a genuine option for human beings to reject this claim of a divine incarnation. If we have evidence that things have happened like that, as in the Christian story of the life, death, and Resurrection of Jesus of Nazareth, and if we also have reason for believing the stated moral and other assumptions to be true, then that all confirms the claim that there is a God, for God has reason for bringing about such a state of affairs—namely, the good of humans. The analogy of the very good parent or very good spouse who makes a supreme sacrifice to save the lost child or spouse also suggests that some such thing is to be expected. Of course God may have reasons for not bringing about such a state. There may be alternative states of affairs that he has good reason to bring about instead. But there are certainly alternative states of affairs that he would appear to have overriding reason not to bring about—for example, the human race left eternally to make itself miserable through the original bad half-conscious choices of

sinner in the centuries before Christ. Hence the occurrence of events of the type described is more probable if there is a God than if there is not, and so their occurrence would be evidence for his existence.

There is clearly some historical evidence for all this, including the testimony of witnesses to the crucial event of the bodily resurrection of Jesus, which, if it occurred, would be beyond reasonable doubt a violation of natural laws. There is no space in this book to discuss whether it is very good evidence, nor is there space to discuss the moral and metaphysical assumptions that are required if that event (if it happened) is in its turn to be evidence of the central doctrines of the Christian religion.⁶ My concern here has been merely to point out what sort of evidence is relevant to this issue and to point out that such evidence would also be evidence for the existence of God, would provide a good C-inductive argument for his existence. It is, of course, very unlikely that I or the reader would think that God was very likely to do the sort of thing described unless we had had some contact with the Christian tradition or some religious tradition with similarities to the Christian one, and had thus come to believe that the condition of the human race was poor, and that there was a need for atonement and example. But that is no reason for supposing that what we think is not true. Unless I had been brought up in the tradition of Western mathematics, I would be unlikely to believe that there is no greatest prime number; for I would not even have the concept of a prime number. But once I have derived from tradition the relevant concepts, I am in a position to assess the proof that there is no greatest prime number. Likewise, in order to come to believe the Christian (or any other) religious system, we need first to be taught what the system claims; only then are we in a position to assess whether or not it is true.

Conclusion

I argued at the beginning of this chapter that there are certain events in human history of the occurrence of which there is no doubt that are such that God has some reason to bring about, and that are more likely to occur if there is a God than if there is not. As far as we can

⁶ I have discussed the historical evidence for the Incarnation and the Resurrection in my book *The Resurrection of God Incarnate* (Clarendon Press, 2003).

tell, their occurrence may be caused by the normal operation of laws of nature. But, in so far as, if there is no God, there is no reason to suppose that natural laws would have this character, that provides further evidence of the existence of God, making in my view a rather small addition to the probability of the existence of God. But I have argued that there are other events, whose occurrence is disputed, that are such that, if they occurred, they would have been violations of laws of nature of a kind that God would have reason to bring about; and so such that significant evidence for their occurrence would certainly confirm theism. I have not, however, for reasons of space, discussed whether any of these alleged events really occurred. And, in discussing one particular phenomenon—the life, death, and alleged Resurrection of Jesus of Nazareth—I have passed over crucial moral and philosophical questions, as well as questions of historicity. These issues are, alas, too big and controversial to be discussed within the covers of this book. There are inevitable crucial gaps in the argument of this chapter, which there are not in earlier chapters.

However, we need some result to carry forward to the next chapter. So I suggest that, in view of the fact that some few occurrent events (which are not necessarily violations of natural laws) do confirm theism, and there is clearly some slight evidence (in the form of much testimony) of the occurrence of violations of natural laws (of a kind that God would have reason to bring about), we say that there is a weak C-inductive argument from various particular events in history to the existence of God. If any reader after more detailed consideration of the historical evidence thinks that there is a stronger C-inductive argument or no good C-inductive argument to be had here, he will need to scale up (or down, as the case may be) my subsequent conclusions.

The Argument from Religious Experience

If there is a God, one might well expect him not merely to concern himself with the progress of the human race by providing opportunities for humans to do worthwhile things, or providing a revelation at a particular moment in history, or to concern himself with particular individuals by fulfilling their prayers; but also perhaps to show himself to and speak individually to at any rate some of the people whom he has made and who are capable of thinking about God and worshipping him. Certainly one would not expect too evident and public manifestations, for the reason that I gave in Chapter 11. If God's existence and intentions became items of evident common knowledge, then our freedom to choose between good and evil would be vastly curtailed. However, one might expect certain private and occasional manifestations by God to some people, although perhaps not to everyone, again for the reason that I gave in Chapter 11. The argument from religious experience claims that this has often occurred; many have experienced God (or some supernatural thing connected with God) and hence know and can tell us of his existence.

The Nature of Religious Experience

Let us begin by investigating the premiss. What are the 'religious experiences' whose occurrence is supposed to be evidence for the existence of God? An experience is a conscious mental event. It may be described in such a way as to entail the existence of some particular external thing apart from the subject, beyond the stream of his

consciousness, normally the thing of which it is an experience; or it may be described in such a way as to carry no such entailment. Thus 'hearing the coach outside the window' is not unnaturally described as an experience; but if I have such an experience, if I really do hear the coach outside the window, then it follows that there is a coach outside the window. Yet, if I describe my experience as 'having an auditory sensation that seemed to come from a coach outside the window', my description does not entail the existence of anything external of which the experience was purportedly an experience (or anything else external). The former kind of description I will call an external description; the latter an internal description. Now when people talk about religious experiences, they often give external descriptions of them. Such external descriptions may be fairly precise—'I talked to God last night', or 'I saw Poseidon standing by the window', or, rather more vaguely, 'I became conscious of a timeless reality beyond myself'. The trouble with taking any external description as the premiss of an argument from religious experience is that there is going to be considerable doubt about the truth of the premiss; but, once you accept the premiss, you are quite obviously most, if not all, of the way to your conclusion. If you accept that Joe talked to God last night, then, of course, there is a God—it hardly needs an argument to show it. If you accept that Joe became conscious of a timeless reality beyond himself, then, admittedly, that does not demonstrate the existence of God, but you are quite a lot of the way towards such a demonstration. So, it seems natural to say that, if they are to be useful, all arguments from religious experience must be phrased as arguments from experiences given internal descriptions. There are various ways of giving internal descriptions of one's experiences, but in the case of most experiences, including those that the subject believes to be of something outside himself, a normal way is to describe how things 'appear' or 'seem' to the subject—one may say 'the room seemed to be going round and round', or 'the carpet appeared to be blue', or 'he appeared to be moving away from me'. Or, in giving such descriptions, one may use verbs that describe how things seem to the subject, the use of which is confined to reporting the deliverances of particular modalities of sense, verbs like 'looks' or 'feels' or 'tastes'—I may say 'It looked as if the coach was moving away from me' or 'it felt smooth' or 'it tasted of pineapple'.

There is a crucial distinction due to Chisholm between the epistemic and the comparative uses of such verbs as 'seems', 'appears',

'looks', etc.¹ To use such words in their epistemic use is to describe what the subject is inclined to believe on the basis of his present sensory experience. If I say 'the ship appears to be moving', I am saying that I am inclined to believe that the ship is moving, and that it is my present sensory experience that leads me to have this inclination to belief. If I am using 'looks' in this way when I say 'the penny looks elliptical', I am saying that I am inclined to believe that it is elliptical, and that my inclination to belief arises from my present visual experience. By contrast, to use 'looks', etc., in the comparative use is to compare the way an object looks with the way other objects normally look. In this use 'the penny looks elliptical' means 'the penny looks the way elliptical things normally look'. The speaker is not saying and does not imply that he is inclined to believe that the penny is elliptical; he may know very well that it is not. Again, in the comparative use, 'from here it looks red' means 'from here it looks the way red things normally look'. When I describe an experience in terms of the way things seem (epistemically) to the subject, I shall say that I describe it epistemically. A full internal description of a subject's experiences would seem to involve both kinds of internal description. When I look at a penny on the table from an angle, it usually looks (in the comparative sense) elliptical, and (in the epistemic sense) circular—and my experience is of it looking to me in both of these ways.

So much for what an 'experience' is and the ways in which we can describe it. But what constitutes a 'religious experience'? The concept of a 'religious experience' in ordinary use has as fuzzy a border as the concept of a religion, and, in order to talk about arguments in this field, we need to make it moderately precise. For our present purposes it will be useful to define it as an experience that seems (epistemically) to the subject to be an experience of God (either of his just being there, or of his saying or bringing about something) or of some other supernatural thing.² The thing may be a person,

¹ R. M. Chisholm, *Perceiving* (Cornell University Press, 1957), ch. 4. Chisholm attempted to distinguish a third ('non-comparative') use of such verbs, but there is some doubt about whether there is such a use, and I am concerned only with the epistemic and comparative uses.

² This definition rules out a lot of what has often been called 'religious experience'. For example, much 'religious experience' does not purport to be an experience of anything external. While many experiences in the religious traditions of Christianity, Judaism, and Islam are what Ninian Smart, following Otto, calls 'numinous', many deep experiences called 'religious' in the Buddhist tradition are not. (Smart defines a numinous experience as 'an experience of a dynamic external presence'. See his article on 'History of Mysticism'

such as Mary or Poseidon; or Heaven, or a 'timeless reality beyond oneself', or something equally mysterious and difficult to describe. For most of the discussion I shall be concerned with experiences that seem to be simply of the presence of God and not with his seeming to tell the subject something specific or to do something specific. But where it is relevant I shall contrast experiences of God with experiences of other supernatural beings. I will come to discuss more specific experiences at the end of the chapter.

The crucial feature of the definition to which I draw attention is that what makes an experience religious is the way it seems to the subject. This definition captures those experiences that are of most importance for the purpose of this book. They are supposed to wear their religious origin on their face. What is it for the subject to be right, in fact to experience God, that is, to be aware of God, and so in a very general sense to perceive God (believing that he is so doing)? (I talk of such awareness of God as a perception without implying that the awareness is necessarily mediated via the normal senses. 'Perceive' is the general verb for awareness of something apart from oneself, which may be mediated by any of the ordinary senses—for example, it may be a matter of seeing or hearing or tasting—or by none of these.) It seems to me, for reasons that others have given at length, that the causal theory of perception is correct—that *S* perceives *x* (believing that he is so doing)³ if and only if an experience of its seeming (epistemically) to *S* that *x* is present is caused by *x*'s being present.⁴ So *S* has an experience of God if and only if its seeming to him that God is present is in fact caused by God being present.

in P. Edwards (ed.), *Encyclopaedia of Philosophy* (Collier-MacMillan, 1967). R. M. Gimello ('Mysticism and Meditation' in S. T. Katz (ed.), *Mysticism and Philosophical Analysis* (Sheldon Press, 1978)), claims (p. 193) that 'it is certain that Buddhists do not ontologize the contents of their mystical experiences, nor people the cosmos with mystical entities, since their very purpose in having them is to "discern" their illusoriness'. Peter Moore ('Mystical Experience, Mystical Doctrine, Mystical Technique' in Katz (ed.), *Mysticism and Philosophical Analysis*) also emphasizes the variety of 'mystical claims', that they include both 'subjective' and 'existential' claims. Only religious experiences of the kind that my definition picks out have apparent evidential value in pointing towards the existence of God, and that is why I am concerned with them alone.

³ There is a use of 'perceive' and other verbs of perception (e.g. 'see', 'hear') in which a subject may be said to perceive something that he does not believe that he is perceiving, e.g. I may be said to have perceived John without realizing that it was John whom I was perceiving. I am not concerned with perception of this kind, but only with perception of things that the subject believes that he is perceiving.

⁴ The best presentation of the theory known to me is that by P. F. Strawson in his 'Causation in Perception' in his *Freedom and Resentment* (Methuen, 1974). However,

Before going on to describe the different kinds of religious experience, I need to make two further preliminary points. First, I need to distinguish between public and private perceptions. An object x may be such as to cause all persons rightly positioned with certain sense organs and certain concepts who pay a certain degree of attention to have the experience of it seeming to them that x is present.⁵ In that case we shall say that the perception of x is a public perception. Almost all our perceptions—for example, my seeing a material object such as a desk—are in this sense public perceptions. For a desk is such as to cause all persons close to it (without any material objects between it and them) whose eyes are pointing at it, who are attentive, and who have normal vision and the concept of a desk, to have the experience of its seeming to them that there is a desk there. But there may be objects o that cause certain persons to have the experience of its seeming to them that o is present without their having that effect on all other attentive persons who occupy similar positions and have similar sense organs and concepts. This could be just because the causal chains that bring about perceptions of o are not at all deterministic—for example, the laws of optics might be such that there was no guarantee that a suitably equipped observer would always see what was there. Or it could be because o is a person who can choose whom to cause to have the experience of its seeming to them that o is there. o may be a normally invisible person with the power of letting you, but not me, see him. If S has the experience of its seeming to him that o is there, but, either because of o 's choice or for some other reason, not every other attentive person rightly positioned and equipped would necessarily have the experience, then S has, I shall say, a private perception of o . If religious experiences are of anything—that is, are perceptions—they are normally

having given the conditions stated above, he argues that they are still insufficient, although necessary, for perception. He claims (pp. 79–80) that there are further restrictions for different senses, e.g. that 'one can only see what is within one's arc of vision' or 'however loud the report of the cannon, if it is far enough away, it will be out of earshot'. But such restrictions seem only to be correct if we suppose that the meaning of such expressions as 'within one's arc of vision' is defined by them. If there is an independent criterion of (e.g.) 'arc of vision'—say a geometrical one—then the stated restriction seems in no way obviously a necessary truth. There is nothing incoherent in supposing that some people can see round corners. I suggest, therefore, that my analysis gives sufficient, as well as necessary, conditions for perception.

⁵ Of course what will constitute being 'rightly positioned' and what sense organs and concepts are needed will vary with the sense and the kind of object, and are matters for empirical inquiry; and so too is the degree of attention that is needed.

private perceptions. When one person has a religious experience, his neighbour equally attentive and equally well equipped with sense organs and concepts normally does not. The religious person's explanation of this is that God or gods give such experiences to those to whom they choose to give them, not to all and sundry.

Secondly, I call attention to a relation between experiences. One often perceives one thing in perceiving something else. In seeing a man dressed in such-and-such a way, I may see John Smith. In seeing the print of such-and-such a shape in the sand, I may see the footprint of a bear. In seeing an especially bright star in the sky, I may see Venus near to the earth. In these cases my very same visual or other sensations (described comparatively) that bring about my perceiving the first thing also bring about my perceiving the second thing. In perceiving the second thing one does not see anything extra in the sense of a new item that had escaped one's notice before; rather one perceives the first thing as the second thing. In these cases one person may perceive both things, and another person perceive only the first thing and yet both have the same visual sensations. This relation that holds between perceptions may also hold between experiences described epistemically. In seeming (epistemically) to see the man dressed in such-and-such a way, I may seem to see John Smith. In such cases the same sensations (described comparatively) that bring about the first experience also bring about the second. Two people may both have the same visual or other sensations (described comparatively) (for example, a bright spot in the middle of their visual fields) and through having those sensations one may have a certain experience described epistemically (for example, seeming to see a lighthouse in the distance) and the other may not. Or, of course, the same visual or other sensations may give rise to totally different experiences (described epistemically) in different people.

Five Kinds of Religious Experience

With these points in mind it will be useful to classify the different kinds of religious experience. In due course I shall make similar points about all of them, but it is worthwhile at this stage pointing out the diversity of experiences that fall under our definition. First, we have experiences that seem (epistemically) to the subject to be experiences of God or something else supernatural, but where she

seems to perceive the supernatural object in perceiving a perfectly ordinary non-religious object. Thus someone may look at the night sky, and suddenly 'see it as' God's handiwork, something that God is bringing about (in the way in which someone may see a vapour trail in the sky as the trail of an aeroplane). She has, it may be said, an experience of contingency. Secondly, there are the experiences that people have in perceiving very unusual public objects. (The occurrence of the unusual object may or may not constitute a violation of a natural law.) The experiences had by those who witnessed 'the Resurrection appearances of Jesus', or the 'appearance of Mary' at Fatima, or (as far as the great light that shone round about him, according to Acts 22: 9) St Paul's experience on the road to Damascus are in this category, if the accounts of these events are in any minimal way reliable. Take the appearance of the risen Jesus to the disciples as described in Luke 24: 36–49. A man looking and talking like Jesus who had been crucified three days earlier suddenly turned up among them and ate some fish (looking and talking like Jesus in the comparative sense—that is, looking and talking the way that Jesus used to look and talk). Yet, in perceiving this public event, the disciples had the religious experience of taking the man to be the risen Jesus Christ. Their religious experience was that he looked like Jesus in the epistemic sense, and so they believed him to be. A sceptic might have had the same visual sensations (described comparatively) and yet not had the religious experience.

The other three classes of religious experiences are ones that do not involve taking public phenomena religiously. In them the divine is apprehended via something private to the subject. In the third place we have cases where the subject has a religious experience in having certain sensations private to himself, sensations of a kind describable by the normal vocabulary used for describing the sensations that result from the use of our five senses. In his dream described in Matthew 1: 20–1. Joseph dreamed that he saw an angel who said to him certain things. Here there were no public phenomena, but Joseph had certain private sensations that he might have been able to describe by means of normal sensory vocabulary—for example, he had the visual sensation like the sensation that he would have had if he had been looking at a man dressed in white, and the auditory sensations that he would have had if someone had been saying such-and-such to him. (He might have been able to tell us the actual words that the angel seemed to be saying to him.) What made the dream a

religious experience was that, in having the sensations, *and* after he had woken up, it seemed to Joseph that an angel was talking to him—that is, he took the man-in-the-dream to be a real angel and not a mere angel-in-a-dream, and the words-in-the-dream to be words uttered by the angel. (What the biblical author meant by saying that it was a dream is presumably that the experience was one had while the subject was by normal public criteria asleep, and that the experience was not of a public phenomenon but that, at the time, though not afterwards, it seemed to Joseph that it was of a public phenomenon.)

Fourthly, we have the case where the subject has a religious experience in having certain sensations private to himself, yet these are not of a kind describable by normal vocabulary. The subject has some sensation analogous to sensations of normal kinds—for example, visual or auditory sensations, but only analogous—such that, if his experience was of a public phenomenon, we might say that it was the experience of a sixth sense.⁶ Presumably mystics and others who find it difficult if not impossible to describe their religious experiences, and yet feel that there is something to be described if only they had the words to do the describing, are having experiences of this kind. Fifthly and finally we have religious experiences that the subject does not have by having sensations. It seems to the subject, perhaps very strongly, that he is aware of God or of a timeless reality or some such thing, and yet not because he is having certain sensations; it just so seems to him, but not through his having sensations—just as it may seem to me strongly that my hand behind my back is facing upward rather than downward, yet not because of any sensations. Many mystics who claim to experience God via ‘nothingness’ or ‘darkness’ may be making the point that their experience of God is not mediated via any sensations. More ordinary cases, however, also fall into this category. Someone may be convinced that God is telling him to do such-and-such (for example, follow such-and-such a vocation), and yet there are no auditory or other sensations occurring.

If the subject is asked ‘What was it about your experience that made it seem to you that you were having an experience of God?’ in

⁶ For ways of distinguishing between senses and so grounds for saying that we have a new sense, see H. P. Grice, ‘Some Remarks about the Senses’, in his *Studies in the Way of Words* (Harvard University Press, 1989).

the case of experiences of the third and fourth kinds (as with experiences of the first two kinds), there is a partial answer, though in the case of experiences of the fourth kind we may lack the vocabulary to give it. In the case of experiences of the third kind, the partial answer will be 'because of such-and-such auditory or visual or other describable sensations that I had'. The answer will be partial because the mere fact that one was having such-and-such sensations does not make the experience seem to be of God; someone else could have those sensations without thereby having a religious experience. In the case of experiences of the fourth kind, the answer to the question will be 'because of the very unusual and virtually indescribable sensations that I had'. But in the case of experiences of the fifth kind, the answer to 'What was it about your experience that made it seem to you that you were having an experience of God?', will be 'It just did. There were no visual, auditory, or any other sensations that made it seem thus to me.'

So much for my classification of religious experiences. It is, I believe, both exclusive and exhaustive. For clearly an experience that seems to be of God may or may not be mediated by something sensory (that is, there may or may not be an answer to the question 'What was it about your experience that made it seem to you that you were having an experience of God?'). If it is mediated by something, the something may be public or private. If it is private, it may or may not be describable by normal sensory vocabulary. If it is public, it may be a common, well-known phenomenon; or something very odd, the occurrence of which may be disputed. However, even though the classification is exclusive and exhaustive, it may sometimes be by no means obvious, even to the subject, into which class a given experience falls. For example, suppose I am alone and seem to see and talk to a figure dressed in white, which I take to be an angel. The correct classification of the experience depends on what others would have experienced if they had been there—this I may not know or have any means of finding out. If others also would have seen a figure dressed in white, then the experience is of the second kind; if not, it is of the third kind.

There is no doubt at all that millions of human beings down the centuries have had religious experiences of one or more of the above kinds. Indeed, that statement rather underplays the situation. For many people life is one vast religious experience. Many people view almost all the events of their life not merely under their ordinary

description but as God's handiwork. For many people, that is, very many of the public phenomena of life are viewed religiously and so constitute religious experiences of the first type. What is seen by one person as simply a wet day is seen by another as God's reminding us of his bounty in constantly providing us with food by means of his watering plants. What is seen by one person as merely a severe illness is seen by another as God's punishing him for the sins of his youth. That God is at work is no inference for these people but what seems (epistemically) to be happening.

John Hick has called our attention to this phenomenon in various of his works but especially in *Faith and Knowledge*. He observes that:

The Old Testament prophets, for example, experienced their historical situation as one in which they were living under the sovereign claim of God, and in which the appropriate way for them to act was as God's agents; whereas to most of their contemporaries, who were 'experiencing as' in a different way, the situation did not have this religious significance. The prophets' interpretation of Hebrew history, as this is embodied in the Old Testament, shows that they were 'experiencing as' in a characteristic and consistent way. Where a secular historian would see at work various economic, social, and geographical factors bringing about the rise and fall of cities and empires, the prophets saw behind all this the hand of God raising up and casting down and gradually fulfilling a purpose. When, for example, the Chaldeans were at the gates of Jerusalem, the prophet Jeremiah experienced this event, not simply as a foreign political threat but also as God's judgment upon Israel. . . . It is important to appreciate that this was not an interpretation in the sense of a theory imposed retrospectively upon remembered facts. It was the way in which the prophet actually experienced and participated in these events at the time.⁷

As well as such experiences of the first kind, very many people, both those who are much of the time religious believers and those who are not, have had many religious experiences of the other kinds.⁸

The question must now be faced as to the evidential value of all this. Is the fact that all these religious experiences have oc-

⁷ J. Hick, *Faith and Knowledge* (2nd edn., MacMillan, 1967), 142–3.

⁸ For some modern 'religious experiences' (although ones of subjects almost entirely from an English Protestant background) see, for example, the volume summarizing a thousand reports of such experiences, provided in response to a public appeal—T. Beardsworth, *A Sense of Presence* (The Religious Experience Research Unit, Manchester College, Oxford, 1977). Some are religious experiences on my criterion and some are not. The author is inclined to draw a conclusion from these experiences that is disputable in the light of some of the arguments put forward in this chapter.

curred evidence for the existence of God (or some other supernatural reality)?

The Principle of Credulity

In discussing religious experience philosophers have sometimes made the claim that an experience is evidence for nothing beyond itself, and that therefore religious experience has no evidential value. That remark reflects a philosophical attitude that those philosophers would not adopt when discussing experiences of any other kind. Quite obviously having the experience of it seeming (epistemically) to you that there is a table there (that is, your seeming to see a table) is good evidence for supposing that there is a table there. Having the experience of its seeming (epistemically) to you that I am here giving a lecture (that is, your seeming to hear me give a lecture) is good evidence for supposing that I am here lecturing. So generally, contrary to the original philosophical claim, I suggest that it is a principle of rationality that (in the absence of special considerations), if it seems (epistemically) to a subject that x is present (and has some characteristic), then probably x is present (and has that characteristic); what one seems to perceive is probably so. And similarly I suggest that (in the absence of special considerations) apparent memory is to be trusted. If it seems to a subject that in the past he perceived something or did something, then (in the absence of special considerations) probably he did. How things seem to be (in contingent respects),⁹ that is how we seem to perceive them, experience them, or remember them are good grounds for a belief about how things are or were. The more forceful the experience, the stronger the memory, the more probable it is that what we seem to perceive or remember is true—other things being equal. Memory, of course, is less forceful than present experience, and sometimes so weak as only to make it a bit probable that we seem to remember is true. This principle, which I shall call the Principle of Credulity, and the conclusion drawn from it seem to me correct. It seems to me, and I hope to my readers, intuitively right in most ordinary cases, such as those to which I have

⁹ For a more thorough account of the Principle of Credulity and an explanation of why the phrase ‘in contingent respects’ is required, see my *Epistemic Justification* (Clarendon Press, 2001), 141–50.

just been referring, to take the way things seem to be as the way they are. From this it would follow that, in the absence of special considerations, all religious experiences ought to be taken by their subjects as genuine, and hence as substantial grounds for belief in the existence of their apparent object—God, or Mary, or Ultimate Reality, or Poseidon.¹⁰

Note that the principle is so phrased that how things seem positively to be is evidence of how they are, but how things seem *not* to be is not such evidence. If it seems to me that there is present a table in the room, or statue in the garden, then probably there is. But if it seems to me that there is no table in the room, then that is only reason for supposing that there is not, if there are good grounds for supposing that I have looked everywhere in the room and (having eyes in working order, being able to recognize a table when I see one, etc.) would have seen one if there was one there. An atheist's claim to have had an experience of its seeming to him that there is no God could be evidence that there was no God only if similar restrictions were satisfied. But, given that my rejection in Chapter 11 of 'the argument from hiddenness' is correct, there are no good grounds for supposing that, if there is a God, necessarily the atheist would have experienced him.

I shall now argue that attempts to restrict the principle in ways designed to rule out its application to religious experience are unsuccessful. I shall consider two such attempts to argue that, while its appearing to me that there are before me tables, chairs, houses, etc. is good grounds for supposing that there are (that is, its seeming to me that I am seeing them is good grounds for supposing that I am), its appearing to me that the world before me is being sustained by God,

¹⁰ C. D. Broad argues in this way for the *prima facie* justification of claims of religious experience in 'Arguments for the Existence of God' in his *Religion, Philosophy and Psychological Research* (Routledge & Kegan Paul, 1953). Many philosophers have made the obvious point that no experience entails the existence of its purported object, but they seem to ignore the question whether it is *prima facie* evidence for it. Those who do discuss this question usually conclude that it is not *prima facie* evidence. For example, T. Penelhum in *Religion and Rationality* (Random House, 1971) claims (p. 168) that 'an argument beginning with the occurrence, as psychological fact, of a given experience or set of experiences and ending with the ascription of them to a divine cause is either a poor explanatory hypothesis or a circular argument'. Such writers do not seem to me to be aware of the sceptical bog in which failure to accept the Principle of Credulity for other experiences will land them. And, if it is all right to use it for other experiences, they need a good argument to show that it is not all right to use it for religious experiences.

or that there are present angels or Ultimate Reality, is not good grounds for supposing that things are thus.

The first argument is that our supposing that the way things seem is the way they are is not a fundamental principle of rationality, but itself requires inductive justification, and that that inductive justification is available in the ordinary cases but not in the religious cases.¹¹ More particularly, a philosopher may claim that the fact that it appears that *x* is present is good grounds for supposing that *x* is present only if we have evidence that, when in the past it appeared that *x* was present, it proved so to be; or at any rate the assumption that *x* is present has proved a successful assumption from which to work. Hence, the philosopher might argue, it is all right to take what looks like a table as a table, because our past experience has shown that such appearances are not misleading; but he might go on to question whether we had the kind of inductive evidence that was necessary to justify taking religious experiences seriously.

One difficulty with this view is that it is ordinarily supposed that people are justified in taking what looks like a table to be one even if they do not at the same time recall their past experiences with tables, and even if they cannot immediately do so. So the principle would have to say that our justification for taking what looked like a table to be one was that we could remember such past experiences if we tried hard enough. It will not do to say that our merely having had the past experiences suffices to justify our present inference, whether or not we can remember those experiences. For, if a claim is to be justified inductively, we must in some sense 'have' the evidence of past performance in order to be justified in making the inference. But then, an induction from past experiences to future experiences is justified only if we recall our past experiences correctly. And what grounds have we got for supposing that we do? Clearly not inductive grounds—an inductive justification of the reliability of memory claims would obviously be circular. Here clearly we must rely on the principle that things are the way they seem, as a basic principle not further justifiable; that we seem to have had such-and-such experiences is in itself good grounds for believing that we had. If you require that other people also shall have had the experience of it

¹¹ 'In order to infer the divine from an apparition we should have to have experience of a connection between them in the way in which we do have experience of the connection between smoke and fires' (A. MacIntyre, 'Visions', in A. Flew and A. MacIntyre (eds.), *New Essays in Philosophical Theology* (SCM Press, 1955), 257).

seeming to them that there was a table present before you are justified in trusting the deliverances of your own sense on this matter, then what is your justification for believing that other people have had such experience? Clearly (as well as assuming that other people probably tell you the truth) you again rely on the principle that things are as they seem to be (that other people seem to you to have said that they had these experiences gives you justification for believing that they did say this). The principle that the rational person supposes that, in the absence of special considerations in particular cases, things are the way they seem to be or to have been can be given inductive justification on the basis of past experiences only if these latter are held to be trustworthy solely on the basis of the Principle of Credulity itself. The Principle of Credulity is a fundamental principle of rationality. And, while the degree of probability that experience or apparent memory confers on what is apparently experienced or remembered varies with the strength of the experience or memory, clearly, unless most experiences and the stronger memories count (in the absence of counter-considerations) as probably true, we would have very little of the ordinary knowledge of the world that we think we have. Unless we allow this principle to have considerable force, we quickly find ourselves in a sceptical bog, in which we know hardly anything.

Another difficulty with the view of the first argument is that its suggested principle clearly needs modification to deal with cases where the subject has no past experience of x 's but does have experience of properties in terms of which x is defined. Thus a centaur is defined as a being with the head, trunk, and arms of a human, and the body and legs of a horse. A subject has seen humans and horses, but not centaurs before. It then appears to him that a centaur is present. Is that good reason to suppose that it is? Surely yes. So the principle behind the first argument had better be modified to read: the fact that it appears that x is present is good grounds for supposing that x is present only if we have evidence that when in the past it has appeared that x or any properties by which x is defined is present they have proved so to be, or at any rate the assumptions that they were present proved successful assumptions from which to work. But then the argument is quite inadequate to rule out taking religious experiences seriously. For 'God', like 'centaur', is defined in terms of properties of which most of us have had experience. He is defined as a 'person' without a 'body' who is unlimited in his 'power',

'knowledge', and 'freedom', and in terms of other similar properties, of all of which we have had mundane experience. Someone might well, through visual, auditory, tactual, etc. experience of recognizing persons of various degrees of power, knowledge, and freedom be able to recognize when he was in the presence of a person of unlimited power, knowledge, and freedom. Indeed, it is plausible to suppose that someone might be able to recognize extreme degrees of these qualities, even if he could not so easily recognize lesser degrees straight off without inductive justification. So, once the inevitable modification is made to the first argument, whatever its merits, it has no force against the claims of religious experience.

The second attempt to restrict the application of the Principle of Credulity allows that the principle holds for 'sensible' characteristics and relations (without needing inductive justification) but denies that (in the absence of inductive justification) it holds in any other cases. One writer who has thus restricted the principle is Chisholm. He claims that, whenever we take something to have a certain sensible characteristic (in my terminology 'property') or relation, we have adequate evidence for the claim that it does have this characteristic (or relation); but that, whenever we take something to have some non-sensible characteristic (or relation), that is not in itself adequate evidence to suppose that it does. And what are these 'sensible' characteristics and relations? Chisholm writes:

The characteristics include being blue, red, green, or yellow; being hard, soft, rough, smooth, heavy, light, hot, or cold; and that of sounding, or making-a-noise. The relations include: being the same, or different with respect to any of the characteristics in question; being more like one object than another with respect to any of the characteristics, or with respect to hue, saturation, and brightness, or with respect to loudness, pitch, and timbre. The class of characteristics and relations also includes the 'common sensibles'—that is, 'movement, rest, number, figure, magnitude'—as well as what is intended by such terms as 'above', 'below', 'right', 'left', 'near', 'far', 'next', 'before', 'after', 'simultaneous', and 'to last', or 'to endure'. In short, the characteristics and relations in question are co-extensive with what Aristotelians have traditionally referred to as the 'proper objects of sense' and the 'common sensibles' and what Reid described as the objects of 'original' perception.¹²

So, according to Chisholm, if something seems (epistemically) to *S* to be brown or square or solid, that is good grounds for believing that it

¹² Chisholm, *Perceiving*, 83.

is. But if something seems to be a table, or a Victorian table, or a ship, or a Russian ship, that is in itself not good grounds for believing that it is. You can have good grounds for believing that something is a table only in terms of it looking brown, square, and solid and in terms of things that look like that having appeared (in the past) to be used for writing on (the notion of 'writing' perhaps being spelt out in terms of 'sensible' characteristics).

Let us say that, if its seeming that an object (or characteristic) x is present is grounds for supposing that it is without need for further justification, then you have a *real experience* of x . But, if this does not hold, then its seeming that an object x is present is an *interpretation* of your experience that stands in need of justification. If you have a real experience of x and if in fact x causes your experience, then you *really perceive* x ; if you conclude that x is present without really experiencing x , then (even if your conclusion is correct and justifiable) you merely *infer* x . Attempts to draw such lines as Chisholm draws between real experience and interpretation, real perception and mere inference, are, of course, as old as the empiricist tradition in philosophy. It is admitted by most of those who draw a line of this kind that even real experience may mislead. You may have a real experience of x , and therefore be justified in supposing that x is present, that x causes your experience, and so that you perceive x , when in fact x is not present at all. In that case you have a delusion, hallucination, or illusion of x , or are merely dreaming that x is present, or some such thing. However, such cases are, on this view, to be distinguished from cases where, although it seems that an object x is present, and you take it so to be when it is not, your mistake is one of misinterpretation of experience—a mistake that you would have been justified in making only if you had other grounds for believing the object x to be present.

That there is such a line to be drawn is a common and seldom argued assumption in many discussions of religious experience. Once the line is drawn, the consequences are evident. For the line always leaves the typical objects of religious experience as matters of interpretation rather than as true objects of real experience. It follows that, even if it seems to you strongly that you are talking to God or gazing at Ultimate Reality, this fact is no reason in itself for supposing that you are. You are having an experience that is properly to be described in a much more mundane way—for example, as the experience of hearing certain noises—which you *interpret* as the voice of God, but

which you have no good reason for so doing unless further evidence is available.

However, no such line as the one that Chisholm attempts to draw can be drawn between real experience and interpretation. For clearly we are justified in holding many perceptual beliefs about objects having non-sensible characteristics that cannot be backed up in terms of beliefs about objects having 'sensible' characteristics. Few would doubt that I am justified in believing that a certain woman whom I see at the other side of a room is my wife. Yet, if asked what it is about the woman I take to be my wife that makes me believe that she is my wife, I would be utterly unable to give a satisfactory answer. I could give only a very vague description of the Chisholmian 'sensible' characteristics by which I recognize her, a description that would fit tens of thousands of other women whom I would not for one moment mistake for my wife. That one can recognize does not entail that one can describe; nor does it even entail that (even if one cannot describe them) one knows what the features are by which one recognizes. I may be justified in claiming that you are tired or angry, just by looking at your face, and yet not know what it is about your face that makes you look tired or angry. Again, I can recognize my wife's voice over the telephone, although I certainly cannot say what it is about the noises that come through the telephone receiver that are especially characteristic of her voice. For senses such as smell and taste most of us have no vocabulary for describing sensible characteristics, other than in terms of the objects that cause them (for example, as 'the taste of tea' or 'the smell of roses'). Asked about the liquid we are drinking 'What is it about it that makes it taste like tea?', we would be stuck for an answer. But that fact casts no doubt on our justification for believing that we are drinking a cup of tea. The fact that it tastes like tea is good reason in itself for supposing that it is—whether or not we can say in more primitive terms what it is about it that makes it taste like tea.

Humans differ in the kinds of objects and properties that they learn to pick out. Sometimes they can pick out and even describe the 'sensible characteristics' of those objects and sometimes they cannot; and, even if they can, the recognition of objects of some kind and their more sophisticated properties may be a more natural process than the description of their sensible characteristics. There is no reason of principle why we should not grow so adept at spotting Russian ships, or Victorian tables, or blue-dwarf stars, or elliptical

galaxies that we can recognize them straight off, without being able to say what it is in the way of Chisholmian sensible characteristics about what we see that makes us identify them as we do.

So this second argument against the original Principle of Credulity fails, and the principle stands. If it seems (epistemically) to *S* that *x* is present, that is good reason for *S* to believe that *x* is present, in the absence of special considerations—whatever *x* may be. And it is good reason too for anyone else to believe that *x* is present. For, if *e* is evidence for *h*, this is a relation that holds quite independently of who knows about *e*. However, how things seem to *S* is clearly something of which *S* knows without inference, whereas others need *S*'s testimony about how things seem to him in order to learn of his experiences. Our justification for relying on *S*'s testimony about his experiences is an issue to which I shall come later in the chapter.

From all this it follows that, if it seems to me that I have a glimpse of Heaven, or a vision of God, that is grounds for me and others to suppose that I do. And, more generally, the occurrence of religious experiences is *prima facie* reason for all to believe in that of which the reported experience was purportedly an experience.

Special Considerations that Limit the Principle of Credulity

It is time to list the special considerations that operate in particular cases and give to the subject or to others grounds for holding that, although his experience was that it seemed to him that *x* was present (and so he is inclined to believe that *x* was present), really *x* was not present. They are considerations that, when added to the report of the experience, prevent it from making it probable that *x* was present. Put symbolically, with *e* as 'it seems to *S* that *x* is present', *h* as '*x* is present', and *k* as irrelevant background knowledge, they are considerations *c* such that although $P(h|e \& k) > 1/2$, $P(h|e \& k \& c) \leq 1/2$. If it seems to *S* that *x* is present, *S* is inclined to claim, at any rate to himself, that he perceives *x*. I shall describe what I am doing as listing the considerations that defeat the perceptual claim that *S* is inclined to make to himself. Having listed these considerations, we can then see whether they will normally be able to show that religious experiences are not to be taken at their face value.

There are basically four kinds of special consideration that defeat perceptual claims. The first two show that the apparent perception

was of a kind with others that proved in the past not to be genuine perceptions. First, one may show that the apparent perception was made under conditions or by a subject found in the past to be unreliable. Thus one may show that *S*'s perceptual claims are generally false, or that perceptual claims are generally false when made under the influence of LSD, which is good inductive grounds for believing that a particular new perceptual claim made by *S* or made under the influence of LSD is false. Secondly, one may show that the perceptual claim was to have perceived an object of a certain kind in circumstances where similar perceptual claims have proved false. Thus, if it seems to *S* that he has read ordinary-size print at a distance of a hundred yards, we can test him on a number of other occasions and see if he is able to read what is written at that distance; and if he is not we have good inductive evidence that the original claim was false.

One variant of the second kind of consideration is where one shows that *S* has not had the kind of experience that has been found empirically necessary to make a probably true perceptual claim of the kind in question. One might, for example, have evidence that only those who have actually tasted tea before and been told what it was that they were tasting can (except by accident) make true perceptual claims to be tasting tea. It is not at all evident *a priori* what kinds of experience people need in order to make probably true perceptual claims of different kinds; and clearly people vary enormously in this respect. Perhaps some of us who have only smelled tea but never tasted it before could in consequence recognize it by its taste at first tasting. Some of us can recognize people by descriptions. Some of us can recognize people only when we have seen them before.

The third and fourth considerations are ones concerned with the particular perceptual claim that do not involve inductive inference from the failure of similar claims. Since to perceive *x* is to have one's experience of its seeming that *x* is present caused by *x*'s being present, one can challenge a perceptual claim to have perceived *x* either by showing that it is very very probable that *x* was not present or by showing that, even if *x* was present, *x*'s presence probably did not cause the experience of its seeming that *x* was present. The third consideration then that defeats a claim to have perceived *x* involves showing that on background evidence it is very very probable that *x* was not present. Now I suggest that in this case it is not enough that the background evidence makes it more probable than not that *x* was

not present. It has to make it very very improbable that x was present if it is to outweigh the force of S 's experience sufficiently for it to remain more probable than not that S was not present. For, after all, most of the things that we think that we see are on background evidence less probable than not. It may seem to me, when I go to London, that I see Jones walking along the other side of Charing Cross Road. I may believe *a priori* that it is more probable than not that he is in Dover, where he lives; and that, even if he is in London, the odds are against his being in Charing Cross Road at that particular moment. But my experience suffices to outweigh this background evidence. We would indeed be imprisoned within the circle of our existing beliefs, if experience did not normally have this force. However, background evidence may make it very, very improbable that x is present—for example, because it makes it very improbable that x exists at all, or very very probable that he is somewhere else. If it is very probable on background evidence that John is dead, then it is very, very improbable that he is walking along the other side of Charing Cross Road at this moment; and my experience does not by itself suffice to push the latter into the category of the probable.

A similar point arises with respect to anyone who claims to have observed an entity of a kind rather different from those already known. If you claim to have seen a dodo on Mauritius, then, if *a priori* it is probable, although not very, very probable, that dodos became extinct in the seventeenth century, your perceptual claim remains overall probable. But, if you claim to have seen a man 20 feet tall getting out of a space machine, then what you claim to have seen has such prior improbability that your claim needs backing up before it becomes probable.

There are various ways in which it can be shown that very, very probably x was not present. One may show that very probably x does not exist, or was in some other particular place at the time, or show that very, very probably x was not at the place in question more directly, in particular by showing that other observers who would very probably have had the experience of its seeming to them that x was there, if x had been there, did not do so. If I claim to have seen John in the corridor, my claim may be defeated by showing that, although there were many others in the corridor with eyes functioning correctly, who were looking out for John and knew what he looked like, they did not have the experience of its seeming that John was there. The application of this test requires us to know what sense

organs and training you need, and how attentive you need to be to perceive the object in question. We normally have or can get this evidence as a result of seeing what is needed for the detection of similar objects. Yet, even when we know what kind of observers would very probably have had the requisite experience if x had been there, their not having the experience does not prove conclusively that he was not there. There will always be a doubt about whether the observers were sufficiently attentive, were their sense organs working properly, etc. But clearly, the more observers apparently rightly positioned with apparently the right sense organs and concepts who fail to observe x , the less probable it is that x was there. But, if several observers have the requisite experience, even if many others do not, that makes it probable that x was there. If x is an object very different in kind from others that have been investigated, we shall not have knowledge of which sense organs, concepts, and degree of attentiveness are needed for perceiving x ; and so we cannot have evidence against the claim, from those who would have had certain experiences if x had been present, but did not have such experiences. But clearly, in so far as some do have the requisite experience, that makes it likely that x is there, even if some others do not have the experience, and we do not know why they do not have the experience (so long as we do not have positive reason for supposing that they would have had it if x had been there). The possibility always exists in such cases that the object is not a public object at all.

However, I am inclined to add that, if we do not know which observers could have been expected to have had an experience apparently of x if x had been there, that *somewhat* lessens the evidential force of an apparent perception—but only somewhat. This is because in that case we cannot have as confirming evidence in favour of the perceptual claim the fact that all witnesses who would have had the requisite experience, if the object was present, did so. If your claim could have been disconfirmed by certain phenomena, but the phenomena are shown not to occur, that very fact confirms the claim. If we do not know which observations count against a claim, there cannot be a failure to make such observations, which counts in favour of the claim. However, I stress the words ‘only somewhat’. For clearly, if three witnesses saw the man in the distance, or the rainbow, or heard the high note, or felt the tremor, and three did not, that is substantial evidence in favour of the occurrence of the object reported, even if we do not know why others were unable to detect it.

Fourthly, *S*'s claim to have perceived *x* may be challenged on the grounds that, whether or not *x* was there, *x* was probably not a cause of the experience of its seeming to *S* that *x* was there. One obvious way in which this can be done (without casting any doubt on other of my perceptual claims) is by showing that (probably) something else caused the experience. We challenge the claim by producing a causal explanation of why it seemed to *S* that *x* was there, which does not involve *x* at any stage. If you show me the actor who was dressed up to look like John and who walked down the corridor, I realize that the experience of its seeming to me that I had seen John was probably caused by the actor, and so that I have no grounds for believing that John was in the corridor.

These four ways that I have listed are all ways in which we or others challenge ordinary perceptual claims. Let us take another example of how they work. Suppose that I wake up startled and, because it so looked to me, I claim to have seen a man dressed in a toga looking up the chimney, who disappeared straight away when I looked at him. You may deny my claim by pointing out (1) that I have become addicted to LSD, which has led me to claim to see things that are not there quite a bit recently; or (2) that tests on other occasions show that I cannot recognize a toga when I see one; or (3) that my wife was awake, but did not see the man, and so probably he was not there; or that I have good inductive reason from long experience to believe that men cannot just vanish, and so that it is unlikely that things happened as I claimed; or (4) that there was on the walls a pattern of shadows that could naturally have been interpreted as a man in a toga, and so there was a cause of my experience of 'its looking to me as if there was a man in a toga' other than a man in a toga. Yet, if none of these challenges can be made, my claim ought to be accepted.

If one of these challenges works, the onus of proof shifts. The subject now needs to prove his experience to be genuine (that is, that he perceives what he thinks he does). He can do this by producing positive evidence in support of his perceptual claims of the types adduced in any of the other possible challenges. Thus, to revert to the toga example, I have good evidence from long experience to show that men in togas do not just vanish up chimneys. So a challenge of type (3) counts substantially against my perceptual claims. But I may be able to show (1) that I am in general a very reliable witness, (2) that I know what togas look like, and that I am reliable in reporting things immediately on waking up, (3) that my wife also saw the man,

and (4) that there was nothing else that might plausibly have been mistaken for a man in a toga. In such cases we must weigh the conflicting evidence. Our criteria for doing so (for example, for how much weight we should give to considerations of type (4) against considerations of type (3)) are by no means clear. There is a large border region of possible cases in which it is unclear which way the balance of probability tilts, but clearly, if there is enough positive evidence of the kind described, it suffices to outweigh an initially successful challenge to a perceptual claim.

Note that some experiences are very much more forceful than others—some experiences are very clear and unavoidable and leave a very strong impression; it may have seemed that an object was very close and that I could not be mistaken about it. In so far as an experience has such a character, clearly it needs more in the way of challenge to defeat it. If I am really convinced that I saw John in the corridor, it needs to be *very* probable that it was the actor who caused my experience before my perceptual claim becomes improbable. And, at the other end of the spectrum, some apparent memories are so faint that they yield only a probability much less than $\frac{1}{2}$ that what is apparently remembered is true; they need supporting evidence.

Challenges from these Special Considerations to Religious Experiences

How far are the above challenges available to defeat the claims of those who claim to have experienced God, or Poseidon, or Ultimate Reality? The first challenge may defeat a few such claims, but it is hardly generally available. Most religious experiences are had by people who normally make reliable perceptual claims, and have not recently taken drugs.¹³ The second challenge would consist in showing that normally religious perceptual claims were unreliable. If there was a good proof of the non-existence of God or anything similar, then of course that could be done. But the point here is that the onus of proof is on the atheist; if he cannot make his case, the claim of religious experience stands.

¹³ For detailed justification of my assertion that this first challenge is not generally available to defeat the claims of religious experience, see C. Franks Davis, *The Evidential Force of Religious Experience* (Clarendon Press, 1989), ch. 8.

The issue arises with the claims of religious experience as with the claims of miracles, whether there is a general proof of their unreliability in the fact that so many of them conflict with each other. Thus, Flew:

Religious experiences are enormously varied, ostensibly authenticating innumerable beliefs many of which are in contradiction with one another . . . The varieties of religious experience include not only those which their subjects are inclined to interpret as visions of the Blessed Virgin or senses of the guiding presence of Jesus Christ, but also others more outlandish presenting themselves as manifestations of Quetzalcoatl or Osiris, or Dionysus or Shiva.¹⁴

Now, of course, devotees of different religions describe their religious experiences in the religious vocabulary with which they are familiar. But in itself this does not mean that their different descriptions are in conflict—God may be known under different names to different cultures (as both Old and New Testaments acknowledge—see Exodus 6: 2–3 and Acts 17: 23). Likewise a Greek's claim to have talked to Poseidon is not necessarily in conflict with a Jew's claim to have talked to the angel who watches over the sea; it is so only if to admit the existence of Poseidon is to commit one to a whole polytheistic theology, and there is no need to suppose that generally it is.

Admittedly, sometimes the giving of one description to the object of religious experience does carry commitment to a doctrine regarded as false by devotees of another religion. Claiming to have experienced the heavenly Christ commits one to a belief in an Incarnation that an orthodox Jew would not admit. But in these cases, if the opponent of the doctrine can produce good grounds for regarding the doctrine as false, that is reason for the subject of the experience to withdraw his original claim. Among those grounds may be that others have had conflicting experiences and that their experiences are more numerous and better authenticated; but there may be many grounds of other kinds as well. The subject of the religious experience need not in such a case withdraw his original claim totally; he need only describe it in a less committed way—for example, claim to have been aware of some supernatural being, not necessarily Dionysus (as originally claimed). The fact that sometimes (and by no means as frequently as Flew suggests) descriptions of the object of a religious experience are in conflict with descriptions of the object of another

¹⁴ A. Flew, *God and Philosophy* (Hutchinson, 1966), 126–7.

religious experience means only that we have a source of challenge to a particular detailed claim, not a source of scepticism about all claims of religious experience. Babylonian astronomers reported the movements of holes in the firmament; Greek astronomers reported the movements of physical bodies in the heavens. The conflict between them did not mean that there were no things in the sky of which both groups were giving further descriptions. But it did mean that the perceptual claims of each group constituted arguments against the perceptual claims of the other group; and, given that the perceptual claims of both groups were equally weighty in number and conviction, that further arguments were needed to adjudicate between them. Eventually Babylonian astronomers had to admit that they had somewhat misdescribed what they saw. But this process need hardly lead to general scepticism about astronomical observation; nor need the similar process in religion.¹⁵

However, it does follow that, if there were a substantial number of religious experiences that entailed the non-existence of a particular supernatural being, that would cast significant doubt on the credibility of claims to have perceived that being. There certainly does not exist evidence of this kind sufficient to cast significant doubt on the

¹⁵ There was a tradition in the description of mystical experience, typified by W. T. Stace (see his *Mysticism and Philosophy* (MacMillan, 1961)), which claimed that all mystical experiences are in essence the same. It was simply that Christians, Muslims, Buddhists, etc. read into them their own different interpretations (or sometimes described them without interpretation). We have seen (p. 308) that any distinction between the real experience and the interpretation superimposed upon it is going to be hard to justify, if the experience is supposed to be of some external object. However, the natural way to interpret Stace's claim is as a claim that all subjects have in essence the same kinds of sensory experiences, i.e. the same kinds of experience described comparatively (see pp. 294–5), which give rise to different kinds of experience described epistemically, i.e. led the subjects to hold different beliefs. Stace distinguished mystical experiences from other religious experiences, such as voices and visions, but even so his claim is disputable. (See the discussion by S. T. Katz, 'Language, Epistemology, and Mysticism', in Katz (ed.), *Mysticism and Philosophical Analysis*.) Yet, even if the claim were correct, there would only be a situation of conflicting experiences, giving rise to the situation described in the text, if the subjects were to describe their experiences as experiences apparently of objects that could not both exist together (e.g. an omnipotent God and an omnipotent Devil). The fact that one subject describes his experience solely in comparative terms ('I had the sort of experience which you have if you look at a very bright light'), and another describes his experience in a more committed way ('I had an experience apparently of a very Pure Being'), does not give rise to this situation. We can accept the claims of the subjects of both experiences. The Principle of Credulity suggests that we should draw an ontological conclusion from the second description, but not from the first—i.e. it suggests that there is some object that the first subject fails to recognize, but that the second subject does recognize.

credibility of claims to have perceived God. Religious experiences in non-Christian traditions are experiences apparently of beings who are supposed to have similar properties to those of God, or experiences apparently of lesser beings, or experiences apparently of states of affairs, but hardly experiences apparently of any person or state whose existence is incompatible with that of God. If there were vastly many experiences apparently of an omnipotent Devil, then that sort of evidence would exist; but there are not such experiences.

Another general objection under the heading of this second challenge might be that those who make claims of religious experience have not had the kind of experience that is needed to make claims of this kind that are probably true. The argument might be that your claim to have recognized a person is likely to be correct only if you have previously perceived that person (and been told who he is) or if you have previously been given a detailed description of his appearance (appropriate to the modality of sense by which you claim to recognize him—for example, you would need previously to have been given a description of his visual appearance, before your claim to have seen him is likely to be correct).¹⁶ But this argument seems clearly mistaken. I can come to recognize people whom I have never perceived before after being given descriptions of them that can hardly be regarded as descriptions of their appearance appropriate to the modality of sense involved. Thus I may be told that Smelinowski is the only Ruritanian with a really English sense of humour, or that General Walters is the most commanding personality whom I am ever likely to meet in a lifetime; and these descriptions may be perfectly adequate for me to be able to recognize Smelinowski or General Walters. The description of God as the one and only omnipotent, omniscient, and perfectly free person may indeed suffice for someone to recognize

¹⁶ In the case of visions of persons now dead who formerly lived on earth, such as the saints, one must be careful not always to suppose that the person will continue to look and otherwise appear as she did on earth, in her more superficial aspects. No doubt character and some memory may be expected to continue, as more intimately connected with personal identity; but one would not expect dress necessarily to be the same and linguistic competence to be confined to the tongue that the person spoke on earth. Thus sceptics are apt to deny the claim of some Portuguese peasant to have seen Mary, on the grounds that the description that the peasant gives of the way Mary was dressed does not fit the way Mary used to dress in Palestine, but corresponds closely to the way she is pictured on the walls of Portuguese churches. That seems to me to count not at all against the peasant's claim. For, if Mary has survived death, what reason is there to suppose that she has now to dress the way she did in Palestine? If she is to manifest herself in bodily form, the obvious way for her to dress is the way in which she would be recognized by those to whom she appears.

him—by hearing his voice, or feeling his presence, or seeing his handiwork, or by some sixth sense. And, as I claimed earlier, even if some of us are not very good at recognizing power, or knowledge, or freedom in the human persons whom we meet, we might well be able to recognize extreme degrees of these qualities when we cannot recognize lesser degrees. Nor, of course, if someone has the ability to recognize something, does it follow that he can imagine in advance what the experience of recognition would be like. What you tell me about an entirely new colour may enable me to recognize that colour when I see it, even if I cannot visualize the experience in advance. This objection certainly has no compelling force, but it does seem to me to have some small force. Great power, knowledge, and freedom are not characteristics that we easily learn to recognize by hearing a voice, or by seeing some object that might be an agent's handiwork, or by feeling. And *some* mild suspicion is cast on a subject's claim to have recognized an agent with these qualities by the qualitative remoteness of his previous experiences from what he claims to have detected—but, for the reasons that I have given, only *some* mild suspicion.

The third challenge to a claim of religious experience would consist, in the case of a purported experience of God, in a demonstration that very, very probably God was not present to be perceived, and so the subject could not have perceived him. But, if there is a God, he is everywhere. He is only not present if he does not exist. So, to use this challenge (barring a consideration below), you have to prove that very, very probably there is no God, and, as stated above, the onus is on the atheist to do so. (For religious experiences of lesser persons or things, for example, of Mary or Poseidon, you need only to show that very very probably the person is not where the subject said that he was). As we have seen, it will not do to show that some people with similar equipment and concepts to those who do have experiences of God do not have such experiences. For we do not know that all persons with certain equipment and concepts could be expected to have experiences of God, if he was there. Clearly, if he so chose, an omnipotent God could cause a private experience, in the way that a table could not. Clearly too some people with similar equipment and concepts do have experiences apparently of God. But, as we have also seen, the fact that the absence of apparent perceptions of God by others has no tendency to show that God is not present has the consequence that the original perceptual claim is, on its own, somewhat less evidence of the existence of God.

The fourth challenge would consist in showing that the religious experience probably had a cause other than its purported object—for example, God. But this is a particularly awkward challenge to apply when we are dealing with a purported experience of God—as opposed to, say, Mary or Poseidon. My religious experience may or may not be caused immediately by some brain event. In either case, an apparent experience of x is an experience of x if x belongs to the causal chain that brings about those events by its presence where it appears to be. Since Mary is not omnipresent, she will appear to be here rather than there; and, since she is not the sustainer of the world, she can be responsible only for some of the causal processes within it. It is possible to show that the causal chain that produced my experience involved events only at places other than where she appeared to be or involved causal processes that would have operated whether or not she had been there. It is possible to show that she was not involved in the process, without in any way tending to disprove her existence. But, if there is a God, he is omnipresent and all causal processes operate only because he sustains them. Hence any causal processes at all that bring about my experience will have God among their causes; and any experience of him will be of him as present at a place where he is. And so, if there is a God, any experience that seems to be of God, will be genuine—will be of God. He may bring about that experience either by intervening in the operation of natural laws (producing an event other than natural laws would ordinarily produce) or by sustaining their normal operation. These latter natural laws would be ones that produce religious experiences in people with certain beliefs or brain states in certain circumstances—for example, when fasting, or in dark churches. If fasting or whatever has a tendency to cause religious experiences, that may be either because, like some eye salve, it helps one to see what is there or because it makes one seem to see what is not there. The mere fact that a religious experience apparently of God was brought about by natural processes has no tendency to show that it was not veridical. To show this, you need to show that God did not cause these processes. That can be attained only by showing that there is no God—for, if he exists as defined, clearly he is responsible both for the normal operation of natural laws and for any occasional violation.¹⁷

¹⁷ This simple point is well made in W. J. Wainwright, 'Natural Explanations and Religious Experience', *Ratio*, 15 (1973), 98–101.

The upshot of all this is that there are two qualifications that somewhat diminish the evidential force of religious experience apparently of God. One is the qualitative remoteness of subjects' past experiences from what they claim to have recognized—namely, God. The other is the fact that the absence of apparent perceptions of God by others, however positioned, has no tendency to show that a particular subject's apparent perception of God was not veridical. But for these qualifications, I would have concluded that a religious experience apparently of God ought to be taken at its face value unless it can be shown on other grounds that very, very probably God does not exist. The qualifications lead me to modify the 'very, very probably', and to suggest that a religious experience apparently of God ought to be taken as veridical unless it can be shown on other grounds significantly more probable than not that God does not exist.

It will be much more difficult to resist challenges to the claims of religious experiences apparently of supernatural persons or things less powerful than God. As we have seen, they are open to the fourth challenge. And they are far more open to the third challenge, that very, very probably the being apparently experienced does not exist; since there is no natural theology available to give some probability to the existence of supernatural beings other than one who sustains the whole universe. Likewise, claims to have experienced an impersonal source of all things are, given my arguments that a personal cause of the universe has greater simplicity and so greater intrinsic probability than an impersonal one, more open to the third challenge than are claims to have experienced God.

To return to the main thread—the issue with regard to the evidential force of experiences apparently of God is just how improbable is the existence of God, on background evidence, that is on that general knowledge of the world, its operations, and its history, together with reports of miracles, considered in Chapters 7–12. Is the existence of God so improbable that an experience apparently of God should not be taken at its face value?

Even if the answer to this question were yes, there is a crucial further consideration. We saw that any apparent perception of what is on background evidence too improbable to be believed (that is, for what is apparently perceived to be improbable overall on all the evidence) may become credible if backed up by positive evidence that the experience is genuine. This positive evidence can take the

form of others' having corroborating experiences. It is a further important principle of rationality that, in the absence of reason for challenge, we should believe what people tell us about their experiences.

The Principle of Testimony

The Principle of Credulity is concerned with a subject's grounds for believing that things are as they seem to him. Clearly in ordinary life we use also a wider principle. Other things being equal, we believe that what others tell us that they perceived probably happened. By 'other things being equal' I mean in the absence of positive grounds for supposing that the others have misreported or misremembered their experiences, or that things were not in fact as they seemed to those others to be. Clearly most of our beliefs about the world are based on what others claim to have perceived—beliefs about geography and history and science and everything else beyond immediate experience are thus based. We do not normally check that an informant is a reliable witness before accepting his reports.

The assumption that things are (probably) as others claim to have perceived them has two components. One is the Principle of Credulity—that (in the absence of special considerations) things are (probably) as others are inclined to believe that they have perceived them. The other component is the principle that (in the absence of special considerations) the experiences of others are (probably) as they report them. This latter principle I will call the Principle of Testimony.¹⁸ I used this principle in claiming (on the basis of what they tell us) that very many people have religious experiences. The special considerations that lead us to doubt a subject's reports of his experiences are evidence that generally or in matters of a particular kind he misremembers or exaggerates or lies. But, in the absence of such positive evidence, we have good grounds to believe what others tell us about their experiences.

In general there are no special considerations for doubting what subjects report about their religious experiences, although sometimes

¹⁸ See my *Epistemic Justification*, 123–7, for discussion of the issue of whether the Principle of Testimony is a fundamental *a priori* principle, or whether it is a consequence of the application of other more fundamental inductive criteria to contingent evidence about what people say when.

there are such considerations. There may be evidence from what he says about other matters on other occasions that a subject is a habitual liar, or tells a lie whenever he can gain attention by so doing, or exaggerates, or misremembers. In these cases his reports of his religious experiences are to be viewed with scepticism. But this is not the normal situation.

One ancient test that may be used where there is doubt about the veracity of a subject's report of some religious experience is to see whether the subject's lifestyle has undergone a change. Suppose that Jones claims to have had an overwhelming experience that strongly seemed to him to be of God. If he really did have that experience, one would expect his faith in God to be much deeper and this to make a great difference to his way of living. Our grounds for this expectation are that, if it really seems to you that you have seen x , then you will believe that x exists. If you believe that x exists, that will make a difference to your behaviour in appropriate circumstances. If it really seems to you that you have talked with God, then it will be much more natural for you to act as if there is a God; (unless you have a strong desire to spite God) prayer, worship, and self-sacrifice will be more natural occupations.

Since (probably) others have the experiences that they report, and since (probably) things are as a subject's experience suggests that they are, then (with some degree of probability) things are as others report. However, the degree of probability is less in the conclusion than it is in either of the premisses. If p is evidence for q , and q is evidence for r , then p is normally less evidence for r than it is for q . If the fingerprint is evidence of Jones's presence at the scene of the crime, and Jones's presence at the scene is evidence that he did the crime, the fingerprint is less evidence for Jones's having done the crime than it is for his presence at the scene of the crime. Hence, if S reports that it seems (epistemically) to S that x is present, then that is reason for others also to believe that x is present, although not as good a reason as it is for S if in fact he is having the experience that he reports. However, clearly it is quite a good reason. As we have seen, our whole system of beliefs about the world beyond our immediate experience is based on trusting the reports of others. And, of course, in so far as a number of others give similar reports, that greatly increases their credibility. There are large numbers of people both today and in the past who have had religious experiences apparently of the presence of God and that

must make it significantly more probable that any one person's experience is veridical.

More Specific Experiences

Many people's experiences are experiences not just of the apparent presence of God or some other supernatural thing, but of God or that thing apparently doing something or telling them something. As before, the Principle of Credulity applies. These perceptual claims, like all others, are to be believed in the absence of special considerations. As before, the first and second kinds of consideration are only very likely to be available to defeat some such claims, and not all claims. But the third and fourth considerations are likely to be far more generally available. There may be arguments to show that it is very, very improbable that, if there is a God, he would have said or done what the subject claims to have experienced him saying or doing. This would be normally because saying or doing those things would be (very probably) incompatible with the perfect goodness of God. Claims that God told subjects to lie, rape, or torture will be immensely improbable. And the counter evidence to a claim of one subject that God has revealed to him some truth may well consist in the claim of another subject that God has revealed to her some truth incompatible with the former. And the fourth consideration is now also available. We do not need to hold that God does not exist in order to hold that causal processes that he sustains lead to people having false beliefs about what he has told them to do. Maybe they come to have these beliefs as a result of a free decision of their human religious instructors to give them a certain understanding of what God is like that causes them to have experiences that God has told them to do certain things. Then God's telling them to do certain things would not be part of the cause of their having the experience of him having told them to do these things.¹⁹

¹⁹ While a perfectly good God cannot lie, and so will not utter incompatible propositions (e.g. telling one person that he was incarnate in Christ, and another person that he was not incarnate in Christ), it does not follow that he will not give different people commands, both of which cannot be executed successfully. The mere fact that a subject *A* has an experience apparently of God telling him to do *X*, and subject *B* has an experience apparently of God telling him to stop *A* doing *X*, does not automatically entail that one or other experience is not veridical. Thus a Muslim may have an experience of being told by

I drew attention earlier to the lower evidential worth (in comparison with experiences apparently of God) of apparent experiences of supernatural beings and things other than God; and I have now drawn attention to the lower evidential worth of more specific religious experiences. Nevertheless, experiences of both these kinds may have indirect force in providing evidence for the existence of God. An experience apparently of God doing something contains an element (the apparent presence of God) that is not open to the stronger objections to its veridicality. And some of the experiences of other beings and things are experiences of persons or things that are more likely to exist if there is a God than if there is not. Thus an experience of the Blessed Virgin Mary is an experience of a person as present, who is much more likely to be present if there is a God who has preserved her in existence after death than if there is no God. In so far as an experience is of this kind, its occurrence is evidence for the existence of something that is in its turn evidence of the existence of God—but clearly much less evidence than is an experience apparently of God. Bearing these points in mind, I now revert to the main kind of religious experience, an experience apparently of the presence of God (which may be an element in a more specific experience).

Conclusion

One who has had a religious experience apparently of God has, by the Principle of Credulity, good reason for believing that there is a God—other things being equal—especially if it is a forceful experience. We have seen that the only way to defeat such a move is to show

God to defend Jerusalem against the infidel, while a Christian has the experience of being told by God to attack it. This might be explained compatibly with God having really issued both commands by the fact that, in the course of human history as a result of factors for which humans are much to blame, Muslims and Christians have come to have different and very limited understandings of God. God is very anxious that human understanding of God should develop through human experience, effort, and cooperation, and should not always be the result of divine intervention; yet he is also very anxious that at any time in history people should live and die by the ideals that they then have. He therefore tells the Muslim and the Christian each to live by his current ideals, knowing that the experience of so doing may lead them to a deeper understanding. On a human level, a sage might well sometimes give to each of two persons who sought his advice the advice to oppose the other, thinking it for the good of both that they should seek to develop their independence and authority.

that it is significantly more probable than not that there is no God. If the prior probability of the existence of God on the evidence discussed in Chapters 7–12 is significantly less than half, then an individual's religious experience needs backing up by the testimony of others to having had similar experiences. That testimony is evidently available. Only if the prior probability of the existence of God is very low indeed will this combined weight of testimony be insufficient to overcome it. Everything turns on just how improbable on background evidence is the existence of God. To that issue I will come in the final chapter. One who has not himself had an experience apparently of God is not in as strong a position as those who have. He will have less evidence for the existence of God; but not very much less, for he will have testimony of many who have had such experiences.

My conclusion about the considerable evidential force of religious experience depends on my Principle of Credulity that apparent perceptions ought to be taken at their face value in the absence of positive reason for challenge. This principle is a very fundamental and very simple principle for the interpretation of experience. It is because of its fundamental and simple character that I did not need to interpret the issues of this chapter via the apparatus of Bayes's theorem. However, since by the argument of this chapter, with h as 'there is a God', e as the evidence of very many religious experiences, and k the phenomena considered in previous chapters, $P(h|e \& k) > P(h|k)$, we can use the relevance criterion in reverse to infer that $P(e|h \& k) > P(e|k)$, and so $> P(e|\sim h \& k)$. Religious experiences apparently of God may be caused by the normal operation of laws of nature, deriving from the most fundamental laws. That the fundamental laws are such as to produce such religious experiences under certain circumstances will then be something 'too big' for science to explain. If these experiences are not produced by the operation of laws, they will be something 'too odd' for science to explain. There is a certain probability that (in an orderly world containing humans and animals, etc.) the laws of nature might be such as to lead to people having such experiences anyway. But if there is a God, there is a greater probability that they will have such experiences, either through God making the laws such as to produce such experiences or by his causing the experiences directly. There is this greater probability because, if people have a basic ability to detect how things are, it is more likely that they will have an experience

apparently of x (that is, an experience inclining them to believe that x is present), if x is present than if x is not present. And that they do have this basic ability is what in effect the Principle of Credulity is claiming. And also, we saw at the beginning of this chapter, there is some reason to suppose that God would give to some people experiences apparently of himself.²⁰

²⁰ William Alston's *Perceiving God* (Cornell University Press, 1991) constitutes a sustained defence of the rationality of believing what we seem to experience (in a wide sense, that is, perceive) in the course of 'Christian mystical practice' (CMP)—that is, in the course of a life of prayer guided by Christian beliefs, which enable us to detect which experiences are veridical and which illusory. Alston wishes to locate religious experiences within a particular practice because of his commitment to reliabilism, the view that a belief is justified if the 'doxastic practice' (e.g. normal sense perception, induction, memory, etc.) that produces it usually produces true beliefs; and so he has to identify the practice before the question of justification becomes resolvable. On this view we can only show that a particular belief is justified if we can show that the practice that produced it usually produces true beliefs. That we cannot do for CMP or for any other practice without relying on the practice itself (and perhaps also to a lesser extent on other practices) to show which beliefs are true. But he claims that the fact that CMP is a socially accepted practice, which we do not have sufficient reason to show to be unreliable, means that it has earned its right to be taken seriously. And this means, he claims, that trusting its deliverances is 'reasonable' or 'rational'; that is, it is reasonable to believe that its deliverances are epistemically justified. But, crucially, the rationality is of 'a kind that doesn't entail a likelihood of truth' (ibid. 181). So, despite some things that he seems to be saying, Alston's approach has not provided any grounds for supposing that religious experiences (one's own, let alone those of others) are in any way likely to be true. His basic mistake, I suggest, is to suppose that there is one true theory of 'justification'—reliabilism. I have argued in *Epistemic Justification* that there are many different theories of 'justification'; but that there are correct *a priori* criteria for what makes what probable captured by an objective internalist theory, and that these include the Principle of Credulity, which allows us to reach a conclusion (in the light of background evidence) about how probable it is that religious experiences are veridical.

The Balance of Probability

In previous chapters I have urged that various occurrent phenomena are such that they are more to be expected, more probable, if there is a God than if there is not. The existence of the universe, its conformity to order, the existence of animals and humans with moral awareness, humans having great opportunities for cooperation in acquiring knowledge and moulding the universe, the pattern of history and the existence of some evidence of miracles, and finally the occurrence of religious experiences, are all such as we have reason to expect if there is a God, and less reason to expect otherwise. For each of these phenomena $e_n P(e_n|h \& k) > P(e_n|k)$, where h is the hypothesis of theism, k are the phenomena previously taken into account (i.e. tautological evidence where e_n is the existence of the universe; the existence of the universe where e_n is its conformity to order, and so on). Hence, by principles of probability, which I discussed in Chapter 3, for each $e_n P(h|e_n \& k) > P(h|k)$ and so each argument from e_n to h was a good C -inductive argument for the existence of God. I also argued that one phenomenon—the existence of morality—that has been considered to be confirming evidence of the existence of God is not such evidence. I discussed in Chapter 11 the main argument against the existence of God from the existence of evil, and also the argument from the hiddenness of God. I argued there that the existence of the amount and kind of evil that there is in the world (evil of kinds that would be evil whether or not there was a God) were such that a perfectly good God would allow it to occur only if he also provided compensatory life and death, and (perhaps) became incarnate to share our suffering. The fact that the evil (e_n) required additional hypotheses to be added to the hypothesis of theism (h) to save it from disconfirmation meant that the evil

lowered the probability of theism as such (bare theism) from its probability on the evidence taken into account previously (k)— $P(h|e_n \& k) < P(h|k)$. The fact of divine hiddenness did not, however, count against the existence of God.

I believe that I have stated in outline what is the main evidence for and against the existence of God. However, only for the evident general public evidence have I been able to analyse the evidence in detail and assess its force. With respect to the important evidence from phenomena of the kind considered in Chapter 12, I was able to consider only what would be the evidential force of various phenomena if they occurred, not whether in fact the detailed historical evidence (of witnesses, etc.) shows that they did occur. Crucial among these phenomena whose evidential force I had no space to consider was the evidence about the central miracle of Christianity, the Resurrection of Jesus Christ. So I could not reach a final conclusion on just how strong a C-inductive argument was to be had here. I also noted in Chapter 13 that within limits the exact force of an argument from religious experience would depend on whether the subject had himself had religious experiences and on their strength. Nevertheless, I believe that I have explored with some rigour the evidential force of the most evident relevant phenomena on the question of the existence of God, and it is now time to draw my threads together to reach a conclusion. The crucial question remaining for discussion is just how probable all the evidence that I have considered makes the hypothesis of theism. Where all the relevant factual evidence is included in e , h is the hypothesis of theism, and k is mere tautological evidence, what is the value of $P(h|e \& k)$? We may not be able to give it an exact numerical value, but the important issue is whether $P(h|e \& k) > P(\sim h|e \& k)$ and so $> 1/2$. Do we have a good P-inductive argument to the existence of God?

It seems fairly evident to me that an argument from the occurrence of all the phenomena that I have described taken together to the existence of God is not a good deductive argument, for the same reason as the reasons that I gave for a cosmological argument not being a good deductive argument (pp. 136–7) and for a teleological argument not being a good deductive argument (p. 155). The reason is simply that the description of a world in which all the phenomena described occur but there is no God seems, with apparent obviousness, to be a coherent description, to contain no buried

self-contradiction. And in that case there is no valid deductive argument from the occurrence of these phenomena to the existence of God. True, what seems to be coherent may in fact not be. But the description of such a world seems to be a coherent description. Attempts to discover self-contradictions within it notoriously fail, and it seems easy enough to spell out one or more ways in which our world could be such a world. So, assuming that there is no good deductive argument to the existence of God, I return to the question of just how probable all the evidence taken together makes the hypothesis of theism.

I have taken for granted our ordinary criteria for what confirms what (that is, increases its probability), and for what makes what probable overall, criteria that seem to us to be intuitively right; and I have been concerned to apply these criteria to investigate the probability of theism. I have derived our ordinary criteria by meditating upon what we think it right to say about cases in science, history, or other ordinary areas of discourse. For example, I urged that a theory is more likely to be true in so far as it is simple. I reached this conclusion by pointing out from scientific and other examples that we regard a theory h_1 as more probable than another theory of equal scope (i.e. giving us a similar amount of information about the world) h_2 where both are equally successful in leading us to expect the phenomena that we observe ($P(e|h_1 \& k) = P(e|h_2 \& k)$) but where h_1 is simpler than h_2 (for example, in the number of entities or the mathematical relations between them that it postulates). A study of what seems intuitively the right thing to say in many cases when evidence confirms this theory, or confirms this theory more than that one, enables us to extrapolate criteria that we can apply to the issue of theism. It was that which enabled me to conclude that this and that confirmed theism. However, when we come to judgements about whether a hypothesis is more probable than not, there is a shortage of suitable examples from science and other areas from which to extrapolate criteria for when evidence makes a theory more probable than not.

There are certainly cases where it is intuitively obvious (and almost all would agree) that evidence makes a hypothesis overwhelmingly probable or overwhelmingly improbable, at any rate where that hypothesis is a hypothesis about a particular past or future occurrence. On the evidence available to us, it is overwhelmingly probable that the sun will rise again within the next twenty-four hours, that

my study will not suddenly disintegrate, that there really was a Roman Empire, and so on. Juries find prisoners guilty, when the evidence points to their guilt 'beyond reasonable doubt'. Similarly there are cases where hypotheses are wildly improbable. But, when we get anywhere near the border between the probable and the improbable, there is a shortage of examples: examples of hypotheses that, it is intuitively obvious, are more probable than not, without being overwhelmingly probable; or less probable than not without being overwhelmingly improbable. Once there is genuine doubt about, say, some historical theory, historians will seldom agree about whether the theory is more probable than not. And when there is genuine doubt about the correctness of the predictions of a scientific theory, scientists will seldom agree about whether those predictions are more probable than not. The situation is even worse when we are dealing, not with hypotheses about particular past or future events, but with universal scientific theories—that is, theories making claims about all regions of space and time, including ones infinitely distant. Is Quantum Theory more probable than not? Or is the General Theory of Relativity? The answer is in no way clear.

Some philosophers have claimed that all universal scientific theories have zero probability on any evidence. There is, however, no compelling argument for this claim and it seems to run counter to what scientists and others naturally wish to say about scientific theories. Most scientists of the eighteenth century would have said that on the evidence then available Newton's theory of gravitation was overwhelmingly probable. Today perhaps the majority of scientists would say that they think that on present evidence Quantum Theory is probable. Further, almost all scientists are prepared to say that, among theories compatible with the evidence, some are more probable than others; and often that one such is more probable than any other. Yet clearly one theory can be more probable than a second only if the first does not have zero probability. Almost all scientists affirm that they believe that Quantum Theory is true, that the evidence points to its truth rather than its falsity. Yet you can believe something only if you also believe that that thing is not totally improbable (that is, if you believe that its probability is not zero); and you can hold that the evidence points to the truth rather than to the falsity of a theory only if you hold that on the evidence the theory is more probable than not. Talk of scientific theories being probable

or improbable does seem to be embedded in or implied by the way in which we naturally talk about these matters.¹

Mary Hesse has argued that, although scientists do talk about scientific theories being probable, such talk ought not to be taken literally. Talk about a scientific theory being probable ought rather to be construed as talk about the theory probably operating in our spatio-temporal region. However, Mary Hesse's discussion of positive arguments for her thesis is very brief and seems mainly to rely on an appeal to the claim that 'it is not reasonable to suppose that any lawlike generalization . . . in current or any future science will remain forever unqualifiedly true in every instance.'² But why is it not reasonable? No adequate answer is given. Science has been a serious pursuit for only a few hundred years, and we may not yet have found the true laws of nature that hold over endless space and time. But why is it unreasonable to suppose that in a few million years we may find them, and that they or some of them may be probable on the evidence available from the finite spatio-temporal region then accessible. I do not know of good positive arguments why we should understand our talk about theories being probable in the way that Hesse suggests. Even if a reader does accept Mary Hesse's view, that provides no ground for taking a similar view about the

¹ The claim that all universal scientific theories have zero probability on any evidence is in no way forced upon us if we suppose as I suppose that the axioms of the probability calculus and so Bayes's theorem have application to talk about the probability of scientific theories. True, if one assumed that each of an infinite number of incompatible universal scientific theories had equal prior probability, one would have to conclude that the prior probability of each was zero, or infinitesimal. This is because (by the Principle of Countable Additivity) the probability of a number of exclusive (i.e. mutually incompatible) and exhaustive (i.e. one of them must be true) alternatives must add up to 1. And, if an infinite number of such theories each has the same probability, that must be less than any finite number—that means 0; or, if we allow infinitesimals, an infinitesimal value. (On this consequence of the principle of countable additivity, and whether we should say that the probability of each of the disjuncts is 0 or infinitesimal, see my *Epistemic Justification* (Clarendon Press, 2001), Additional Note G.) From that it follows, given the calculus, that the posterior probability of each theory, on any evidence at all, would also be 0 or infinitesimal (except in cases where, if we allow infinitesimals, the prior probability of the evidence $P(e|k)$ is itself infinitesimal). But there is no need to make the implausible assumption that each of such theories has the same prior probability. If simplicity is evidence of truth as I have claimed in this book, the simpler a theory the greater its prior probability; and hence prior probabilities differ as simplicity differs. And, an infinite number of prior probabilities, each of which has a finite value and which are not all equal to each other, can add up to 1. For example, the infinite series $1/2 + 1/4 + 1/8 + 1/16 \dots$ etc. adds up to 1.

² Mary Hesse, *The Structure of Scientific Inference* (MacMillan, 1974), 182.

probability of theism. For her claim amounts to the claim that for every scientific theory T that we discover to hold within a limited spatio-temporal region S , there is always a more fundamental theory T' that holds in a wider region and explains the operation of T within the narrow region. But, if theism explains phenomena in one region, it explains phenomena in all regions (if there is an eternal, omnipotent being at one place and time, there will be one at all places and times); and, if it is true, nothing further explains why it is true (see p. 108). And so, if it probably holds in our spatio-temporal region, it probably holds universally.

The point, however, remains that, given the difficulty in reaching a conclusion about whether any scientific theory is more probable than not, any difficulty about reaching a conclusion about whether theism is more probable than not is not in any way to the special discredit of theism. Yet the situation is by no means hopeless; and so let us proceed, conscious, however, of the considerable difficulty of making judgements of this kind. The reader will recall that, by Bayes's theorem, the probability of a hypothesis h on empirical evidence e and background knowledge k is a function of its explanatory power and its prior probability.

$$P(h|e \& k) = \frac{P(e|h \& k)}{P(e|k)} P(h|k).$$

The explanatory power of a hypothesis is $P(e|h \& k)$, which I term its predictive power, divided by $P(e|k)$, the prior probability of the evidence. $P(e|h \& k)$ is a matter of how likely it is if h is true (and k holds) that e will occur. $P(e|k)$ is a matter of how likely e is to occur at all, whether or not h is true. $P(h|k)$ is the prior probability of h , how likely h is to be true *a priori*—that is whether or not e holds.

We now take h as the hypothesis of theism, 'God exists'. Let k be mere tautological knowledge, and so $P(h|k)$ be the intrinsic probability of theism. We saw in Chapter 3 that prior probability depends on simplicity, fit with contingent background knowledge, and scope. Where the prior probability is intrinsic probability, the second factor does not play any role—for k does not include any contingent background knowledge; there are not any accepted theories for neighbouring fields with which h ought to fit. We saw also, by the example of Newton's theory, that, where we are dealing with a theory of large scope, scope is of far less importance than simplicity in

determining prior probability. The intrinsic probability of theism seems to depend mainly on just how simple a theory theism is.

I argued in Chapters 4 and 5 that theism is an extremely simple theory. It postulates that all explanation is reducible to personal explanation, explanation in terms of the free intentional agency of God or of some finite agent permitted by God to exercise such agency. According to theism, explanations of the other kind—that is, scientific explanations—are reducible to personal explanations in the sense that the operation of the factors cited in the scientific explanations is to be explained in personal terms. Newton's laws work because God keeps them in operation. There are planets because God is responsible for the operation of the law that brought about their evolution from pre-existing matter, and so on. Theism is simple in postulating that in this way complete explanation is all of one kind.

Further, as the cause of the existence and powers of all other things, theism postulates the existence of just one person, the simplest kind of person that there could be. A person is a being with powers, beliefs, and intentions. The theist postulates that God has powers as great as they logically can be. He is infinitely powerful—that is, omnipotent. The power of a person is the power to do some action intentionally. Such a person is omnipotent if and only if whatever (logically possible) action he intends to do, he succeeds in doing. That there is an omnipotent God is a simpler hypothesis than the hypothesis that there is a God who has such-and-such limited power (for example, the power to rearrange matter, but not the power to create it). It is simpler in just the same way that the hypothesis that some particle has zero mass or infinite velocity is simpler than the hypothesis that it has a mass of 0.34127 of some unit, or a velocity of 301,000 km/sec. A finite limitation cries out for an explanation of why there is just that particular limit, in a way that limitlessness does not. There is a simplicity about zero and infinity that particular finite numbers lack, a simplicity recognized by scientists as evidence of truth in the judgements that they make about the relative probability of scientific theories.

According to the theist, God's beliefs have a similar infinite quality. Human persons have some few finite beliefs, some true, some false, some justified, some not. In so far as they are true and justified (or at any rate justified in a certain way), beliefs amount to knowledge. It would seem most consonant with his omnipotence that an

omnipotent being have beliefs that amount to knowledge. For without true beliefs, about the consequences of your actions, you may fail to realize your intentions. True beliefs fail to amount to knowledge only if they are true by accident. But, if (see below) the divine properties are possessed necessarily, God's beliefs could not be false, and so could not be true by accident. And, if an omnipotent being has knowledge, the simplest supposition is to postulate that the omnipotent being is limited in his knowledge, as in his power, only by logic. In that case he would have all the knowledge that it is coherent to suppose that a person could have—that is, he would be omniscient.

A person could be omnipotent in the sense that, whatever (logically possible) action he formed the intention to do, he would succeed in doing, and also omniscient so that he knew what were all the (logically possible) actions available to an omnipotent being in his situation, and yet be predetermined to form certain intentions. His intentions might be determined by factors outside his control, or at any rate, as are those of humans, greatly influenced by them. But, if a person is predetermined (or has an inbuilt probabilistic tendency) to act in a particular way, this means that a tendency to act in a particular way is built into him; he does not choose how to act solely on the basis of the goodness of the possible actions he could do. A person with an inbuilt detailed specification of how to act is a much more complex person than a perfectly free one. And theism makes the simplest supposition, that God is not merely omnipotent, and omniscient (and so knows what are all the possible actions) but also perfectly free. He chooses between actions solely in virtue of *a priori* considerations about whether they are good actions, and how good they are. It is clearly simpler to suppose that the ultimate principle of explanation, the final source of things, has always been the same—rather than to suppose that only, for example, in 4004 BC did God come to be and reign—and so to suppose that God has existed eternally. Theism thus postulates one person of an incredibly simple kind—one with such powers, beliefs, and intentions that there are no limits (apart from those of logic) to his powers, to the extent of his justified true belief, and to his choice of intention; and no limits of time to his existence. It is simpler to suppose that these properties of eternal omnipotence, omniscience and perfect freedom belong essentially together—for that removes the need to explain why God exists now in terms of his having existed a long time ago and subsequently choosing to retain his omnipotence or whatever.

I argued in Chapter 5 that the other defining properties of God, and especially the property of perfect goodness, are possessed necessarily by a being who possesses the ones that I have just considered. I conclude that the hypothesis of theism is a very simple hypothesis.

It remains to me, as to so many who have thought about the matter, a source of extreme puzzlement that there should exist anything at all. And maybe for that reason $P(h|k)$ is low. But there does exist something. And, if there is to exist something, it seems impossible to conceive of anything simpler (and therefore *a priori* more probable) than the existence of God. The simplest kind of scientific explanation of our data would be, I argued in Chapters 7 and 8, that everything began from an unextended physical point. Much less simple would be hypotheses that everything began from an extended volume of matter-energy, or that there always had been an extended volume of matter-energy. Let me allow for the sake of argument (despite my doubts on this) that the unextended point hypothesis is as simple as theism; and that one or two of the other scientific hypotheses are not too much less simple than theism. The issue between them and theism as complete (or ultimate) explanations of phenomena will now turn on their relative explanatory power. This is a matter of how much more probable it is that the evidence would occur given the hypothesis than it would be otherwise.

Our evidence e are the phenomena that I have described in Chapters 7–13; but I shall find it convenient to exclude for the moment the evidence of religious experience that I described in Chapter 13, and confine myself to the evidence of the existence and general character of the world described in Chapters 7–11 together with the evidence of the occurrence of certain undoubted historical events, together with some evidence confirming the occurrence of violations of natural laws described in Chapter 12. Let me summarize that evidence. There is a universe, in which, throughout the vast spatio-temporal region of which we have knowledge, physical bodies have identical powers. They attract or repel each other in identical ways on earth and on Mars, and, as far as we can tell, on the most distant galaxy, this year and last year and millions of years ago. The powers of things are identical. The universe is a beautiful universe. In it there are many conscious beings, and among them humans, agents of limited power and knowledge and, as far as we can tell, limited free will in the sense in which I have defined this. They have the power to

grow in these qualities or to abandon them. They are capable of marvelling at the order of nature and of worshipping God. Humans are subject to desires, including biologically useful ones; these desires are sometimes bad ones, which means that humans are subject to temptation to do what is morally wrong (or otherwise bad). They are, therefore, able to choose for themselves whether to do the morally right action, and may in the course of time develop a morally good character. Humans are greatly interdependent, capable of increasing each other's power, knowledge, and freedom, making each other happy or unhappy, and in this way, being subject to birth and death, capable of influencing distant generations. Humans are thus, in my terminology, apparently humanly free agents. The world is thus a providential world both in the respect that humans are often able to work to satisfy their bodily needs, and more importantly in the respect that provision is made for their deeper needs. The world contains much evil, but the evil is not endless and it is either evil brought about by humans, or evil of a kind that is necessary if humans (and animals) are to have the various benefits described in Chapters 10 and 11. All of this forms the general character of the world in which we live—the human condition. The world is a do-it-yourself world, one that humans can mould and in which they can make themselves, and make themselves fit for a different kind of world if such there be. Within our world there are further relevant particular phenomena—there is the work of prophets encouraging humans to the worship of God, and to knowledge of God and the universe, and to the service of their fellows, encouraging them towards an 'upwards march' especially when selfishness dominates. And there is too some slight evidence of violations of natural laws from time to time in religious contexts for good and religious ends. With this *e*, what is the predictive power of theism? How likely is it if theism is true that things will be thus?

Now the existence of God does not entail that there will be a universe of the kind described. I argued in Chapter 6 that, whether or not his goodness required him to make a universe of some kind, his perfect goodness certainly did not make it inevitable that he would make a universe of this kind. God could not make a best of all possible worlds, for there could not be such a world. But perhaps he could make a world of a best kind. Any such world would contain all the kinds of goodness; and other worlds would be better than it only in containing more instances of each of these kinds or more

intense degrees of these kinds. Now, in having freedom to choose between good and evil, humanly free agents have a unique kind of goodness that God himself does not have. All the other kinds of goodness in the world (consciousness, power, beauty, rationality, and so on) are possessed by God himself; or, if some of these kinds (love and cooperation) require other beings with whom God interacts, they could be instantiated in the world by God creating beings quite different from ourselves. But, of course, to give humanly free agents this freedom to choose between good and evil is to risk the evil that would result from them taking the wrong choices. In consequence of this disadvantage I suggested that to bring about a world in which, as well as other things, there were humanly free agents would be an equal best kind of act, equal best with the action of bringing about a world of an otherwise best kind without such agents. But, if there is to be a world containing such agents, they must be embodied, and the universe must be governed by simple natural laws. And one would expect all the other general features of our universe that I have just described. The orderliness of the universe is likely to lead to it being a universe beautiful in the ways in which the physical objects that it contains behave. God will be expected to bring about a beautiful universe, and in this respect and also because of the plants that it contains, our universe is a beautiful universe. If God creates such free moral agents, they will be beings capable of knowing him and God will to some extent (without unduly limiting their freedom) seek to make himself known to them; and, if they make many bad choices and become corrupt, God will need to help them to recover their route to goodness. Hence we might expect the pattern of history and some evidence of violations of natural laws of the kind described in Chapter 12. Although much evil is necessary for the attainment of many of these purposes, there is a limit to the amount of evil that God ought to allow humans (and animals) to suffer for the sake of the good that it makes possible. Even so, if he allows us to suffer as much as we do, he would need to provide a compensatory period of afterlife for any who suffer too much and perhaps to become incarnate to share our suffering. To add to the hypothesis of theism that he does these things complicates it, but not—I claimed in Chapter 11—very much. In the absence of this extra evil, the probability that God would create a universe such as I have described would be, I claimed, an equal best kind of act, and so there would be a probability of $1/2$ that he would do so. But, if we need to complicate

theism somewhat in order to account for the amount of evil, we must put the probability that God would bring about our kind of universe a bit lower. With e as all the evidence listed so far, k as tautological background evidence, and h as the hypothesis of theism, let's say $P(e|h \& k) = \frac{1}{3}$. More loosely and accurately, my point is simply that our universe is of a kind such that, given God's character, he might well choose to create. (I argued that, although the existence of animals is—despite the evil necessary to make possible some of their good states—a good state that God has reason to bring about, his doing so would not be, or be part of, a best or equal best kind of act, and so we cannot give such a high probability to God creating animals.)

What finally (with e and k as above) of the intrinsic probability of e , the probability that e would occur whether or not theism is true? It follows from the axioms of the probability calculus that the probability of any proposition on any other proposition is the sum of the probabilities of each of the different ways in which the first proposition could be true on the evidence of the second proposition—for example, $P(p|r) = P(p \& q|r) + P(p \& \sim q|r)$. Examples make this intuitively obvious. The probability on some evidence that Jones did the crime is the sum of the probability that Jones did the crime together with Smith, and the probability that Jones did the crime but not together with Smith. So, for our stated h , e , and k , $P(e|k)$ is the sum of the intrinsic probability that e holds and there is a God, and the intrinsic probability that e holds and there is no God. $P(e|k) = P(e \& h|k) + P(e \& \sim h|k)$. Now $P(e \& h|k)$ is the product of the two factors whose value we have already considered, $P(e \& h|k) = P(h|k)P(e|h \& k)$. This simply repeats the top line of the right-hand side of Bayes's theorem. So everything depends on the other component of $P(e|k)$, $P(e \& \sim h|k)$ which equals $P(e|\sim h \& k) P(\sim h|k)$. So we may rewrite Bayes's theorem as:

$$P(h|e \& k) = \frac{P(e|h \& k)P(h|k)}{P(e|h \& k)P(h|k) + P(e|\sim h \& k)P(\sim h|k)}.$$

$P(e|\sim h \& k)P(\sim h|k)$, the probability that there would be a universe of our kind and no God, is the sum of the probabilities of the various ways in which this could come about. Let h_1 be the hypothesis that there are many gods or limited gods; h_2 be the hypothesis that there is no God or gods but an initial (or everlasting) physical

state of the universe, different from the present state but of such a kind as to bring about the present state; and h_3 the hypothesis that there is no explanation at all (the universe just is and always has been as it is). Then

$$P(e|\sim h \& k)P(\sim h|k) = P(e|h_1 \& k)P(h_1|k) + P(e|h_2 \& k)P(h_2|k) \\ + P(e|h_3 \& k)P(h_3|k).$$

Now the hypothesis of theism differs from other hypotheses that purport to give a complete (or ultimate) explanation of our data in that we do not have to build into God a propensity to bring about a universe of a certain kind rather than any other kind. There is a significant probability that he will do so merely in the virtue of his essential properties, and especially his perfect goodness. If some lesser god is to have the same propensity to bring about this universe as a God would have, we have to suppose him to have a certain character that would not be entailed, as God's is, by the extent of his knowledge and freedom. And, quite apart from that point, I have argued that the hypothesis of theism is a very simple hypothesis indeed, simpler than hypotheses of many or limited gods. I argued in Chapters 3 and 7 that our normal criteria of probability give a very simple hypothesis greater intrinsic probability than a disjunction of many alternative more complex hypotheses. In that case, theism is going to be more probable than h_1 , the disjunction of hypotheses of many or limited gods; and there is much less reason why they should bring about a universe at all or one of our character—they may not be able to do so, and not being perfectly good may not have much reason to do so (unless we complicate these hypotheses further by building into them the requisite propensity). $P(e|h \& k)P(h|k) \gg P(e|h_1 \& k)P(h_1|k)$.

Let us turn now to h_2 . This is the hypothesis that there is no god or gods, but an initial or everlasting physical state of the universe, different from the present state but such as to bring about the present state. But there is no particular reason why an unextended physical point or any of the other possible starting points of the universe, or an everlasting extended universe, should as such have the power and liability to bring about all the features that I have described. Indeed I have shown chapter by chapter that, for each such feature, it is rather improbable that any physical mechanisms at all (described in

fairly general terms) would bring about those features rather than any other ones. It will only become at all probable that there will be a universe of our kind if we build into the hypotheses an enormous amount of complexity. And then by the above principle a disjunction of such very complex hypotheses will be intrinsically less probable than one very simple hypothesis. Either $P(e|h_2 \& k)$ is going to be much lower than $P(e|h \& k)$ or $P(h_2|k)$ is going to be much lower than $P(h|k)$.

And that our universe should have all the characteristics described (above all, the overwhelming fact that each particle of matter throughout vast volumes of space should behave in exactly the same way as every other particle codified in 'laws of nature') without there being some explanation of this is beyond belief. While $P(e|h_3 \& k) = 1$ (the universe being this way unexplained entails it being this way), $P(h_3|k)$ is infinitesimally low.

I conclude that

$$P(e|h_1 \& k)P(h_1|k) + P(e|h_2 \& k)P(h_2|k) \\ + P(e|h_3 \& k)P(h_3|k) = P(e|\sim h \& k)P(\sim h|k)$$

will not exceed $P(e|h \& k)P(h|k)$, the top line of the right-hand side of Bayes's theorem. And so, $P(h|e \& k)$, the posterior probability of theism on the evidence considered so far, will not be less than 1/2. I stress again that it is impossible to give anything like exact numerical values to the probabilities involved in these calculations. I have attempted to bring out the force of my arguments by giving some arbitrary values that do, I hope, capture within the roughest of ranges the kinds of probabilities involved. But in reality all that my conclusion so far amounts to is that it is something like as probable as not that theism is true, on the evidence so far considered. However, so far in this chapter I have ignored one crucial piece of evidence, the evidence from religious experience. I concluded the last chapter (p. 326) with the claim that, unless the probability of theism on other evidence is very low, the testimony of many witnesses to experiences apparently of God suffices to make many of those experiences probably veridical. That is, the evidence of religious experience is in that case sufficient to make theism overall probable. The argument of Chapter 13 was that the testimony of many witnesses to experiences apparently of God makes the existence of God probable if it is not

already on other evidence very improbable. I believe that I have shown that that condition is more than adequately satisfied and hence the conclusion of Chapter 13 applies. On our total evidence theism is more probable than not.³ An argument from all the evidence considered in this book to the existence of God is a good P-inductive argument. The experience of so many people in their moments of religious vision corroborates what nature and history show to be quite likely—that there is a God who made and sustains man and the universe.

³ I argue in *The Resurrection of God Incarnate* (Clarendon Press, 2003) that, when we take into account the detailed historical evidence of the life, death, and Resurrection of Jesus, the probability that there is a God becomes very much greater than that.

Additional Note 1

The Trinity

The arguments of this book are purportedly ‘arguments to the existence of one God’, defined as necessary, necessarily eternal, omnipotent, omniscient, and perfectly free (from which the other divine properties follow). I call such a being a divine being or person. There is an argument on p. 98–9 to the effect that there can be only one such being: that, if there was a second divine being, and so the first divine being depended for his existence on a second divine being (who permits him to exist), that second being could not depend for his existence on the first being (for one cannot have causation in a circle), and the first being would not be divine. Duns Scotus gives this argument.¹ If, however, God’s eternity is understood as his everlastingness (existing at every period of time), and literally simultaneous causation is logically impossible (i.e. causation takes time), and so a cause ‘acting at a time’ to bring about an effect is to be understood in the way analysed on p. 77 n. 3, the argument does not work. For a first divine being could cause (actively, or passively—i.e. permit to exist) at each period of everlasting time the existence of a second divine being at a virtually identical but slightly later period of time, who in turn necessarily caused (actively or passively) the existence of the first divine being at a yet slightly later time, and so each could depend everlastingly on the other for its existence. The causation would have to be necessary—a consequence of their natures that each thus caused the existence of the other. The perfect freedom of a divine being is always limited by his perfect goodness, and the perfect goodness of each would ensure that neither destroyed the other.

¹ *Opus Oxoniense*, I, Dist. 11, Q1, 7th way, trans. in *Duns Scotus: Philosophical Writings*, ed. A. Wolter (Hackett, 1987), 89–91.

There is, however, a further objection to the possibility of the existence of two divine beings, also raised by Scotus, that 'if two omnipotent beings exist, each will make the other impotent, not indeed by destroying the other, but because one by his positive will could keep non-existent what the other wills should exist'. Thus one could will that the Earth move clockwise round the sun at the same time as the other wills that it move anti-clockwise round the sun.

This, however, is avoidable if one divine being sees as part of his perfect goodness that he should confine his causation to one sphere of activity, while the other confines his causation to a different sphere. That can be ensured if one of the divine beings actively causes the other to exist at each period of time on that condition, while the second being only passively causes the existence of the first being (that is, does not stop him existing), at each period of time. In virtue of this asymmetry of dependence, the second being recognizes the authority of the first being to delimit their spheres of activity. The existence of more than one divine being is possible given an asymmetry of dependence, and I cannot see how else it is possible. For only thus would each divine being recognize who had the right to define the spheres of activity of each. Clearly this same process could be repeated to allow for more than two divine beings. For a fuller account of this issue, see my *The Christian God* (Clarendon Press, 1994), esp. chs. 7 and 8.

The arguments of the present book are, therefore, more accurately described as arguments to the existence of one divine being, a person on whom all else depends. That is theism, the view common to Christianity, Judaism, and Islam. Theism is, however, compatible with the Christian doctrine of the Trinity that there is more than one divine person, given that any other divine person depends on the first divine person for his existence and for the delimitation of the sphere within which it would be good for him to act. In describing my arguments as arguments to 'the existence of God' I am therefore oversimplifying, for the Christian God consists of Father, Son, and Holy Spirit. In Christian terms my arguments are arguments to the existence of God the Father, though compatible with some of his activity being mediated through God the Son and God the Holy Spirit.

In *The Christian God* I argue for the truth of this Christian doctrine on *a priori* grounds. I argue there that his perfect goodness requires a first divine being to produce a second and in cooperation

with the second to produce a third; but that there is no necessity to produce a fourth. Any fourth divine being would therefore not exist necessarily (in the sense examined on p. 96 of the present book), and so would not be a divine being. So there will be three and only three divine beings. The arguments of the present book do not, however, depend on these latter claims. If I am right in supposing that the existence of one divine being entails the existence of two more, then, in arguing in the present book to the existence of a divine being, my arguments are in effect arguments to the existence of the Christian Trinity. But, if I am mistaken in this supposition, then they are arguments only to one divine being, which Judaism and Islam suppose to be the only divine being. I do, however, oversimplify (on p. 7 and subsequently) in defining God in the Jewish–Islamic way. This can easily be rephrased by the Christian as follows: God is the whole Trinity consisting of three divine persons; the arguments are arguments to the existence of one divine person, on whom the other two necessarily depend.

Additional Note 2

Recent Arguments to Design from Biology

From time to time various writers have claimed that evolution by natural selection of small variations cannot explain the emergence of some characteristic, or more generally of some species. Most recently Michael Behe has claimed that there are in nature ‘irreducibly complex systems, that is, systems composed of several well-matched interacting parts that contribute to the basic function [of the system] and where the removal of any one of the parts causes the system to effectively cease functioning’. In these systems each of the parts by itself conveys no evolutionary advantage on the organism; only if all are present, is there that advantage. Yet it would require a number of mutations of separate genes to occur simultaneously to produce the system, and this is extremely improbable. The occurrence of such systems, Behe argues, shows the inadequacy of scientific explanations of the evolutionary process, and the need to invoke a designer.

Now there is considerable doubt whether Behe has shown that there are such systems. It may well be that in the examples that he cites of such irreducibly complex systems, such as the bacterial flagellum, the separate parts do convey an evolutionary advantage (although not necessarily the same advantage as that of the whole system) on organisms that have them.¹ But even if there are such systems, that fact may make no difference to the cogency of the argument from spatial order (that is, the argument from the

¹ On these detailed biological issues, see the papers by Michael Behe and Kenneth R. Miller in Neil A. Manson (ed.), *God and Design* (Routledge, 2003).

existence of human and animal bodies) set out in Chapter 8; it may not either add to or diminish its strength.

If physical laws apart from those concerned with the human brain (see p. 170) were totally deterministic (contrary to the more usual interpretation of Quantum Theory), then—given the boundary conditions of the universe (for example, the initial conditions at the time of the Big Bang), it would have been inevitable that human and animal bodies would evolve on Earth, and perhaps on other planets too. All that the discovery of irreducibly complex systems would show is that the mechanism through which this was brought about was not always that of natural selection of small variations of evolutionary advantage. It would show that the initial conditions and laws are such that sometimes there are simultaneous mutations that throw up whole complex systems (not merely single mutations causing small differences) that then get selected for the advantage that they give to their possessors in the struggle for survival. But that would make no difference to the degree to which the universe was very fine-tuned, and so it would make it no more or less probable that it was designed by a powerful designer.

If, however, physical laws are merely probabilistic, then it is just possible that the discovery of such systems could be relevant. But for this to be the case two other theses would need to be established. First, we would need to show that—given the actual physical laws, and given the initial conditions of life on Earth (that is, the exact arrangements of fundamental particles in the first organisms to live on Earth), and given all the influences on organisms from their environment, it is improbable that there would ever occur simultaneous mutations throwing up irreducibly complex systems possessing a selective advantage. This seems to me implausible. Of course normally no doubt mutations are not correlated in this way, but it is not improbable that occasionally this would happen. And, secondly, it would need to be shown that, without such correlated mutations, it is improbable that human and animal bodies would evolve. It would need to be shown, not merely that certain species would not evolve, but that (even if the same species did not evolve), the evolutionary process leading to human bodies would not have occurred. That seems to me rather dubious. But in so far as evidence made both of these theses probable, and given that, as I suppose, evidence makes it probable (for reasons of Quantum Theory or other reasons) that

physical laws are indeterministic, then discoveries of ‘irreducibly complex systems’ (if such discoveries were established) would be crucially relevant to the argument from spatial order.

For, first, it would have shown that, barring supernatural intervention, it is not probable that human and animal life would evolve and so that the fine-tuning of the universe does not provide such a strong argument for the existence of God. For the fact that the laws and boundary conditions of the universe do not rule out the evolution of humans and animals is not enough to show that the universe is the kind of universe God would have made in order to produce humans and animals; for, if that was his purpose, he could have made a universe much more conducive to that evolution. And so, secondly, it would have shown that we need discoveries of irreducibly complex systems as evidence of divine intervention in a process that otherwise would be most unlikely to produce intelligent life. Then such discoveries would indeed then provide a good C-inductive argument to the existence of God. But in my view current science seems to suggest that neither of the two theses of the previous paragraph (and especially the first thesis) is very probable.

I should add that Behe, and other writers in the same spirit, claim that their argument shows only design by a designer and not by God. But any phenomenon can be evidence of design only if it is the kind of thing a designer would have sought to bring about. Different designers have different purposes. There could be a designer who had (from our point of view) the strangest purposes—for example, to produce ugly chaotic universes. It is only if we have some *a priori* reason to suppose that certain sorts of designer with certain kinds of purpose are more likely to exist than are others that any argument from design can get going. I have argued that *a priori* the existence of a God with the traditional properties is a simpler hypothesis than is the hypothesis of any other kind of designer, and so that it is more probable *a priori* that there will be such a designer. He will (to simplify the argument of Chapter 6) seek to bring about very good states of affairs. So, given—as I have argued—that the existence of humans is a very good state, we would expect to find such a state if there is a God, and so the existence of such a state is evidence of the existence of God. No phenomenon shows design as such without at the same time showing something about the sort of being who designed it. Behe’s argument must be construed either as an argument to the existence of God, or as an argument to the existence of a

less powerful but basically benevolent designer. But it will show only the latter rather than the former, if *a priori* (because of reasons of simplicity and scope, or for reasons of background knowledge) the existence of a lesser deity (without the existence of God) is more probable than the existence of God. I have argued to the contrary in Chapter 6.

Additional Note 3

Plantinga's Argument against Evolutionary Naturalism

Alvin Plantinga has developed recently an argument that makes some of the same moves as does my exposition of the argument from consciousness.¹ He writes that he understands by naturalism (N) 'the belief that there aren't any supernatural beings—no such person as God, for example'.² He understands by evolution (E) 'the belief that we humans have evolved in conformity with current evolutionary doctrine'—that is, presumably, the Darwinian theory of natural selection of random variations in genes that cause phenotypic characteristics (including behaviour) conducive to the survival of the organism. He understands by the 'reliability of our cognitive processes (R) the claim that most of the beliefs which we acquire about the present, the future, and the past, and the nature of the world are more-or-less true'. He then claims that, given only N and E, the probability of R is 'low or inscrutable',³ and so that an evolutionary

¹ Plantinga's most recent statement of the argument and fullest defence against a large number of objections is contained in the collection J. Beilby (ed.), *Naturalism Defeated?* (Cornell University Press, 2002). My quotations from Plantinga and objections to him are all from this volume.

² *Ibid.* 1–2.

³ In *ibid.* 5 n. 11, Plantinga allows the possibility that we may not be able to calculate this probability (or others that he discusses). Hence the 'or inscrutable'. If he affirms that the probability is inscrutable, he can perhaps say that, given only N and E, we have no reason to believe that R; but not that R is probably false. But the former allows the possibility that we have, beside N and E, plenty of other reason for believing R. Plantinga's argument has significant force only against naturalism, and so against the existence of God, if we take him to be affirming that the probability of R, given N and E, is low (and that other probabilities that have consequences for this probability have values such that they

naturalist cannot rationally believe that her cognitive faculties are reliable, and so cannot rationally believe any belief that they produce including the beliefs N and E. (N & E) is a self-defeating belief; no one can rationally believe that N and E are both true. So, although Plantinga does not put it in this way (presumably because of his doubts about E), given R, E makes it probable that N is false. If we regard E as established, there is a correct P-inductive argument from E and R to the falsity of N, and so to the existence of God (or at least some lesser god), who created our cognitive faculties and made them reliable. As originally presented, the argument had four steps, but it can easily be reduced to two.⁴ First, Plantinga argues, evolutionary naturalism as such gives no reason to suppose that we would have beliefs whose content guides our 'desires' to cause our behaviour. (Plantinga means by 'desire' what I mean by intention or purpose. I shall follow his terminology in this Additional Note.) That is, that we do what we do in part because of what we believe. For example, when we see a tiger and run away, we run away because we believe that the tiger is dangerous and we 'desire' not to be killed. It is far more probable, given evolutionary naturalism, that our behaviour would be caused by brain events themselves uncaused by beliefs, and that any beliefs would be mere epiphenomena. So there would be no selective advantage in having true beliefs rather than false beliefs; and, since there are far more possible false beliefs than true ones, it is improbable that, whatever processes throw up beliefs would throw up largely true ones. This stage of the argument seems to me correct for reasons that are in essence the same as ones that I gave in Chapter 9. *A priori* it is improbable (barring theism) that brain events would produce beliefs (about our environment) at all—because of the complexity of the laws that are inevitably involved therein. *A priori* and for the same reason, it is improbable (barring theism) that (together with purposes) beliefs would affect brain events and thereby behaviour. A naturalistic theory of evolution that rules out theism makes it very improbable that there would be beliefs correlated in these ways with brain events.

The second stage of Plantinga's argument is to claim that, even if beliefs (with 'desires') did cause behaviour, and so organisms that

have the consequence that this probability is low). For this reason I shall in future ignore the 'or inscrutable' clause.

⁴ And was so reduced by Plantinga (*ibid.* 9–10).

had beliefs (with 'desires') that caused behaviour conducive to their survival would have a selective advantage and so tend to survive, there is no reason (in the absence of God or gods) to suppose that this would give a selective advantage to organisms having true beliefs. For false beliefs could be combined with 'desires' in such a way as to produce behaviour that gave the same selective advantage. Plantinga gives two very different kinds of example to show how this could occur. In the first kind of example, belief-forming processes cause a false belief instead of the true belief that (we believe) they normally cause, and so, if we had had our normal desire, we would behave in a way not conducive to survival; but a desire-forming process at the same time forms an unusual desire that (together with the false belief) leads to the same behaviour as does (what we believe to be) the true-belief-normal-desire pair.

Paul sees a tiger and moves away, but this need not be caused by his belief that here is a dangerous tiger and his desire to stay alive. Perhaps, writes Plantinga, instead Paul 'thinks the tiger is a large, friendly cuddly pussycat and wants to pet it; but he also believes that the best way to pet it is to run away from it . . . or perhaps he thinks he is about to take part in a sixteen-hundred metre race, wants to win and believes the appearance of the tiger is the starting signal . . .'.⁵ And so on. All this is logically possible. But it would require a very complicated mechanism of belief and desire production by our brains to bring it about. By far the simplest mechanisms (unsimple though they are for reasons given in Chapter 9) for producing beliefs and desires will be two separate mechanisms, one of which produces beliefs and the other of which produces desires. So, for given beliefs different desires would lead to different behaviour; and, conversely, for given desires different beliefs would lead to different behaviour. Plantinga-type scenarios are not compatible with this. For, if Paul ceases to desire to pet the animal in front of him and holds the beliefs that Plantinga postulates, he will be eliminated. His beliefs would have to change at the same time (for example, he would have to cease to believe that here is a large, friendly, cuddly pussycat) if he is to survive. And similarly if he ceases to want to win the race. The mechanisms that cause beliefs would have to cause totally different beliefs as our desires change. But, given (largely) independent mechanisms causing beliefs and desires and given that we have a desire to

⁵ *Naturalism Defeated?*. 8. Plantinga here quotes from an earlier work of his.

survive (as the evolutionary naturalist supposes that normally we do), we are by a long way more likely to survive if we have true beliefs about how to fulfil our other desires compatibly with surviving than if we have false ones. It is much more probable that, if biochemical processes cause beliefs and desires and these cause behaviour, those processes would throw up simple mechanisms than that they would throw up very complicated ones in which beliefs vary with desires in such a way as to cause behaviour that mimics that caused by simple mechanisms. Hence, given mental–physical interaction, it is—given (N & E)—probable that our belief-forming processes will be reliable. I argue in Chapter 9 that it is most improbable that our belief forming processes will be reliable unless we use the correct inductive criteria of what is evidence for what.

Plantinga's second kind of example to show that false beliefs could be linked with desires in such a way as to produce behaviour with a selective advantage is of cases where beliefs about mundane matters expressed in terms of a false metaphysical theory have the same consequence for behaviour as those expressed in terms of a true metaphysical theory: 'Perhaps Paul is a sort of early Leibnizian and thinks everything is conscious (and suppose that is false); furthermore, his ways of referring to things all involve definite descriptions that entail consciousness, so that all of his beliefs are of the form "that so-and-so conscious being is such-and-such". Perhaps he is an animist and thinks everything is alive.'⁶

Now of course this sort of thing can and does happen. But there is no sharp dividing line between beliefs about the mundane and beliefs about the metaphysical. The latter are merely beliefs at the end of a spectrum of beliefs about larger and deeper matters. There is, we have seen, a selective advantage in having mechanisms that make inferences on mundane matters in accord with correct criteria of what is evidence for what. The higher animals, as well as humans, can make the predictions licensed by simple generalizations from observable data of particular interest to them. If food has been provided each morning in this place for many days past, they expect to find it here today also. But humans are capable of more sophisticated reasoning; and they have a selective advantage in being able to add to their stock of true beliefs by reflection and experiment. The continuity of subject matter between the mundane and the metaphysical will lead us to use

⁶ Ibid. 9.

the same criteria in metaphysics. As I illustrate in Chapter 3, beliefs about the unobservable are judged correct by the same criteria as beliefs about the observable; and fundamental laws and ultimate explanations are judged correct by the same criteria as explanations of some event in terms of some immediately precedent cause. So it is to be expected (given N & E) that we will seek to improve our metaphysical beliefs by correct criteria—as the present book is trying to do; and so that, in the course of time, they will indeed improve. Natural selection encourages the emergence of true metaphysical beliefs—not because they themselves provide any selective advantage, but because they are produced by mechanisms of the same kind as produce true mundane beliefs. Yet, the more metaphysical the belief, the more our application of the criteria needs to be conscious, careful, explicit, and sophisticated. Simple generalization often requires little thought. And humans who form false metaphysical beliefs are not likely in consequence to suffer quick elimination. On the contrary, having false metaphysical beliefs can sometimes confer a short-term survival advantage. If in Roman times you believed that confessing the truth was unimportant, or alternatively believed that Caesar was a god, you would offer incense to Caesar. Thereby you would be more likely to survive than those who did not hold such false beliefs. So false metaphysical theories are likely to flourish for far longer than false simple mundane generalizations, as Plantinga's example illustrates. But humans have the correct inductive criteria to weed them out, and, when combined with a desire in general to hold true beliefs—without which humans would not survive for long—humans will have a tendency in the course of time to acquire true metaphysical beliefs.

I conclude that Plantinga's argument against evolutionary naturalism provides no good additional reason for theism beyond that provided by my argument from consciousness.

Concordance

	Revised edition	Second edition
Introduction	[1]	1
1. Inductive Arguments	[5]	4
2. The Nature of Explanation	[22]	23
3. The Justification of Explanation	[51]	52
4. Complete Explanation	[70]	73
5. The Intrinsic Probability of Theism	[90]	93
6. The Explanatory Power of Theism: General Considerations	[107]	110
7. The Cosmological Argument	[116]	133
8. Teleological Arguments	[133]	153
9. Arguments from Consciousness and Morality	[152]	192
10. The Argument from Providence	[180]	219
11. The Problem of Evil	[200]	236
12. Arguments from History and Miracles	[225]	273
13. The Argument from Religious Experience	[244]	293
14. The Balance of Probability	[277]	328
Appendix A	[293]	various places
Appendix B	[300]	172–88
Index	[323]	357

This page intentionally left blank

Index

- Abraham 273
actions, basic 35–6
 best kind of 102
 done for reason 100–5
 equal best kind of 102–3
 intentional 35–45
 mediated 35–6
Adams, R. M. 155 n.
agnosticism, honest 267–72
Alston, W. P. 327 n.
altruism 216
ancestors, responsible for our
 evils 239
angels 119 n.8, 130, 239–40,
 262
animals 118, 120, 131, 188–9,
 200–1, 231–3, 236, 243–5,
 250–2, 255–6, 262,
 337–8
animate being, defined 117
Anselm, St 8 n.
Aquinas, St. Thomas 9 n.5, 90–2,
 100, 114 n.3, 117, 135–6,
 154–5, 282 n.
argument from design, defined 153
arguments, a posteriori 8–12
 a priori *see* ontological arguments
 cumulative 12–14, 17–19
Aristotle 100, 116 n., 136
Armstrong, D. M. 32 n.7
atomic theory of chemistry 84–5
atonement 289
Augustine, St. 285
Babylonian astronomy 317
background evidence,
 knowledge 16, 60–1
 irrelevant to theism 65, 145–6
Barnes, J 8 n.
Barrow, J. D. 173 nn.10–11, 174 n.,
 176 n.19, nn.21–2, 179 n.,
 180 n.
Bayes's Theorem 66–72, 332 n.1,
 339–41
Beardsworth, T. 302 n.8
beauty of universe 121, 188–91
Behe, M. 346–9
being-of-use 223, 259–61
beliefs, basic, defined 253
 true, selective advantage
 of 352–4
Bennett, J. 148 n. 19
best of possible worlds 114–15
Big Bang, *see* universe, beginning of
'big' phenomena 74–5, 142–3, 160,
 235, 277, 326
biologically useful sensations 221,
 237
bodies, nature of 50–51, 169–70
 need for 123–31, 182–4
boundary conditions of
 beginningless
 universe 179–80
brain transplants 197–200
Broad, C. D. 304 n.
Brown, P. 90 n., 91 nn. 10–11
Buddhism 295 n.2

- Carroll, J. W. 32 n.6
- C-inductive arguments,
defined 6–7
- carbon-based life 173–6
- causal powers, *see* laws of nature,
S–P–L account of
- causes, accidentally ordered 90–2
essentially ordered 90–2
full, defined 24
of collections of states 141–2
partial, defined 24
see also explanation
- character formation 222, 225
- charity, principle of 61–2
- Chisholm, R. M. 294–5, 307–10
- Christ, Jesus 75, 239 n. 2, 266 nn
8–9, 273–4, 290–2, 299,
324 n., 329
- Christianity 7, 286, 288–91, 295 n.
2, 344–5
- Church, Christian 274
- Clarke, S. 136
- coherence, internal, of theism 1
- Collins, R. 173 n. 9, 174 n. 13,
175 nn. 14–18, 177 n. 23
- comparative use of appearance
verbs, defined 294–5
- compassion 240–1, 243
- completest fallacy 76
- confirm, defined 6
- confirmation, theory 2–3, 14–19
see also Bayes's Theorem
- consciousness, argument
from 192–212, 350–1
- control, region of basic 124–31
- converting events 275, 277
- copresence, regularities of, *see* order,
spatial
- cosmological argument 18, 133–52
- cosmological constant 175–6
- counterfactuals 31
- Craig, W. L. 138 n. 10
- create, defined 119 n. 7
- creation *e nihilo* 49–50
- Credulity, Principle of 170, 253,
283, 303–15, 322, 326–7
- curve-fitting problem 58–9
- Danto, A. C. 35 n. 10
- Darwin, C. 5, *see also* Evolution
- Davidson, D. 38–43
- Davis, C. F. 315 n.
- death, value of 228–31
see also life after death
- decay, world of 226, 253
- deductive arguments 4, 6–7, 136–7,
155, 329–30
- deism 94
- depravity 243
- Descartes, R. 8 n.
- design, arguments from 18, 153,
346–9
see also teleological arguments
- desires 42–3, 104, 217–8, 221,
226–7, 233, 268
- determinism of brain, *see* free will
- Devil, the 318
see also angels
- Dicke, R. H. 176 n. 20
- divine, defined 96
- Dionysus 316
- Dionysian principle 117–8
- dualism, event 195–6, 199
explanatory 38–45, 199
substance 196–9
- Earman, J. 181 n. 28
- Einstein, A., *see* Relativity, Theory of
electromagnetism 205
- emotions 226–8
see also compassion *and* desires
- environment, responsibility for 233
- epistemic distance 269
see also hiddenness
- epistemic use of appearance
verbs 294–5

- evangelization 267, 271
 event, defined 40–1
 Everything, Theory of 201–2
 evidence, background, *see*
 background evidence
 tautological 18, 67–8
 evil, defined 219 n.
 moral defined 236
 natural defined 236
 problem of 219–72
 quantity of 263–7
 Evolution 171, 188–91, 206–9, 216,
 217 n., 346–9, 350–4
 Exodus, Book of 316
 experience, interpretation of,
 defined 308, 317 n.
 real, defined 308, 317 n.
 explanation, 23–5
 absolute 79, 148
 axiarchic 47 n.
 complete, justification of 77–8,
 80–92
 deductive-nomological 26–7
 full, defined 25–6, 76–8
 more than one of a
 phenomenon 45–7
 partial, defined 25–6
 personal 21, 23, 35–45, 61–6,
 87–9
 of God's action 47–9
 scientific 21, 26–35, 52–61, 80–7,
 88–9
 substances-powers – and
 liabilities account 33–5,
 44–5
 ultimate, defined 78–9
 explanatory power, defined 56–8,
 65, 66–7

 faith 2
 fawn caught in fire 244, 256, 261
 fine-tuning 172–88
 Flew, A. G. N. 12 n., 316 n.

 four forces, *see* standard theory
 free will, human 113, 119, 169–70,
 234, 257, 272
 free will defence to moral evil, *see*
 evil, problem of

 Galileo, G. 5, 69, 75, 78, 81, 140 n.
 Gildenhuys, P. 216 n. 7
 Gimello, R. M. 295 n. 2
 God, defined 7, 93–6, 345
 creator 7, 48–50, 94, 99, 107–8,
 112–23
 eternal 7, 80, 96, 99, 335
 logically contingent 79, 148
 necessary being 95–6, 99, 148
 necessary (essential) properties
 of 95–6, 335–6
 omnipotent 7, 94–5, 334–5
 omnipresent spirit 7, 50–1, 93–4,
 99
 omniscient 7, 97–8, 100, 334–5
 perfectly free 7, 98, 105, 335
 perfectly good 7, 95, 99–106,
 113–4
 person 7 n. 1, 97, 344–5
 right to inflict harm 238, 257–63
 signature of 286
 simple, *see* theism, simplicity of
 uniqueness of 98–9, 343
 see also theism
 Goldman, A. 42 n. 15
 goodness, moral 100–5, 113–4
 see also moral truth
 Grand Unified Theory 201–2
 greater good, evils serving,
 defined 237–8
 Grice, H. P. 300 n.
 'grue' 54

 Halley, E. 83 n. 7.
 Hampshire, S. 100
 Harré, R. 33
 Hawking, S. W. 176 n. 20

- Heaven 266 n. 8, 296
 see also life after death
 Hempel, C. G. 26–9, 30 n. 3
 Hesse, M. 332–3
 Hick, J. 302
 hiddenness of God 130, 231, 253,
 267–72, 275, 293, 304
 higher-order goods defence 240–5
 history, arguments from 273–7
 humanly free agents 118–31,
 169–71
 Hume, D. 2, 29, 33, 58 n., 76,
 134 n., 136, 140–1, 145,
 165 n., 167–8, 282 n., 283–4,
 288.

 identity of indiscernibles 135 n. 3
 incarnation 264–5, 288–91
 inductive inference 4–7, 12–21
 correct criteria of 52–72, 208,
 330, 353–4
 normal 245–53, 255–6
 infinitesimals 178 n.
 inflation theory 180–1
 intention 35, 39 n., 40–3
 irreducibly complex systems 346–9
 Islam 7, 266, 286, 288, 295 n. 2,
 324 n., 344–5
 Israel, *see* Judaism

 Jeremiah 153, 302
 Joseph's dream 299–300
 Judaism 7, 286, 295 n. 2, 344–5

 Kane, G. L. 181 n. 29
 Kant, I. 2, 10–12, 18 n., 58 n., 133,
 136, 147 n., 148 n.19, 212,
 214
 Katz, S. T. 317 n.
 Kepler, J. 45–6, 58–9, 69
 kidnapper example 156–7
 Kretzmann, N. 117
 Kripke, S. 95–6

 laws of nature 5, 26–35, 44, 153–4,
 277–88
 fundamental 140 n., 157–60,
 221, 278–9, 347
 Hempelian model of, 26–30, 43
 probabilistic, statistical 27–30
 see also Quantum Theory
 regularity (Humean) account
 of 29–31, 160
 S–P–L account of 33–5, 162–4
 universal (deterministic) 26–9
 universals account of 32–3,
 161–2
 Lazarus, parable of Dives
 and 266 n. 8
 Leibniz, G. W. 114, 136, 143, 147–9
 Leslie, J. L. 47 n., 184 n.
 Lewis, D. 30–1
 light 205
 life, goodness of 259
 life after death 231, 261–2, 264–6,
 268
 liked, desire to be 268–70
 Limbo 266 n. 8
 Locke, J. 192–3
 Lovejoy, A. O. 116 n.

 McGrew, T. 178 n.
 machine rooms 127–9, 220
 MacIntyre, A. 12 n., 305 n.
 Mackie, J. L. 70
 Madden, E. H. 33
 Mary, Blessed Virgin 296, 299, 304,
 318 n., 319–20, 325
 Mellor, D. H. 131–2
 memory, apparent 283, 305–6
 mental events, properties and
 substances, defined 193–4
 metaphysical concepts, humans
 need 211–2
 metastability 173, 183
 Miller, K. R. 346 n.
 miracles 18, 274–92

- monism, substance 196
 moral awareness, argument
 from 215–8
 moral goodness, *see* goodness,
 moral
 moral truth, argument from 212–5
 Mostevin, J. 181 n. 28
 Muhammad 273
 multiverse 134, 165, 185–7

 necessity, factual 96, 148 n. 19
 logical 8, 148 n. 19
 physical 28–9, 31–3, 148 n. 19,
 156
 negligence, opportunities for 226
 Newton, I. 5, 45–6, 56–7, 60, 69, 75,
 76, 78, 81–4, 87, 140 n., 178,
 331, 333–4
 normalizability problem 178 n.
 numinous experience 295 n. 2

 obligation, moral 101–2
 occasionalism 107 n.
 ‘odd’ phenomena 74, 209, 281, 287,
 326
 ontological arguments 8–9, 148
 n. 19
 Oppenheim, P. 26
 opportunity to provide for
 others 222–31
 for self 220–2
 order, spatial 153–4, 167–89
 temporal 153–66
 Otto, R. 295 n. 2

 P-inductive arguments 6, 330–3
 pain, nature of 242, 248
 Paley, W. 167
 Parfit, D. 47 n.
 Paul, St 299
 Pauli principle 174, 184
 Peebles, P. J. E. 176 n. 20
 Penelhum, T. 304 n.

 perception 296
 region of basic 124–31
 person, defined 21 n. 16
 phenomenal laws,
 regularities 157–60, 221
 physical events, properties, and
 substances, defined 193–4
 physical universe, defined 133–4
 beginning of 138–40, 162–4, 167,
 171–2, 175–6, 180
 complexity of 150–1
 God’s reasons to create 121–2,
 127
 Plantinga, A. 211 n., 239 n. 3,
 350–4
 Plato 116 n.
 poltergeists 63–4, 71, 88
 polytheism 146–7, 149, 339–40
 Pope, A. 83 n. 7
 Poseidon 296, 304, 316, 319–20
 prayer, answers to 274, 276, 285
 prediction (in literal sense) 69–70
 predictive power, defined 56
 probability calculus 67; *see also*
 confirmation theory
 comparative 68
 epistemic 16 n. 12
 inductive 14–20
 inscrutable 350 n. 3
 intrinsic, defined 67–8
 logical, defined 15–16
 physical 15, 29, 31–3
 prior 53, 65, 67, 240
 subjective 16 n. 12
 zero of scientific theories
 331–3
 property, defined 40
 prophets 273–6, 286–7
 providence, argument from
 219–35
 psychophysical theory 201–9
 punishment for sin 238–9
 purpose, *see* intention

- Quantum Theory 29, 30, 169–70,
172, 174, 179, 184, 201, 279,
281, 331
- quasi-violation, defined 280–1
- rational agent, defined 21
- rationality, value of 255–6
- Relativity, theory of 19, 30, 78, 83,
140 n., 172, 201, 279, 331
- relevance criterion 70, 110
- religions, different, conflicts
between 288, 316–17, 324 n.
- religious experience, argument
from 293–327, 341–2
kinds of 298–303
more specific 324–5
nature of 293–8
private, defined 297–8
public, defined 297–8
- responsibility for others 119–20,
222–26
- revelation 286, 290
- Rice, D. H. 47 n.
- Rowe, W. L. 136 n. 6, 244 n.
- Salmon, W. C. 27 n.
- sceptical bog 304 n., 306
- Schellenberg, J. 267 n.
- science, opportunity to
pursue 221–2, 254–5
- sciences, integration of 205–6
- scientific concepts, humans
need 211–12
- scope of theory 53, 57, 59–60,
333–4
- Scotus, J. Duns 90–2, 343–4
- semi-divine beings 119
- simplicity 53–66, 70–71, 146,
177 n. 24, 178–9, 185 n.
- singularity, initial 106, 139, 163–4,
179
- sloth 226
- Smart, R. N. 295 n. 2
- Sober, E. 216 nn. 7–8
standard theory 172, 177–9
- Stace, W. T. 317 n.
- statistical mechanics 206
- Stawson, P. F. 296 n. 4
- string theory 181 n. 29
- substances, defined 40, 194
- succession, regularities of, *see* order,
temporal
- suffering, limit to 229–30, 261–5
- Sufficient Reason, Principle
of 147–9
- suicide 259, 261
- supererogation 102
- Swinburne, R. G. 138 nn 8–9
Epistemic Justification 16 n. 12,
53 n., 54 n., 67 n. 8, 70 n. 9,
178 n., 303 n., 322 n.,
327 n., 332 n. 1
Faith and Reason 2
Is There a God? 21 n. 17
*Providence and the Problem of
Evil* 10, 236 n.
*Responsibility and
Atonement* 102 n., 114 n. 1
Space and Time 91 n. 9
The Christian God 7 n. 3, 77 n. 3,
93 n., 119 n. 7, 344–5
The Coherence of Theism 1, 7 n.
2, 9 n. 6, 21 n. 16, 80 n.,
93 n., 95 n. 2, 100 n. 7,
136 n. 7, 148 n. 19, 214 n. 5
The Evolution of the Soul 170 n.,
201 n.
*The Resurrection of God
Incarnate* 10 n. 8, 266 n. 9,
291 n., 342 n.
- sympathy 241–4
- Taylor, R. 42
- Tegmark, M. 185 n.
- teleological arguments 153–91,
346–9

- temptation 227, 237, 241, 268
see also desires
- Tennant, F. R. 190 n
- Testimony, Principle of 283, 322–4
- theism, defined 7, 344
 explanatory power of 110–32,
 336–9
 intrinsic probability of 93–109,
 333–6
 provides complete
 explanation 80, 98–9, 108,
 334
 provides ultimate
 explanation 96, 108
 scope of 108–9, 332–3
 simplicity of 96–109, 334–6
- thermodynamics 205–6
- time, instants and periods 138
- Tipler, F. J., *see* Barrow, J. D.
- Tooley, M. 32 n. 7
- traces, physical 283
- Trinity, the 118–19, 343–5
- tuning, defined 174
see also fine-tuning
- unextended point, *see* singularity
- universe, *see* physical universe
- unobservables 57–8, 63–4
- violations of laws of nature,
 defined 277
see also miracles
- von Wright, G. H. 35 n.9
- Wainwright, W. J. 320 n.
- weakness of will 104
see also temptation
- ‘what’ of explanation, defined
 24
- ‘why’ of explanation, defined 25
- Wilson, D. S. 216 nn. 7–8
- Wolterstorff, N. 285 n.
- world, possible, defined 114
 n. 2
- world-I 228
- world-II 228
- world-III 228–30
- world-IV 230
- Wynn, M. 146 n., 190