



OXFORD

A World For Us

The Case for Phenomenalistic Idealism

JOHN FOSTER

A WORLD FOR US

This page intentionally left blank

A World for Us

The Case for Phenomenalistic Idealism

JOHN FOSTER

OXFORD
UNIVERSITY PRESS

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 Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

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

© John Foster 2008

The moral rights of the author have been asserted
Database right Oxford University Press (maker)

First published 2008

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 the same condition on any acquirer

British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloguing in Publication Data

Foster, John, 1941 May 5-

A world for us: the case for phenomenalist idealism/John Foster.

p. cm.

Includes index.

ISBN 978-0-19-929713-9

1. Idealism. 2. Realism. 3. Phenomenology. I. Title.

B823.F675 2008

141—dc22

2008001087

Typeset by Laserwords Private Limited, Chennai, India

Printed in Great Britain

on acid-free paper by

Biddles Ltd., King's Lynn, Norfolk

ISBN 978-0-19-929713-9

1 3 5 7 9 10 8 6 4 2

To
George Berkeley

This page intentionally left blank

Preface

Put at its simplest, the aim of this work is to establish that the existence of the physical world is logically sustained by the world-suggestive way in which, under God's ordinance and authority, things are disposed to appear at the human empirical viewpoint. This idealist thesis—a version of what I speak of as *phenomenalistic* idealism—stands in sharp contrast to the commonly accepted realist view, which takes the world to have an existence that is both logically independent of the human mind and metaphysically fundamental. I argue that it is only by accepting the idealist thesis that we can represent the physical world as having the empirical immanence it needs if it is to form a world *for us*.

A world whose existence is logically sustained by how things are disposed to appear at the human empirical viewpoint sounds as if it would be, at best, a virtual reality—a mere experiential simulation of a world. But in the idealism for which I argue what logically sustains the existence of the world is not the world-suggestive system of appearance on its own, but this system as ordained and authorized by God, and I see this as making a crucial difference. The God in question is the God of the Judaeo-Christian tradition, and I conceive of him as having the full perfection of nature that this tradition accords him.

The framework in which I hold this conception is one which represents God as a personal being—as a rational mental subject—and so as an entity of the same general category as ourselves. In particular, within the context of my argument, I take God to be a being who has knowledge and purposes, and who brings things about by the exercise of his will. All this, I am confident, is in line with what most ordinary Judaeo-Christian believers implicitly accept. At the same time, I realize that some philosophers and theologians will find the approach unacceptably crude. They may say, for instance, that while God is indeed a being, the sort of being he is cannot be positively specified in any terms that we can understand. Or they may say that God should not be thought of as *a* being (as an entity among entities), but simply as *being* itself, or perhaps as *active being*. Or again, they may insist that the ontological nature of God is beyond our comprehension altogether. In one way or another, these philosophers and theologians will think that my representation of God as a personal being detracts from his essential mystery, and thereby from the

perfection and unsurpassable greatness that the Judaeo-Christian tradition ascribes to him. I do not discuss these delicate issues in the text of the book, nor shall I try to do so here. I am happy to acknowledge that the conception of God I employ might turn out to be in certain respects too crude, and stand in need of modification. But if it does, this will not, in the end, undermine the force of my argument. Provided that the modification does not prevent me from being able to think of God as ultimately responsible, and in some way *purposively* responsible, for the way in which things are disposed to present themselves at the human empirical viewpoint, the role that God is required to play in my idealist account will not be affected. And a conception of God that did not allow me to assign such responsibility to him would not preserve traditional theism in any recognizable sense.

Some readers will be aware that this is not the first book I have written which is devoted to a defence of phenomenalistic idealism. There are two reasons why *The Case for Idealism*, published in 1982, does not make the present work redundant. The first is simply that while the new book preserves a significant resemblance to the earlier one, it also includes extensive changes—changes in the topics covered, in the nature of the idealist thesis I endorse, in the arguments I develop in order to establish this thesis, and in the objections to it that I consider and try to meet. Someone who has read and understood the earlier book should not, on reading the present one, feel that he has travelled this route before. The second reason is that the whole style of the new book is, I hope, much more accessible than the first: wherever possible, I have avoided technicalities, and quite generally I have done my utmost to make a complex and difficult topic clear and comprehensible. The technical precision of *The Case for Idealism* may have a certain virtue in the abstract, but it made enormous, and arguably intolerable, demands on the patience of the reader.

In writing the present book I have had to work more or less on my own, partly as a result of serious ill health, partly because it is hard to find philosophers who work in this area. So I have not on the whole been able to gain the benefit of friendly criticism and advice. A notable exception is provided by the many discussions I have had over the years with Howard Robinson. He has helped me greatly to clarify and refine my thinking about a number of issues that feature in the work, and I would like to record my thanks to him here.

There is one other philosopher I need to acknowledge, and in a sense thank: George Berkeley, the originator of phenomenalistic idealism as I conceive of

it. Berkeley's version of this idealism is not entirely to my liking. In particular, he has what is, from my standpoint, an impoverished view of the sorts of entity and property that the idealistic world can contain. Sometimes, too, he seems to stray from the path of anything I would count as idealism at all. But it is to his vision of a world that is created by the orderly way in which God brings about our sensory experiences that my own approach can be ultimately traced. Although I do not discuss Berkeley's idealism in any detail, in dedicating this book to him I am signalling my considerable debt to that vision.

John Foster
Brasenose College, Oxford

This page intentionally left blank

Contents

1. The Problem of Perception	1
2. The Inscrutability of Intrinsic Content	42
3. Realism and Phenomenalistic Idealism	83
4. The Refutation of Realism	123
5. The Challenge of Nihilism	164
6. The Issue of Objectivity	199

<i>Bibliography</i>	247
<i>Index</i>	249

This page intentionally left blank

1

The Problem of Perception

I

We ordinarily take it for granted that we have perceptual access to the physical world, and that through it we acquire our information about the world and our situation within it. But how should this perceptual access be understood? What is it for someone (a human subject) to perceive (perceive by his senses) a physical item? There are a number of specific theories of the nature of such perception, but we can classify them all under two rival general views. It is on these rival views, and the issue between them, that I want, in this opening phase of the discussion, to focus. I shall do so, initially, within the framework of the common-sense assumption that the physical world is ontologically independent of the human mind—that it is something whose existence is logically independent of facts about human mentality. The adoption of this framework may seem hardly worth mentioning. How could the world of space and material objects be anything other than mind-independent? But what I shall try to show is that, although the two rival views exhaust the range of possibilities, or at least do so for any given case of perception, neither of them can be made to yield a satisfactory account so long as the assumption of mind independence is retained. It is this that creates what I see as the *problem* of perception.

Much of what I shall say, in elaborating this problem, draws on points that I have developed in more detail in my book *The Nature of Perception*.¹

II

Before I can state the rival views, I need to introduce and explain a key concept. The concept is of something I call *constitution*, and it will play a major role not

¹ John Foster, *The Nature of Perception* (Oxford: Oxford University Press, 2000).

only in the present context, where we are focusing on the topic of perception, but also throughout the rest of the discussion. Indeed, the central topic of the whole discussion concerns a certain thesis of constitution—a thesis about what constitutively sustains the existence of the physical world.

Constitution, in the sense I intend, is a two-place relation that holds within the domain of facts (instances of what is the case) and sets of facts. It can be provisionally defined as follows:

A fact F is *constituted by* a fact F', or by a set of facts S, if and only if two conditions are satisfied, namely:

- (1) The obtaining of F is logically due to the obtaining of F' (the obtaining of the members of S).
- (2) The obtaining of F involves nothing over and above the obtaining of F' (the obtaining of the members of S).

When I speak here of the *obtaining of a fact*, I do not mean to refer to the higher-order fact that the relevant fact obtains. I only mean to refer to the fact itself, but in a way that conveys its nature as an instance of what is the case. So, when I speak of the obtaining of a fact F as logically due to (or as involving nothing over and above) the obtaining of a fact F', this is just another way of saying that, if F is the fact that p, and if F' is the fact that q, its being the case that p is logically due to (involves nothing over and above) its being the case that q.

Thus defined, constitution comes in two forms: a *single-fact* form, in which a fact is constituted by another (single) fact, and a *multi-fact* form, in which a fact is constituted by a set of two or more facts. Strictly speaking, the definition also allows for the case of a fact being constituted by a set containing just one fact. But this case can be ignored, since to speak of constitution by the one-membered set {F'} is just a contorted way of speaking of constitution by its sole member F'. Where a fact is constituted by a set of facts (a set of two or more facts), I shall also often speak of it as constituted by the facts themselves, as a collective plurality. And, to emphasize the plural nature of what is doing the constituting, I shall sometimes speak of the constituted fact as *breaking down*, or *decomposing*, into these facts. Any case of multi-fact constitution can, of course, be recast as a case of single-fact constitution, by simply replacing the relevant set of facts by a fact that conjoins its members; or, at least, this can be done if the number of its members is finite. But all this shows is that a case of single-fact constitution is only *interestingly* single-fact if the constitutive fact in question is not a conjunctive one.

I have described the definition above as *provisional*. This is because it will need to be slightly revised if it is precisely to capture the concept of constitution I have in mind. But before we can appreciate the need for this revision, we have to get clear about the nature of the definition in its present form, and the content of the two conditions for constitution that it prescribes.

One thing which each of these conditions is to be understood as implying is that the obtaining of F' (the members of S) logically necessitates (logically ensures, guarantees, suffices for) the obtaining of F . So, in the single-fact case, if F is the fact that p , and if F' is the fact that q , each of the conditions implies that it is logically necessary that if q , then p ; and in the multi-fact case, if F is the fact that p , and if S is the set of facts that q_1 , that q_2 , ..., each of the conditions implies that it is logically necessary that if q_1, q_2, \dots , then p . By 'logical' necessity, I mean, as is standardly meant, *strict* or *absolute* necessity: it is *logically* necessary that p if and only if there is no possible world of any sort—not even a world with different natural laws—in which it is not the case that p . As Kripke reminded the philosophical world, it is important not to confuse the question of what holds as a logical necessity in this sense with the question of what can be established a priori. We need empirical evidence to establish that water, as ordinarily conceived, is the same substance as H_2O , as chemically defined. But granted that water and H_2O are the same substance, the fact of their identity is logically necessary in the relevant sense, since, if A is the substance water, and B the substance H_2O , there is no possible world of any sort in which A and B are numerically different.

Each of the conditions (1) and (2) is to be understood as implying that the obtaining of F' (the members of S) logically necessitates the obtaining of F . But there are two respects in which such necessitation does not, on its own, suffice for constitution, as I am conceiving of it, and it is these respects which bring to light the further implications of the two conditions.

In the first place, logical necessitation is not, as such, asymmetric. There are cases in which the obtaining of a fact, or set of facts, logically necessitates the obtaining of a fact, or set of facts, *and vice versa*. The obvious example is that in which the first fact, or set of facts, and the second fact, or set of facts, are the same, since, trivially, the obtaining of any fact logically necessitates itself. In contrast with this, I want the relation of constitution to be necessarily asymmetric: I want the relation to be such that, where a fact F is constituted by a fact F' (or a set of facts S), F derives its obtaining from (owes its obtaining to) the obtaining of F' (the members of S) in a way that—on pain, as it

were, of metaphysical circularity—precludes the same relationship holding in reverse. It is this element of asymmetric dependence which, in addition to mere logical necessitation, is expressed by saying, in condition (1), that the obtaining of F is *logically due to* the obtaining of F' (the obtaining of the members of S).

Second, there are cases where one fact is logically necessitated by another fact or set of facts, but where its obtaining has, as it were, an ontological life of its own, that lies outside the obtaining of the fact or set of facts that necessitates it. For example, suppose that, at a certain time t_1 , God prescribes that a certain kind of event will occur at the later time t_2 , and let F1 be the fact that God issues this prescription at t_1 , and let F2 be the fact that the relevant kind of event occurs at t_2 . Granted that God is (in some relevant sense) omnipotent, and is so essentially, the obtaining of F1 logically necessitates the obtaining of F2; but, given that the prescription and the subsequent event are ontologically discrete items, that occur at different times, there is a clear sense in which the obtaining of F2 (its being the case that an event of the relevant kind occurs at t_2) is something separate from—something genuinely additional to—the obtaining of F1 (its being the case that God issues the relevant prescription at t_1). Now I want the relation of constitution to exclude this kind of separateness: I want things to be such that, where a fact is constituted by another fact or set of facts, its obtaining is wholly included in the obtaining of this other fact or set of facts. And it is this inclusion that is expressed by saying, in condition (2), that the obtaining of F involves nothing over and above the obtaining of F' (the obtaining of the members of S).

Examples of constitution, in the sense defined, are not hard to find. One clear-cut range of single-fact cases is marked out by the principle that the instantiation of a generic (determinable) property is always constituted by the instantiation of a property that is more specific (more determinate). For instance, if an object is (generically) red, the fact of its being so is constituted, in the relevant sense, by a more specific fact about its colour, such as its being scarlet or carmine: the obtaining of the generic colour fact is logically due to, and involves nothing over and above, the obtaining of the specific fact. A clear-cut range of multi-fact cases is provided by the relationship between one object's being heavier than another and their individual weights. Thus, if John weighs twelve stone and Mary weighs ten stone, the fact that John is heavier than Mary is constituted by the combination of the facts that John weighs twelve stone and that Mary weighs ten stone: the obtaining of their weight relationship is logically due to, and involves nothing over and

above, the obtaining of these separate weight facts about them. These are *clear-cut* cases in the sense that the relevant claims of constitution are wholly uncontroversial: no one will think of denying that, in the kinds of situation envisaged, the two prescribed conditions for constitution are satisfied in the specified ways. Not surprisingly, the claims of constitution on which we shall be focusing in our philosophical discussion will not be uncontroversial in that way.

I have already indicated that the present definition of constitution will need to be slightly revised. The point of the revision is to allow for cases in which the obtaining of a fact has two independent modes of constitution. To focus on a simple example, let us, again, suppose that John weighs twelve stone and Mary weighs ten stone. We can draw up the following list of facts whose obtaining is implicit in this supposition:

F1: the fact that John weighs twelve stone.

F2: the fact that Mary weighs ten stone.

F3: the fact that either John weighs twelve stone or $2 + 2 = 5$.

F4: the fact that either Mary weighs ten stone or $2 + 2 = 5$.

F5: the fact that either John weighs twelve stone or Mary weighs ten stone.

Whenever it is true that p and false that q , the obtaining of the fact that p or q is logically due to, and involves nothing over and above, the obtaining of the fact that p . So, the obtaining of F3 is logically due to, and involves nothing over and above, the obtaining of F1, and the obtaining of F4 is logically due to, and involves nothing over and above, the obtaining of F2. This means that F3 is constituted by F1 and that F4 is constituted by F2. But now consider how F5 stands to F1 and F2. Granted that F3 is constituted by F1, we want to be able to say that F5, too, is constituted by F1, since the obtaining of F1 bears on the obtaining of F3 and on the obtaining of F5 in exactly the same way. Similarly, granted that F4 is constituted by F2, we want to be able to say that F5, too, is constituted by F2, since the obtaining of F2 bears on the obtaining of F4 and on the obtaining of F5 in exactly the same way. But these are not things that we are able to say under our present definition of constitution. Given that it is logically necessitated by the obtaining of F2, we cannot say, without qualification, that the obtaining of F5 is logically due to, or involves nothing over and above, the obtaining of F1; and, given that it is logically necessitated by the obtaining of F1, we cannot say, without qualification, that the obtaining of F5 is logically due to, or involves nothing over and above, the obtaining of F2. All we can say is

that, *with respect to its necessitation by F1*, the obtaining of F5 is logically due to, and involves nothing over and above, the obtaining of F1, and that, *with respect to its necessitation by F2*, the obtaining of F5 is logically due to, and involves nothing over and above, the obtaining of F2.

The solution is to reformulate the two conditions for constitution in a way that relativizes what they claim to the context of a single source of necessitation. So, the revised definition of constitution will read:

A fact F is *constituted by* a fact F', or by a set of facts S, if and only if the obtaining of F is logically necessitated by the obtaining of F' (the obtaining of the members of S), and, with respect to that source of necessitation:

- (1) the obtaining of F is logically due to the obtaining of F' (the obtaining of the members of S);
- (2) the obtaining of F involves nothing over and above the obtaining of F' (the obtaining of the members of S).

This will allow us to say that F5 is separately constituted by both F1 and F2; and it will accommodate a whole range of cases of a similar kind, in which a single fact has two or more independent modes of constitution.

Although the new definition is needed for the purposes of precision, the issue of independent modes of constitution will not be relevant to the topics of our future discussion, and in what follows I shall normally, for convenience, continue to express the conditions for constitution in their simpler original form, without relativizing their content to a source of necessitation.

III

With the concept of constitution in place, we must now turn to the topic of physical-item perception, and the two rival general views about its nature. In setting out and discussing these views, I shall work on the assumption that we can think of human subjects as perceiving items at points (moments) of time, with zero temporal extent. This assumption does not imply, what would clearly be incorrect, that a momentary instance of perceiving can occur in isolation, without being contained within an extended episode of perceiving, and once that point is understood, the assumption may not seem unreasonable. But, in any case, my reason for adopting it is more one of expositional convenience than philosophical conviction; and certainly nothing of substance will turn on it. If anyone objects to the assumption,

on the grounds that any genuine instance of perceiving must be temporally extended, I am happy for him to read what I say with the adjustments that his position requires. To harmonize with my recognition of momentary instances of perceiving, I shall also recognize certain kinds of momentary physical item to serve as objects, or potential objects, of perception. In particular, among such objects, I shall recognize the momentary stages of persisting physical items, where each such stage is something whose existence covers all and only what is covered by the existence of the relevant persisting object at a particular point in time. Anyone who wants to make adjustments to what I say in the case of momentary instances of perceiving will need to make appropriate adjustments here, too. Strictly speaking, some of the things that I represent as momentary stages of persisting items are really entities whose spatial points, while of zero temporal extent, are spread over a *range* of moments. But, for simplicity of exposition, that is something I shall largely ignore.

To enable me to formulate the two rival views clearly and concisely, it will be helpful if I begin by making terminological provision for certain further key notions, specific to the topic of perception.

First, given a subject *S* and two items *x* and *y* that he simultaneously perceives, I shall say that *S*'s perceiving of *x* (the fact of his perceiving *x*) is *mediated by* his perceiving of *y* (the fact of his perceiving *y*) if and only if (1) *S*'s perceiving of *x* is constituted by the combination of his perceiving of *y* and certain additional facts, and (2), apart from any concern they may have with *S*'s perceiving of *y*, these additional facts do not involve anything about *S*'s perceptual condition at the relevant time. Where a subject's perceiving of one item is mediated, in this way, by his perceiving of another, we can speak of him as perceiving the second item *more immediately* than he perceives the first. I shall illustrate this notion of perceptual mediation presently.

Second, I shall say that a subject *S* ϕ -terminally perceives an item *x* at a time *t* if and only if *x* is a physical item and *S* perceives *x* at *t* and there is no other physical item *y* such that *S*'s perceiving of *x* at *t* is mediated by his perceiving of *y* at *t*. So, the perceiving of an item qualifies as ϕ -terminal just in case the item is physical and there is no other physical item that is, in the context of that perceiving, perceived more immediately. (This leaves open the possibility of there being a *non*-physical item that is perceived more immediately.) I shall take it for granted that physical-item perception is not infinitely regressive, and that whenever a physical item is perceived, there is some physical item which is, in respect of that perception, ϕ -terminally perceived.

Finally, when I speak of a psychological state as *in itself physically perceptive*, I mean that it is logically impossible for someone to be in that state without thereby perceiving a physical item, and when I say that a psychological state is *in itself perceptive of an item x*, I mean that it is logically impossible for someone to be in that state without thereby perceiving x. Obviously, I am here using the term 'state' to mean *type* state, not *token* state. So, a psychological state is something that is capable of realization in different subjects and on different occasions.

The two rival views of perception can now be set out as follows. According to the first view, whenever someone perceives a physical item, there is some psychological state which is not in itself physically perceptive, and the fact of his perceiving that item breaks down into (is constituted by the combination of) two components. One of these components consists in his being in that state. The other comprises certain additional facts, but ones that do not involve anything further about his psychological condition at the relevant time. In practice, the advocate of this view will take these additional facts to concern, or concern amongst other things, the qualitative relationship of the psychological state to the physical item (or to the item ϕ -terminally perceived), and the role of the item in causing the subject to be in that state at that time. I shall call this first view the *decompositional* view. According to the second view, whenever someone ϕ -terminally perceives a physical item, the fact of his perceiving it is something psychologically fundamental. It is not something that breaks down into (is constituted by the combination of) his being in some further psychological state, which is not in itself physically perceptive, and certain additional facts, not involving anything further about his current psychological condition. It does not, at the psychological level, break down into further facts at all, except perhaps (if this is possible) in a purely trivial way, where the perceived item is something complex and the subject's perceiving of it breaks down into the separate perceivings of its parts. In other words, the psychological state that is fundamentally involved in the perceiving of the relevant physical item is one that is in itself perceptive of that item. I shall call this second view the *fundamentalist* view.²

It is obvious that, for any case of physical-item perception, the two views are, with respect to the ϕ -terminal perceiving involved, mutually exclusive: they are explicitly formulated so as to be so. It is also true, though not

² The decompositional and fundamentalist views coincide with what, in *The Nature of Perception*, I respectively labelled the *broad representative theory* and *strong direct realism*, except that the latter positions were defined as ones that explicitly endorsed a realist conception of the physical world.

guaranteed by their mode of formulation, that the views are, case by case, jointly exhaustive. For where an item is ϕ -terminally perceived, the only way in which we could sensibly think of the psychological state that is fundamentally involved as failing to be in itself perceptive of that item would be by thinking of it as failing to be in itself physically perceptive altogether; and if the state is not in itself physically perceptive, the decompositional account of the perceiving is the only one available.

We can best bring out the nature of the two views by focusing on a particular case. Suppose, at a certain time t , Ralph sees an apple on the table in front of him. His seeing of the apple is mediated, in the sense defined, by his seeing of a certain portion of its surface (it is constituted by the combination of his seeing this portion and the fact that this latter item is a portion of the apple's surface), and his seeing of this persisting surface portion is, in an exactly analogous way, mediated by his seeing of a certain momentary stage of it—a stage which, depending on the precise temporal length of the causal process from the surface portion to its visual registering at t , occurs slightly earlier than t .³ This momentary stage of the apple's surface portion is then what, in the sense defined, Ralph ϕ -terminally sees at t —what, relative to the domain of physical candidates, he most immediately sees. It is in terms of his visual contact with this ϕ -terminal object that the two views come into sharp conflict. Let us call this ϕ -terminal object O . The fundamentalist will say that, given that O is the ϕ -terminal object of perception, Ralph's seeing of O is a fundamental aspect of his psychological condition at t : his fundamental psychological state, though wholly a matter of what is occurring within his mind at that time, inherently involves his standing in this awareness relation to this external physical item. In contrast, the decompositionalist will say that, instead of being psychologically fundamental, the fact of Ralph's visual contact with O breaks down into two components. One component will cover all the relevant aspects of Ralph's psychological condition at t —all that obtains or occurs in his mind at that time that in any way logically contributes to the obtaining of the relevant perceptual fact—and the decompositionalist will insist that these aspects, on their own, do not secure perceptual contact with O or suffice for physical-item perception at all. The other component will cover the remaining facts that are relevant to the securing of visual contact with O —facts that do not involve anything further about the

³ This is an example of something that I am representing as a momentary stage, but which in reality, because of the slightly varying length of the causal process from different points on the surface portion, is spread over a range of moments.

subject's psychological condition at *t*. In practice, decompositionalists will take the first component to consist in the fact of Ralph's having a certain kind of visual experience at *t*, and will take the second component to consist in, or centrally involve, facts concerning the qualitative relationship of this experience to *O* and the nature of the causal process from *O* to the occurrence of the experience. Fundamentalists too will take the seeing of *O* to involve the occurrence of a visual experience and will accept that this experience occurs at the end of a causal process starting from *O*. But they will take this experience to be one which, by its intrinsic psychological nature, puts the subject into perceptual contact with *O*. They will not think that either the qualitative relationship of the experience to the external item or the nature of the causal process from this item to the experience constitutively contributes to the securing of this contact.⁴

I have said that both decompositionalists and fundamentalists will recognize the existence of a causal process from the perceived physical item to the visual experience. But the ways in which they will understand the nature of this process are very different. For the decompositionalist, the nature of this process will be, in its general character, relatively straightforward. Light reflected from *O* enters Ralph's eyes, producing a certain pattern of firings in his optic nerves, which in turn brings about a further complex event in the relevant part of his brain, which finally results in the occurrence of the visual experience. At each stage in this process the character of what happens directly causally depends only on the character of what takes place immediately before it and the currently prevailing conditions. So the pattern of firings in the optic nerves directly depends only on the character of the photic input, together with the relevant facts about the structure and prevailing state of the subject's eyes and nervous system; the character of the brain event directly depends only on the character of the optic-nerve firings, together with the relevant facts about the structure and prevailing state of the brain; and, crucially, the character of the visual experience directly depends only on the character of the brain event and prevailing brain conditions, together with the relevant facts about the character of the subject's mind and the form of its cerebral embodiment. In other words, the causal process works in a standard serial way, whereby each stage in the process contributes to the final outcome only in so far as it affects what immediately follows it. But this cannot be how

⁴ They could still, of course, think that this psychologically fundamental contact logically could not have occurred without the experience being qualitatively and/or causally related to *O* in a certain way.

the fundamentalist will see the causal process. This becomes clear when we expand the original example to include a second seeing of the apple. Thus, suppose, after seeing the apple at t , Ralph closes his eyes for a few seconds, and then, on opening them, has another experience of seeing the apple, this time at t' . And let us call the ϕ -terminal object of this t' episode of seeing O' , an item which we can assume to be a correspondingly later momentary stage of the same portion of the apple's surface. Let us also suppose that the causal process from O' to brain is of exactly the same kind as that from O to brain, that the resulting brain events are of exactly the same kind, and that the cerebral and mental conditions in which these events occur are relevantly the same. Despite the qualitative identity of these processes, events, and conditions, the fundamentalist will take the resulting psychological events—the two visual experiences—to be, and to be at the fundamental level of description, of different psychological types, since he will say that one of these events (experiences), by its psychological character alone, puts Ralph into visual contact with O , and that the other, by its psychological character alone, puts him into visual contact with O' , a momentary stage that is later than, and so numerically different from, O . Clearly, then, the fundamentalist cannot, like the decompositionalist, think that the psychological outcomes in the two cases directly causally depend only on the character of the brain events that immediately precede them and the conditions in which these events occur. Rather, he will have to recognize certain additional factors which, by combining with these events and conditions, account for why the resulting outcome is in the one case a seeing of O and in the other case a seeing of O' . And it is not difficult to see what these factors will be. Obviously, from the standpoint of the fundamentalist view, what will combine with the relevant brain event and prevailing conditions to causally ensure a seeing of O at t will be the whole causal process from O to brain, and, in particular, the role of O (as distinct from any other physical item) as the initiator of that process; and what, subsequently, will combine with the relevant brain event and conditions to causally ensure a seeing of O' at t' will be the whole causal process from O' to brain, and, in particular, the role of O' (as distinct from any other physical item) as the initiator of that process. So, in each case, the causal process from the relevant physical item to the brain will play a double causal role with respect to the psychological outcome, being both what brings about the relevant brain event immediately preceding that outcome and what combines with that event, and the other relevant aspects of the prevailing conditions, to determine the specific perceptive character of the outcome. This may seem,

at first sight, a strange way for causation to operate—where the nature of what occurs at a given time is affected not just by what occurs and obtains immediately prior to that time, but also by the sequence of events leading up to it. But, on reflection, we can see that this kind of causation is precisely what fits the distinctive character of the fundamentalist position, in which, on any perceptual occasion, contact with the ϕ -terminally perceived item is taken to be a fundamental aspect of the subject's psychological condition, and where that item occurs earlier than the time of the perceiving. Once it is accepted that the fundamental character of the psychological outcome incorporates the existence of the perceived item and its perception by the subject, there is nothing strange about supposing that the process from the item to the subject's brain has a direct causal influence on it.

There is one other aspect of the fundamentalist view that needs to be underlined here. The fundamentalist will insist that the experience of seeing O and the experience of seeing O' differ in their fundamental psychological character. But he is not obliged to say that they differ in their *subjective* (*phenomenological*) character—in how they introspectively appear to the subject, or would appear if properly scrutinized. Indeed, he is likely to acknowledge that, since the relevant brain events and conditions in the two cases are of exactly the same type, the chances are that the way in which O sensibly appears to Ralph at t and the way in which O' sensibly appears to him at t' will be exactly the same as well, and that, in consequence, the two visual experiences will be, in terms of their introspective appearance, indistinguishable. There is nothing paradoxical in this. Once it has been accepted that the psychological state that is fundamentally involved in any perception is in itself perceptive of a particular physical item, it is inevitable that there is an aspect of what is going on psychologically that cannot reveal itself introspectively.

One reason for underlining this point is that it removes the temptation to attempt an instant refutation of the fundamentalist view by appeal to the phenomenon of hallucination—the phenomenon of experiences which subjectively pose as perceptions of certain kinds of physical item, but are not physically perceptive at all. Thus, on first encounter, we might have thought that we could prove the falsity of the view by advancing the following simple argument:

- (1) For any physical-item perception, there could be a subjectively matching hallucination.

- (2) So, for any physical-item perception, there could be an experience which is exactly the same in its fundamental psychological character, but is not perceptive of any physical item.
- (3) So, the psychological states fundamentally involved in physical-item perception are not in themselves physically perceptive.
- (4) So, the fundamentalist view is false.

But, given the nature of his position, the fundamentalist will simply reject the move from (1) to (2). And he can do so without embarrassment. For if there is no problem in his holding that the experiences of seeing O and O' may be subjectively indistinguishable, while differing in their fundamental psychological character, there is obviously no problem in his holding that experiences that are physically perceptive differ in their fundamental psychological character from those that are not.⁵ Nor, as is sometimes thought,⁶ can we create a problem by focusing on a case where, as well as the perceptive and hallucinatory experiences subjectively matching, the cerebral and mental factors that are causally involved in their occurrence are of exactly the same type. For, as we have seen, the fundamentalist does not suppose that such cerebral and mental factors are the only factors that causally contribute to the psychological character of the resulting experience. If, in the case of Ralph, he can say, without embarrassment, that the role of the environmental item in producing the relevant brain event is causally responsible for ensuring that the resulting visual experience is in itself perceptive of that item, he can also say, without embarrassment, that the presence of an environmental item that plays the right kind of role with

⁵ Distinguishing the fundamental psychological characters of perceptive and hallucinatory experiences is an instance of what is often described as a *disjunctivist* approach to the nature of perceptual experience. The approach originates with the work of J. M. Hinton; see, in particular, his *Experiences* (Oxford: Oxford University Press, 1973). Several philosophers have followed Hinton in offering disjunctivist accounts of the content of perceptive and hallucinatory experience, including Paul Snowdon ('Perception, vision, and causation', *Proceedings of the Aristotelian Society*, 81 (1980–1), 175–92, 'The objects of perceptual experience', *Proceedings of the Aristotelian Society*, supp. 64 (1990), 121–50), John McDowell ('Criteria, defeasibility, and knowledge', *Proceedings of the British Academy*, 68 (1982), 455–79), and Michael Martin ('The transparency of experience', *Mind and Language*, 17/4 (2002), 376–425, 'The limits of self-awareness', *Philosophical Studies*, 120 (2004), 37–89).

⁶ Such a thought is implicit in Howard Robinson's argument against naive realism in his 'The general form of the argument for Berkeleyan Idealism', in J. Foster and H. Robinson (eds.), *Essays on Berkeley* (Oxford: Oxford University Press, 1985), 163–86, an argument that he further develops in his *Perception* (London: Routledge, 1994), 151–62. I, too, once thought that the fundamentalist view could be undermined by considerations of this sort. Thus, see my *Ayer* (London: Routledge & Kegan Paul, 1985), 147–9, 161.

respect to what happens in the subject's brain is causally required if the resulting experience is to be in itself physically perceptive at all.⁷

There would still, of course, be the question of what account the fundamentalist should give of the psychological character of hallucination—a character which he has to take to be fundamentally different from that of perception, while able to simulate it subjectively. It is possible that there would be problems for him on that front. But this is not an issue that we can hope to address until we have considered in more detail his position on the nature of perception. When we do, we shall see that, quite apart from any issue about hallucination, his fundamentalist account of the ϕ -terminal perceptual relationship is untenable.

IV

I have formulated the two rival general views about the nature of physical-item perception, and illustrated them by reference to a particular case. It is also clear that this case is entirely typical, and that the way in which the two views work out in other cases of ϕ -terminal perceiving would be, in all basic respects, the same. The next task must be to try to evaluate the views, and so decide the issue between them. I shall begin by considering the fundamentalist view. This, as we have seen, claims that, in any case of ϕ -terminal perceiving, the perceptual relationship between the subject and the perceived physical item is something psychologically fundamental—something which does not, at the psychological level, decompose into further factors. I have already stressed that the fundamentalist view cannot be refuted by a simple appeal to the phenomenon of hallucination. But there is a quite different area where its capacity to provide an adequate account needs to be scrutinized, and it is on this that I now want to focus.

Whenever someone perceives a physical item, he perceives it, or at least perceives whatever it is that he ϕ -terminally perceives, under a certain sensible appearance. Thus, when Ralph sees the apple, it (or, at any moment, the portion stage of it he ϕ -terminally sees) visually appears to him as a roughly hemispherical patch of a certain size and colouring, located at a certain distance in front of him (or, more precisely, whose various parts are located at various distances and in various directions from him). Likewise, when I

⁷ For a fuller elaboration of this point see my *The Nature of Perception*, 23–43.

take a bottle of wine from the fridge, the bottle (or, at any moment, the portion stage of it I ϕ -terminally feel) tactually appears to me as the hard, cold, smooth, curved surface of something I am holding. Quite generally, in any case of ϕ -terminal perceiving, the perceived item sensibly appears to the subject in a certain way—a way which represents it as an environmental item with a certain sensible character, and, at least typically, assigns to it, in perspective, a more specific environmental location, or set of locations, relative to the subject's own body or current position. The sensible character that the perceived item appears to possess draws its elements exclusively from qualities of spatial and temporal arrangement and qualities that are distinctively associated with the relevant sense realm. The precise nature of this distinctive association need not concern us, but I am thinking of the kind of association which qualities of sensible colour have with the visual realm, qualities of sensible hardness have with the tactual realm, qualities of sensible sound have with the auditory realm, and so on.

Not every way in which a ϕ -terminally perceived item appears to its percipient counts as part of its *sensible* appearance in the relevant sense. Ralph may well see the relevant ϕ -terminal object not only as something of a certain shape, size, and colouring, but also as an apple, or as part of the surface of an apple. But, if so, this further aspect of how the item appears to him does not count as an aspect of its sensible appearance. Likewise, I may feel the ϕ -terminal object of my tactual perception not just as a surface of a certain shape, hardness, texture, and temperature, but also as the surface of a bottle, and, again, this further aspect of how the item may appear to me would not count as an aspect of its sensible appearance. Setting out precisely what it is that makes sensible appearance distinctive is a complicated matter, as I have shown elsewhere;⁸ in particular, it cannot be done solely by reference to the kinds of qualitative element that feature in its content and their relationship to the relevant sense realm. For example, whiteness has the right sort of association with the visual realm to feature in the content of visual sensible appearance. But if someone familiar with its daytime appearance inspects a white flower in his garden at night, when it sensibly appears to him as something grey, it may still, in a certain sense, look white to him (the same sense in which a ripe strawberry may look sweet to someone, or a stagnant pool might look smelly), and this looking white is not an instance of sensible appearance. As we shall shortly see, one of the things that distinguishes sensible appearance from other

⁸ In *The Nature of Perception*, 44–51.

forms of perceptual appearance is its distinctive phenomenological character; but that is not something I can explain until after I have introduced and defined a further key concept. At present, all we need to note is that, in any instance of ϕ -terminal perceiving, the perceived item's sensible appearance is, in a distinctive way, the basic core of its perceptual appearance, which is integral to the perceptual relationship itself. A physical item cannot be ϕ -terminally perceived except under a sensible appearance. But it is possible for there to be instances of ϕ -terminal perceiving without any elements of non-sensible appearance, and, in any particular case of ϕ -terminal perceiving, any elements of non-sensible appearance could be stripped away without affecting perceptual contact with the relevant physical item.

For an item to sensibly appear to a subject in a certain way is for the subject to perceive that item in a certain experiential way. So, whenever a physical item is ϕ -terminally perceived, there is something in the content of the subject's perceptual experience that embodies the item's sensible appearance in its mental aspect—something which forms, as it were, the way in which the subject is sensibly appeared to. I shall refer to this element in the perceptual experience as *phenomenal content*. Phenomenal content need not form the total content of the experience. If Ralph sees the ϕ -terminal item as part of the surface of an apple, that, too, will be part of the content of his perceptual experience. But, in the way that the sensible appearance of a perceived item forms the basic core of its perceptual appearance, the phenomenal content of a perceptual experience forms the basic core of its experiential content. I have introduced this concept of phenomenal content in a context where we are focusing on the case of physically perceptive experience, where the content of the experience forms the experiential manner in which the relevant physical item is ϕ -terminally perceived. But, although this will be the main area of our concern, it will be convenient to allow the concept to apply, additionally, to the phenomenologically corresponding content of hallucination, where the experience has the subjective character of physical-item perception, without being physically perceptive.⁹ In other words, it will be convenient to think of phenomenal content as the way in which the subject is sensibly appeared to, *whether or not there is a physical item that is doing the appearing*. This, I should stress, leaves entirely open the issue of whether the psychological character of phenomenal content is the same in the two cases. And on this we can expect

⁹ I am here diverging from my understanding of this concept in *The Nature of Perception*, but the difference is only one of terminology, not one of substance.

the decompositionalist and the fundamentalist to take opposite positions. A decompositionalist claims that the psychological states that are fundamentally involved in perception are not in themselves physically perceptive, and so is committed to saying that any experiential content that is fundamentally involved in perception could also occur in hallucination. A fundamentalist, in contrast, holds that perceptive and hallucinatory experiences differ in their fundamental psychological character, and, since the phenomenal content of a perceptive experience forms the experiential manner in which the ϕ -terminal contact is achieved, it is almost inevitable that he will take the psychological difference between the perceptive and hallucinatory cases to apply, in particular, to the nature of their phenomenal content.

Any adequate theory of perception has to provide an adequate account of the nature of phenomenal content, as it occurs in perception, and of its intimate involvement in the obtaining of the perceptual relationship between the subject and what he ϕ -terminally perceives. It is the question of whether the fundamentalist can provide such an account that I now want to pursue. Is there an acceptable way of understanding what, in the perceptive case, phenomenal content is, and how it is involved in ϕ -terminal contact, when, in any case of such contact, the perceptual relationship between the subject and the relevant physical item is taken to be psychologically fundamental?

The fundamentalist's simplest approach would be to adopt what I shall call the *presentational view*. Let us say that an item x is *presented to* a subject S , or that S is *presentationally aware of* x , if and only if x is psychologically related to S in a way that satisfies three conditions. First, the relationship is such as to make x available for demonstrative identification by S . In other words, it brings x before S 's mind in a way that allows him to pick it out as 'this item' (of which he is now conscious, and on to which he directs his attention), or, at least, in a way that would allow him to do this if he had the conceptual resources needed for demonstrative thought. Second, the relationship is such as to display, where appropriate in a certain perspective, certain aspects of x 's character, or character and location, in a way that makes them immediately available for cognitive scrutiny—though, once again, S 's capacity to take advantage of this availability depends on his having the requisite conceptual resources. Third, and most crucially, the relationship is wholly non-representational. It does not involve the use of concepts, symbols, or mental images, as a psychological means of registering x 's presence or the relevant (displayed) aspects of its character (character and location). Rather, the item and, in their concretely realized form, these aspects are before S 's

mind in a mode of absolute ontological immediacy, forming, in their own person, the very content of his awareness. What the presentational view then claims is that ϕ -terminal perceiving is a relationship that meets these three conditions. More precisely, it claims that, whenever a subject ϕ -terminally perceives a physical item, this item is presented to him in the sense that these conditions define, and the phenomenal content of the perceptual experience draws its qualitative ingredients from the physical features that are thus displayed. What particularly needs to be emphasized, here, is that the drawing of these qualitative ingredients is an *ontological*, not just a *causal*, matter. It is not just that the presence of an item with a certain sensible character causes the subject to have a perceptual experience with a matching content. It is that the ingredients of the content are themselves the very elements of the external situation made experientially present. The featuring of a quality in the phenomenal content is not something ontologically separate from its external realization in the perceived item—something that merely serves to *represent* that realization—but is that realization itself brought immediately before the mind, without the mediating role of concepts, symbols, or mental images.

Adopting the presentational view would be the fundamentalist's simplest approach. It is both the simplest way of representing the perceptual relationship involved in ϕ -terminal perceiving as something psychologically fundamental, and the simplest way of explaining, within that fundamentalist framework, how the phenomenal content of the perceptual experience forms the experiential manner in which perceptual contact is achieved. There is also something else that might make this view initially attractive to him. For it captures, in the most straightforward way, the phenomenological feel of perception. Thus, when a subject ϕ -terminally perceives an item, his experience carries the subjective impression of being presentational in the relevant sense. It gives him the impression that the perceived item and certain aspects of its character (or character and location) are presentationally before his mind in the relevant mode of ontological immediacy, with these displayed aspects exactly covering the content of how the item sensibly appears. It is this that gives sensible appearance the distinctive phenomenological character I referred to earlier; for other forms of perceptual appearance do not carry the subjective impression that the property or character that the perceived item appears to possess is, in this way, on presentational display. The point I am now making is that, to take ϕ -terminal perceiving to be genuinely presentational, and to think of the phenomenal content involved as drawing its qualitative ingredients from what is presentationally displayed, would

be the fundamentalist's most straightforward way of doing justice to the phenomenological facts.

The presentational view of perception is the simplest version of the fundamentalist view, and is what the phenomenological facts immediately suggest. But it is also, at least when offered as a general theory, open to an obvious and decisive objection. For it cannot accommodate cases of non-veridical perception. I am not thinking, here, of cases of hallucination, where no physical item is perceived: the presentational view does not even purport to offer an account of these cases. I am thinking, rather, of cases in which a physical item is ϕ -terminally perceived, but under a sensible appearance that misrepresents its true character. That such cases occur can hardly be denied—at least on the assumption that we perceive physical items at all. The notorious case of the stick partially immersed in water (in reality straight, but appearing bent) is an obvious example—though if it were thought (surely implausibly) that the presentationalist could handle this case either by assimilating it to the case of seeing veridically but in a special perspective, or by claiming that what is ϕ -terminally perceived is not the relevant time slice of the surface portion of the stick, but the light array it transmits to the subject's eye, then we could switch our attention to such phenomena as astigmatism and colour-blindness, where the distorting physical factors lie within the subject's own visual system. In whatever form they arise, the presentational view cannot accommodate cases of non-veridical perception because in taking the qualitative ingredients of phenomenal content to be directly drawn from the relevant physical item, it excludes the possibility of the sensible appearance of this item being at variance with its true character. Sensible appearance just is, for the presentationalist, the direct bringing of certain aspects of the item's actual character, sometimes in the relevant perspective, before the subject's mind.

The presentational view cannot deal with the phenomenal content that occurs in cases of non-veridical perception, and so cannot be accepted in the general form in which I have formulated it. But this still leaves the option of retaining the presentationalist approach for cases of veridical perception. So, in the case where a straight stick in water looks bent, we are forced to say that the featuring of bentness in the phenomenal content is not the featuring of some physical instance of bentness. But in the case where a straight stick out of water looks (veridically) straight, we could still say that the featuring of straightness in the phenomenal content consists in the instance of straightness in the stick being made present to the mind. But

while this mixture of approaches is an option, it is hardly a plausible one. For it is very hard to suppose, in the sorts of case that we are considering, that the veridicality or non-veridicality of an experience correlates with such a fundamental difference in its nature. We are not now, as in the earlier context, concerned with the distinction between perception and hallucination, where distinguishing the psychological nature of the experiences involved is a serious option. In the present context the veridical and non-veridical experiences we are envisaging are alike in being physically perceptive; and, crucially, they are alike in being perceptive of the concrete physical feature whose sensible appearance is at issue (so that, in the case of the stick in water, the physical instance of straightness is still perceived, though non-veridically). Moreover, they causally originate from the perceived physical item by processes of a broadly similar kind. All this puts us under strong pressure to think of the experiences as amenable to a unitary account. This pressure becomes, to my mind, irresistible when we focus on a case where a shift from veridical to non-veridical perception involves only very slight changes to the qualitative character of the phenomenal content and to the details of the causal process from the relevant physical item. Think, for example, of a situation in which someone first looks at an object through plain flat glass, seeing its shape as it is, and then looks at it through glass whose very slight degree of curvature imposes a correspondingly slight distortion on the way that the shape of the object appears. It would surely be absurd to deny that these two perceptive experiences—and experiences that are alike in being perceptive of the object's shape—are, in their intrinsic character, of the same generic type.¹⁰

Granted that he needs a unitary account of veridical and non-veridical perception, the fundamentalist is obliged to conclude that, in all cases of perception, the qualitative ingredients of phenomenal content are internal to (creatures of) the mind, rather than ontologically drawn from the physical items perceived. Thus, if someone ϕ -terminally perceives a physical item x , and if, in the context of that perceiving, x sensibly appears to him as being F , then the fundamentalist is obliged to say that, even if x is in fact F , and even if its being so is causally responsible for this aspect of its appearance, the featuring of F -ness in the context of this appearance is not the featuring of the x -instance of F -ness (or of any other physical instance of F -ness), but is merely something in the subject's mind that in some way represents that instance. I shall speak of this position as the *internalist*

¹⁰ I elaborate this argument in detail in *The Nature of Perception*, 67–72.

view. I should stress that I am here using this label to refer exclusively to a version of the fundamentalist view. Decompositionalists, too, of course, will regard the ingredients of phenomenal content as internal to the mind; the decompositional view allows for no other possibility. But, in the sense I here intend, the internalist view combines the claim of internality with the acceptance that where there is ϕ -terminal perceptual contact with a physical item, the fact of this contact is something psychologically fundamental. This has to be borne in mind when interpreting the claim of internality itself. Whether a straight stick looks straight or bent, the featuring of this latter shape in the content of the appearance is not, on the internalist view, a featuring of the instance of straightness in the stick. But the internalist will still accept that the physical instance of straightness is, in the context of this appearance, perceived; and, as a fundamentalist, he will insist that its perceptual presence (whether it is perceived veridically or non-veridically) is an aspect of what is psychologically fundamental.

By adopting this internalist view, the fundamentalist avoids the problem that defeated the presentationalist: since the ingredients of phenomenal content are not ontologically drawn from the perceived item, there is no difficulty in understanding how phenomenal content can be at variance with the item's true character. But he now faces problems of a different kind.

The basic problem, as I see it, is that, on the internalist view, the fundamentalist cannot make sense of the way in which perceptual contact and phenomenal content fit together—the way in which phenomenal content embodies the sensible appearance under which the ϕ -terminal object is perceived, and forms the experiential manner in which ϕ -terminal contact is achieved. There is no difficulty, in this respect, for the presentationalist. As he sees it, phenomenal content is precisely what ϕ -terminal perceptual contact automatically supplies by virtue of its presentational character—by the way in which it directly brings before the subject's mind certain aspects of the character and situation of the perceived physical item. There is no difficulty in understanding how these displayed aspects form the qualitative ingredients of the item's sensible appearance to the subject, and so of the experiential manner in which he perceives it. Likewise, there is no difficulty, on this front, for the decompositionalist. He takes perceptual contact to be something that breaks down into the subject's being in a more fundamental (not in itself physically perceptive) psychological state, together with certain additional factors, and, in the case of ϕ -terminal contact, he will take the relevant psychological state to be that of having a perceptual experience with

a certain kind of phenomenal content. Phenomenal content will then form the experiential manner in which ϕ -terminal contact is achieved, and embody the sensible appearance of the perceived item, by being that by which, in the context of the other factors, contact is constitutively secured. The difficulty is in seeing what other option is available—in thinking of some way of making sense of the intimate link between contact and content without invoking either a presentational or a decompositional account. In effect, the problem I see with the internalist view is that its two components pull in opposite directions. Once we take the ϕ -terminal relationship to be something psychologically fundamental, I can see no way of fitting phenomenal content into the picture except as something that draws its qualitative ingredients from aspects of the perceived physical situation. Once we think of the ingredients of content as internal to the mind, and ontologically distinct from their physical counterparts, I can see no way of making provision for perceptual contact except in a form to which content constitutively contributes. If, like the internalist, we think of phenomenal content as neither drawing its ingredients from aspects of the physical situation nor as constitutively contributing to the securing of perceptual contact, we are, as I see it, left with no way of understanding how it could be anything other than an experiential accompaniment of such contact, rather than, as it has to be, the manner of its achievement.

One specific way in which this basic *Us* problem manifests itself is with respect to what I shall speak of as the *appropriateness requirement*. What I mean by this is the fact that, in order for a perceptual experience to be a ϕ -terminal perception of some physical item, its phenomenal content has to be, to an adequate degree, qualitatively appropriate to that item, relative to the conditions of observation. The best way to see this is to focus on a case where all the other conditions associated with physical-item perception are present, but the factor of appropriateness is conspicuously absent. Thus, suppose I am in my sitting room, with my eyes turned towards the clock on the mantelpiece, with nothing obstructing my line of vision, and with all the other external factors favouring the achievement of visual contact. And suppose that light reflected from the clock and its surroundings enters my eyes in the normal way and sets up the appropriate kind of process in my optic nerves, which in turn transmit the appropriate signals to my brain. But then something peculiar happens. My brain responds to the incoming signals in a totally bizarre way, producing a visual experience which is not remotely like the sort of experience that is normal for that kind of photic

input. For instance, it might be that the resulting experience is like that of seeing a football match, or like that of seeing a clear blue sky. Now it is surely clear that, given the extent of the disparity between the real character of the external environment and the phenomenal content of my experience, this experience is not physically perceptive. It is true that the clock and its surroundings play a causal role in producing the experience, and, with respect to the photic input, this role is of the normal kind for the circumstances in question. And we can even suppose that, as in the case of normal visual perception, the brain response preserves a kind of causal isomorphism between elements of the resulting experience and elements of the input, so that, relative to a suitably fine-grained division, different elements in the content of the experience causally trace back to different elements of the relevant portion of the environment. But it would be absurd to suppose that the experience qualifies as an actual seeing of this portion, and that the only way in which its deviant content affects the situation is in making this seeing radically non-veridical. It is just obvious that, in the context of the conditions envisaged, the extent of the non-veridicality precludes visual contact altogether. So, here we have a clear illustration of the point at issue, that there can only be ϕ -terminal contact in cases where the phenomenal content has an adequate degree of qualitative appropriateness to the relevant physical item, relative to the conditions of observation; and it is easy to think of a host of other examples that would illustrate the point in an analogous way. In this kind of example, the failure of the experience to achieve the requisite degree of appropriateness is revealed by its radical non-veridicality. But we should not think of appropriateness as entirely a matter of veridicality. It would, of course, be entirely a matter of veridicality from the standpoint of the presentational view, which excludes non-veridical perception altogether. But now that we have rejected that view—and rejected it precisely because it cannot accommodate non-veridical perception—we must also accept that appropriateness is partly a matter of conformity to what is normal, or normative, for the conditions of observation in question. For example, if a straight stick is partially immersed at an angle in water, the appropriate way for it to sensibly appear to a visual percipient is as bent, and it would be someone who could only see it as straight whose vision would be defective.

It is undeniable that the appropriateness requirement holds. But it creates difficulties for the internalist view in two ways, and ones which a defender of the view cannot, I think, adequately deal with.

In the first place, the internalist does not, it seems to me, have any adequate way of accounting for it. Considering the issue of explanation in the abstract, we can identify two clear-cut ways in which the requirement might be explained. On the one hand, there is the explanation that would be offered by the decompositionalist. This would be to say that a sufficient degree of appropriateness is an essential constitutive factor in the securing of ϕ -terminal contact. Thus, the decompositionalist holds that, wherever there is ϕ -terminal contact, it is constituted by the combination of the subject's having a perceptual experience of a certain phenomenal kind and certain additional facts, and so he will account for the appropriateness requirement by insisting that, for such a constitution to succeed, the presence of an adequate degree of appropriateness is one of the things that the additional facts have to cover. On the other hand, there is the explanation that would be offered by the presentationalist. On his account of perception, the qualitative ingredients of phenomenal content are directly drawn from aspects of the physical environment. There is no room, here, for any degree of inappropriateness (which, for the presentationalist, would be the same as non-veridicality), since it is only insofar as there is an accurate display of the environmental situation that there is phenomenal content at all. Both these accounts of perception would, in their contrasting ways, provide a complete rationale for the appropriateness requirement. But neither of them is available to the internalist, who combines a fundamentalist view of contact with an internalist understanding of the ingredients of content. Nor, it seems to me, can he derive a rationale from any other source. Once it has been accepted that the qualitative ingredients of phenomenal content are internal to the mind, and not directly drawn from aspects of the environment, the only way I can see of making sense of there being a limit on the amount of inappropriateness that ϕ -terminal perception can tolerate is by supposing that a sufficient degree of appropriateness is an essential element in the factors by which ϕ -terminal contact is constituted.

The second point involves something more subtle. Although there is a limit on the degree of inappropriateness that ϕ -terminal perception can tolerate, our concept of it is, surely, not so precise as to ensure that, in any particular case, there will always be an objective answer to the question of whether the requirement is satisfied. As a result, we can envisage cases whose status, as perceptive or non-perceptive, is inherently borderline. Thus, suppose scientists have constructed a device that can be used to distort the visual appearance of the physical scene by sending a stream of radiation

through the subject's visual cortex, the amount of the distortion increasing with the strength of the radiation. And suppose an experimenter is about to use this device on someone who is looking at an apple. At one extreme, with very weak radiation, we can envisage the experimenter producing an effect on phenomenal content so slight that there would be no threat to the continuation of visual contact: the subject would continue to see the apple (at each moment ϕ -terminally seeing a certain momentary stage of a certain portion of its surface), but its apparent shape would be a little warped or its surface colour-pattern look blurred. At the other extreme, with very strong radiation, we can envisage an effect so great that visual contact would clearly be severed: how things appear to the subject would bear no resemblance at all to how things are, and the experience could not, by any stretch of the imagination, be construed as perceptive. But between these extremes we can also, surely, envisage a range of cases that would be inherently borderline—cases whose classification as perceptive or non-perceptive would be a matter for stipulation, rather than an issue of objective fact. To reach such cases we need only envisage a series, from the first extreme to the second, in which the experimenter very gradually increases the strength of the radiation and the resulting degree of the effect on phenomenal content. It is surely clear that, somewhere in the middle, cases would occur where the question of whether the extent of the inappropriateness was enough to sever visual contact with the apple would have no definite answer, even from a God's-eye view, and where the status of the relevant experience would be left indeterminate.

The possibility of these borderline cases is easy enough to explain in the framework of the decompositional view, which takes a subject's ϕ -terminal perceiving of a physical item to be partly constituted by the fact that the phenomenal content of his experience stands in the right sort of qualitative relationship to it. For if a sufficient degree of appropriateness is a constitutive requirement for the obtaining of ϕ -terminal contact, borderline cases could arise in this area in the way that they are liable to arise in any area where the application of a concept constitutively depends on whether the situation achieves a sufficient value along a certain qualitative dimension, but where there is no particular point on this dimension that marks a theoretically crucial division. So, it would be as easy to understand why there might be no objective answer to the question of whether a certain experience has enough appropriateness to count as perceptive as it is to understand why there is sometimes no objective answer to the question of whether

some group of people is sufficiently numerous to count as a crowd, or to the question of whether someone has enough hair on his head to avoid counting as bald. But the situation for the internalist is quite different. Even if he could find some rationale for the appropriateness requirement itself, his commitment to the fundamentalist view would prevent him from even acknowledging the possibility of borderline cases of the sort envisaged. After all, perceptual contact itself, unlike qualitative appropriateness, does not admit of degrees: it is all or nothing. So, if such contact, in its ϕ -terminal form, is taken to be psychologically fundamental—something which does not, at the psychological level, break down into further factors—there is nothing at the psychological level of description which could explain how the question of its obtaining could ever fail to have an objective answer. It is only if the obtaining of contact is constitutively controlled by more fundamental factors, and if the conditions for its obtaining are not fully precise in terms of those factors, that there is an opportunity for borderline cases to occur.

It seems to me, then, that the internalist view does not allow the development of an adequate account of the relationship between perceptual contact and phenomenal content. It does not, as I see it, allow any adequate understanding of how phenomenal content forms the experiential manner in which ϕ -terminal contact is achieved. Nor does it allow any explanation of why, in the case of such contact, there is a limit on the degree to which phenomenal content can be qualitatively inappropriate to the item perceived, relative to the conditions of observation. Nor can it accommodate the fact that, because the requirements for sufficient appropriateness are not sharply defined, we have to allow for cases where there is no objective answer to the question of whether there is perceptual contact or not. In all these ways, as I see it, the internalist view shows itself to be unsatisfactory.

The internalist view is one version of the fundamentalist view of perception, which holds the ϕ -terminal perceptual relationship to be psychologically fundamental; it is that version which takes the qualitative ingredients of phenomenal content to be internal to the mind. The alternative version is that of the presentational view, which takes ϕ -terminal perceiving to be presentational, and takes the ingredients of content to be ontologically drawn from the external physical situation. Since the presentational view has also shown itself to be unsatisfactory, and unsatisfactory for cases of veridical and non-veridical perception alike, and since there are no other fundamentalist

options available, I conclude that the fundamentalist view is mistaken, and mistaken in its application to all cases of perception.

V

The alternative to the fundamentalist view of perception is the decompositional view. This holds that perceptual contact with a physical item, even when it is ϕ -terminal, breaks down into two components, one of which consists in the subject's being in some further psychological state, which is not in itself physically perceptive, and the other of which comprises certain additional facts, but ones that do not add anything further to the subject's psychological condition at the relevant time. Having found fault with the fundamentalist view, we now need to consider whether this alternative approach fares better. Since there is no other approach to perception now available, it will be awkward if it does not.

One question that arises with respect to this view is over the nature of the relevant psychological states. We are told that they are not in themselves physically perceptive; so they are capable of occurrence in both the context of physical-item perception and the context of hallucination. We also know that the states are experiential—states of visual, auditory, tactual, or some other kind of sense-perceptual experience—and that the central component of their experiential content—the component we have called phenomenal content—covers the way in which things sensibly appear to the subject. But we still do not know exactly what psychological character these states are supposed to have; we do not know how best to understand their character from a decompositional standpoint. This is an issue that I have discussed at considerable length in *The Nature of Perception* (the whole of the third part is devoted to it), and I still accept the conclusions I reached there and the arguments I invoked to establish them. But since the issue is only peripheral to my present concerns, and its outcome does not affect the argument I want to develop, I shall only briefly touch on it here.

Deciding on the psychological nature of the relevant experiential states has been a source of a great deal of controversy within the decompositional camp. Until about the middle of the last century the standard approach, at least within the empiricist tradition, was to say that, although these states are not in themselves *physically* perceptive, they do involve, as their central component, the presentational awareness of a *non-physical* item—an item

which is characterized by sensible qualities associated with the sense realm in question, but whose existence is internal to the mind. So, in the case of Ralph, the claim would be that, in seeing the apple, the central component of his visual experience, at any time, consists in his presentational awareness of a certain spatial array of colours, though an array which, unlike the relevant physical item, has no existence outside the realm of experience. These putative mental objects of awareness, often labelled by earlier philosophers as 'ideas' or 'impressions', came to be known in the twentieth century as *sense data*, and the position which postulates them is now usually referred to as the *sense-datum theory*. In representing this theory as claiming that the awareness of a sense datum forms only the *central component* of the relevant experiential states, I am going beyond the explicit requirements of the tradition. But it seems to me that the only form of sense-datum theory worth considering is one which takes each total perceptual experience to combine the presentational awareness of a sense datum with an element of interpretation. The reason why the theory needs to recognize this element of experiential interpretation is that, without it, it could not do justice to the phenomenological character of perceptual experience. For such experience not only makes it introspectively seem to the subject that there is a sensible item before his mind, but also makes it seem to him that this item is an ingredient of his physical environment; the apparent environmental character of what is before the mind is, indeed, an essential aspect of phenomenal content. The only way in which the sense-datum theorist can account for this phenomenological situation is by taking perceptual experience to include an element of interpretation that is directed on to the presented sense datum and which represents it as environmental in the appropriate way. So, in the case of the visual realm, he must think of the element of experiential interpretation as representing the two-dimensional colour array which forms the presented sense datum as a three-dimensional arrangement of colours located at various distances and in various directions from the subject.

In recent years the sense-datum theory has become increasingly unpopular, partly because its ontology of mental objects is seen as inherently problematic, and partly because the acceptance of this ontology is incompatible with the physicalistic approach to human mentality that modern philosophers tend to favour. As a result, most current decompositionists try to understand the nature of perceptual experience without recourse to sense data. One approach, here, is to represent the relevant experiential states as purely conceptual in character. For example, it might be claimed that perceptual experience

is simply the conscious acquisition, in a suitably vivid form, of putative information about the external environment.¹¹ Another approach is to accept that the central component of a perceptual experience is something akin to a sense datum, in that it involves the experiential realization of certain sensible qualities, but insist that these qualities are realized not as features of an *object* of awareness, but as aspects of a *manner* of awareness.¹² The main problem for any alternative approach is in accommodating the phenomenology of perception. I have already stressed that the sense-datum theorist needs to recognize an element of experiential interpretation in order to explain why the presented sense datum poses as an *environmental* object of awareness. These alternative approaches face the more basic challenge of explaining why perceptual experience carries the subjective impression of being an awareness of something at all. If such experience is not inherently perceptive of anything physical—which it cannot be on the decompositionalist view—and if it does not involve an ontology of mental objects of awareness either, why should it introspectively seem to the subject—as it undeniably does—that there is a sensible item before his mind?

In my view, this challenge cannot be satisfactorily met, and the sense-datum approach is the right one for the decompositionalist to adopt. Admittedly, I think that to make this approach intelligible, we need to revise its traditional form in one crucial respect. Traditionally sense data are taken to be private and momentary particulars: each sense datum occurs as the object of a *single* episode of awareness, in the mind of a *single* subject at a *single* time, and its existence is logically restricted to this unique mind–time location. But it seems to me that, to give sense data the right kind of ontological status to serve as genuine objects of awareness, we need to reconstrue them as sensory universals, capable of presentational occurrence to different subjects and on different occasions. The reasons for this, and the more general reasons for thinking that we need some kind of sense-datum theory at all, are matters that I have dealt with in detail in the earlier book,¹³ and I shall not pursue them any further here. They are in any case, as I have indicated, peripheral

¹¹ This is the approach pioneered by David Armstrong and George Pitcher. Thus, see Armstrong's *A Materialist Theory of the Mind* (London: Routledge & Kegan Paul, 1968), ch. 10, and Pitcher's *A Theory of Perception* (Princeton, NJ: Princeton University Press, 1971).

¹² This is the *adverbialist* approach. Its early advocates include C. Ducasse, 'Moore's refutation of idealism', in P. Schilpp (ed.), *The Philosophy of G. E. Moore* (Chicago, Ill.: Northwestern University Press, 1942), 223–51, and R. Chisholm, *Perceiving* (Ithaca, NY: Cornell University Press, 1957), 115–25.

¹³ *The Nature of Perception*, 93–195.

to my present concerns. The issue on which I want to focus, and to which I now turn, is to do with a different aspect of the decompositional view, and the conclusions for which I shall argue do not depend on what account of perceptual experience the decompositionalist adopts.

According to the decompositionalist, perceptual contact with a physical item, even when ϕ -terminal, is secured by the combination of the subject's being in a certain experiential psychological state—a state which is not in itself physically perceptive—and certain additional facts, and it is stipulated that these additional facts do not involve anything further about the subject's psychological condition at the time in question. But if these additional facts do not add anything to the subject's psychological condition at the relevant time, and if the relevant experiential aspect of this condition does not in itself put the subject into perceptual contact with any physical item, it may well be wondered how such contact gets secured at all. Perceptual contact, after all, is an awareness relation: to perceive an item is to be perceptually aware of it. How can an experiential state that is not in itself physically perceptive and factors that contribute nothing further to the subject's psychological condition at the relevant time add up to a genuine awareness of a physical item? How can the subject's awareness reach beyond the boundaries of his mind if there is nothing in his fundamental psychological condition that effects this? It is here, I think, that we encounter a fundamental problem for the decompositional view, and one which cannot be eliminated within the framework of assumptions on which we are currently working.

To bring out the problem, let us consider, once again, the case of Ralph and the apple, and, for simplicity, let us initially assume that, apart from the perceptual experience itself, the decompositionalist restricts the factors that he sees as contributing to the securing of ϕ -terminal perceptual contact to the qualitative and causal relations of the experience to the relevant physical item. Then, from the decompositionalist's standpoint, we can take the basic facts of the case to be these: (i) Ralph's eyes are open and turned towards the apple on the table in front of him; (ii) light reflected from a certain portion of the apple's surface enters his eyes and initiates a neural process that terminates in a visual experience, and all this happens in accordance with the normal mode of functioning of the human visual system; (iii) this experience is not, in itself, perceptive of any physical item, but its content, and, in particular, its phenomenal content, is such as to give Ralph the impression, and thereby invite him to believe, that he sees, at a certain distance in front of him, a roughly hemispherical patch of a certain size, shape, and colouring; and

(iv) the shape, size, and distance of the relevant surface portion of the apple (or, more precisely, the relevant momentary stage of this portion) more or less coincide with those of the notional item that Ralph has the impression of seeing, and the pattern of pigments on this portion (portion stage) more or less matches, at least in spatial character and qualitative structure, the pattern of sensible colours of this notional item.¹⁴ Now the decompositionist wants to say that, in combination, the occurrence of the visual experience, the qualitative fit between its content and the relevant physical item, and the causal role of this item in producing the experience suffice to make it true that Ralph sees the item. But, given our actual concept of seeing, it is hard to understand what could warrant this conclusion. The facts just listed certainly indicate that there is a significant link between the experience and the external item: there is a sense in which, by causing the experience, and by giving it the sort of content that it does, the item conveys its presence to the subject. But it is hard to think of the item as thereby becoming something which the subject genuinely sees—something of which he is genuinely visually aware. What we have here, surely, is not an awareness of an external item, but only an experience which, by its content and its causal origins, provides its subject with some kind of representation of, or with information about, an external item. Nor would things turn out differently if we were to represent the decompositionist as recognizing further facts as constitutively relevant to the securing of visual contact. Granted that the totality of facts that he recognizes as relevant divides into the realization of a psychological state that is not in itself physically perceptible and facts that involve nothing further about the subject's psychological condition at the relevant time, the basic problem will always remain.

The point I am making here is a more general version of something that is already familiar in the case where the decompositionist view assumes its sense-datum form—the form in which the relevant experiential state is taken to involve, as its central component, the occurrence of a mental object of awareness. For it is commonly acknowledged that if what is fundamentally before the subject's mind is something which (or which with respect to that particular presentational occurrence¹⁵) only exists in his mind, and if it only seems to him that he is perceiving something external (environmental)

¹⁴ The reason why I say 'at least in spatial character and qualitative structure' is to leave room for a Lockean account of colour. This account, and the Lockean account of the secondary qualities in general, will be discussed in the next chapter.

¹⁵ This is to allow for the universalist construal of sense data which, as I explained, I favour.

because his experience interprets this item in an externalist way, then the subject's impression of perceiving something external is simply mistaken, and his awareness does not reach to anything beyond the mental item itself. The point, in this special case, is sometimes figuratively expressed by saying that the internal sensible items create a kind of 'veil of perception'—as if their presence before the mind forms an obstructive screen, which blocks perceptual access to the external realm that lies on its further side. But, despite its charm, this image of a veil is misleading. The reason why the sense-datum theory does not allow for perceptual access to the external world is not that sense-data are positively obstructive to such access, but that there is simply nothing in the fundamental situation envisaged by the theory that would suffice to create it. It is in this respect that the problem for the sense-datum version of the decompositional view is only a special instance of a general problem that affects all versions. The basic point, in the case of all versions, is that there is no provision for a genuine awareness of anything external if the relevant experiential state of the subject is not in itself physically (externally) perceptive and if the only additional relevant factors involve nothing further about the subject's psychological condition at the relevant time. The only way in which the adoption of the sense-datum theory affects the situation is by making this point particularly conspicuous—by making it obvious, given the presence of the internal presented items, that the subject's awareness does not reach beyond the boundaries of his own mind.

It might be wondered whether I am working with too narrow a view of what is required for perceptual awareness. In our ordinary thinking we seem to recognize various types of case in which the perceiving of one physical item is in some way mediated by the perceiving of another. For example, we are happy to accept that someone can watch a football match on television, when we know that his visual access to the match is channelled through his access to the patterns on the screen. Likewise, we are happy to speak of a radar operator as seeing the approach of a missile, when we know that his only way of detecting it is by seeing and correctly interpreting some signal on his monitor. Or again, we think nothing of saying that we can hear the approach of an ambulance when the only indicator of this is the sound of its siren. Even with respect to cases, like that of Ralph and the apple, which we take to be paradigmatically perceptive, we recognize that the subject's contact with the object in question is achieved through his contact with one of its parts, or with some momentary stage of a part. All these seem to be cases where we ordinarily think of the subject's perceptual contact with one

physical item as mediated by his perceptual contact with another; and, as I have defined it, perceptual mediation is decompositional: the contact with one item breaks down into contact with the other and certain additional (not relevantly perceptual) facts. But if our ordinary concept of perceiving allows for this kind of decompositional mediation, could we not, after all, see it as permitting the situation envisaged by the decompositionalist? Why should perceptual awareness be able to tolerate the one sort of decomposition but not the other?

This is a fair question. But I think that we only need to examine the cases of supposed perceptual mediation in more detail to see that there is nothing here that could be of any assistance to the decompositionalist. In each case, what we find is that there is no proper analogy between how the case is to be ultimately understood and what is envisaged by the decompositional view.

Take, first, the case of someone following a football match on television. There is no denying that we ordinarily think of such a subject as able to see events on the football pitch, and we also recognize that his visual access to these events is in some way channelled through his visual access to what takes place on the screen. But in order for this to help the cause of the decompositionalist, the channelling has, in line with my definition of mediation, to be of a decompositional kind: it has to be such that whenever the viewer makes perceptual contact with events on the pitch, the fact of this contact breaks down into (is constituted by) the fact of his contact with events on the screen and certain other (not relevantly perceptual) facts. It is here that things start to go wrong. It is true that there is a way of representing the situation in such decompositional terms. For we could claim, and perhaps with some plausibility, that what is ultimately going on psychologically is that the subject visually registers patterns on the screen, but experientially interprets them as scenes from the match. And once this claim is accepted, it will be hard to deny that such contact as the subject has with the match ultimately breaks down into this registering and experiential interpreting, together with relevant facts about the causal process from the stadium to the television. But the trouble with this, from the standpoint of the decompositionalist, is that, once we have represented the situation in such decompositional terms, it is no longer plausible to suppose that the subject's contact with the match is genuinely perceptual. It is no more plausible to suppose that the subject's awareness reaches beyond the patterns on the screen than it was plausible to suppose, in the context of the sense-datum theory, that a subject's awareness reaches beyond the presented items in his mind.

The only way in which we can plausibly think of the viewer as genuinely seeing the match is by taking his visual contact with it to be something psychologically fundamental, and construing the role of the patterns on the screen, and his access to them, as merely causal—as consisting in the fact that his reception of light from the screen is a crucial part of the causal process by which events on the pitch become visible to him. This would accord with how things phenomenologically seem to the viewer himself; and, because of our first-person familiarity with televisual experience, it is how we tend to understand the situation in our ordinary thinking. But, thus construed, the television case would not provide any kind of analogy with what is postulated by the decompositionalist.

The case of the radar operator is equally of no help to the decompositionalist. There is no denying that such contact as the operator has with the approaching missile breaks down into his perception and recognition of the signal, and the nature of the causal link between the signal and the missile itself. And in this case, unlike that of the television viewer, we are not, even in our ordinary thinking, inclined to understand the situation in any other way, since there is nothing in the phenomenology of the radar-monitoring experience that might tempt us to a different conclusion. But the trouble, once again, is that, once we take account of the decompositionist nature of the situation, we are prevented from regarding the contact with the missile as genuinely perceptual. Indeed, even in our ordinary thinking we recognize it as obvious that the operator does not really see the missile, but merely detects its presence by inference from the signal. If in our ordinary usage we are happy to speak of him as *seeing* the missile, this is only because ordinary usage does not always aim to describe things as they strictly are. All this also applies, *mutatis mutandis*, to the case of the ambulance and the siren.

There remain the cases of perceiving a whole object by perceiving a part, and of perceiving a persisting object by perceiving a momentary stage; and at first sight these may seem to be the ideal types of case for the defender of the decompositionist view. On the one hand, there is no disputing the claim that any perceptual contact with the whole or with the persisting item is decompositionally mediated by perceptual contact with the part or the stage. Thus, in the case of Ralph, there is no denying that, at any given time, whatever visual contact he has with the whole apple breaks down into his contact with a certain portion of its surface, together with the fact that this latter item *is* a portion of its surface; and whatever visual contact he has with this persisting surface portion breaks down into his contact with one of its

momentary stages, together with the fact that this latter item *is* a stage of that portion. On the other hand, the decompositional nature of these cases does not make it difficult to accept that the contact with the whole and the persisting items is genuinely perceptual. There is no temptation to say that, because all that is fundamentally visible to Ralph at time *t* is portion stage *O*, he does not really see the persisting portion of which *O* is a stage or the apple of which this portion is a part. On the face of it, then, we have here exactly the right sorts of case for the decompositionalist's purposes—cases in which, by all ordinary standards, there is both genuine perception and decompositional mediation. But on further reflection it becomes clear that these cases could not be less helpful to his cause. For what here allows us to recognize the combination of genuine perception and decompositional mediation is that, in each case, the two putative perceptual objects involved are not, as in the other cases we considered, ontologically separate, and so there is no sense in which the subject's awareness has to get *beyond* the object that features in the mediating perception to reach the object that features in the mediated. Contact with the momentary stage just is contact with the persisting item in a temporally focused form; contact with the part just is contact with the whole item in a mereologically focused form. This cannot provide any analogical support for the decompositional view, where the whole issue turns on how we can think of the subject's awareness as reaching beyond the boundaries of his own mind. Even when that view is developed along its traditional sense-datum lines, where the relevant experiential states involve the occurrence of objects of awareness, these objects are located in the mind, not in the external world, and so are, on a grand scale, ontologically separate from the physical items supposedly perceived.

I have considered three types of case where, in our ordinary thinking, we either recognize, or seem to recognize, instances of perceptual mediation, and which might suggest the possibility that we can, after all, make sense of the situation envisaged by the decompositionalist, and we have seen that none of them affords the decompositionalist the analogical support that he needs. In the case of the television viewer, the only way in which we can represent the situation as one in which the subject genuinely sees the football match is by denying that there is mediation in the relevant decompositional sense, and taking the role of the subject's access to the screen to be merely causal. In the cases of the radar operator and the ambulance, there is no denying that the subject's contact with the remote item (the missile or the ambulance) is decompositionally mediated by his

contact with something more immediate (the signal or the sound of the siren), but there is equally no question of thinking of the contact with the remote item as genuinely perceptual. In the cases of the whole and the part, and of the persisting object and the momentary stage, we must accept that there is both genuine perception and decompositional mediation—or at least we must accept this once we accept that the item which is supposedly more immediately perceived is genuinely perceived—but, unlike anything that might be envisaged under the decompositional view, the two putative perceptual objects are not ontologically separate, and so there is not the same problem in understanding how contact with the one automatically secures contact with the other. Although there are other types of case that we could consider, I cannot think of any that would not, from the standpoint of trying to help the decompositionalist, fail in one of these three ways, where either there is no decompositional mediation, or the supposed remote perceptual object is not genuinely perceived, or the two perceptual objects are not ontologically separate.

VI

We can now see why the topic of perception presents us with a problem. We have identified two rival general views about how physical-item perception should be understood. Thus, according to the decompositional view, whenever someone perceives a physical item, he is in some psychological state which is not in itself physically perceptive, and the fact of his perceiving that item breaks down into (is constituted by the combination of) his being in this state and certain additional facts, where these latter facts do not involve anything further about his psychological condition at the relevant time. According to the fundamentalist view, whenever someone ϕ -terminally perceives a physical item, the fact of his perceiving this item is something psychologically fundamental: it does not break down into his being in some further psychological state, which is not in itself physically perceptive, together with additional facts of the relevant kind; it does not, at the psychological level, break down into further facts at all. For any given case of ϕ -terminal perceiving, these two views are jointly exhaustive: the perceptual relationship must either psychologically decompose or be psychologically fundamental. Yet we have apparently found that neither view offers, for any case, an acceptable account. The fundamentalist view fails because it does not allow an

adequate understanding of the relationship between perceptual contact and phenomenal content. The decompositional view fails because once we accept that the subject's contact with the relevant physical item decomposes in the envisaged way, there is no way of understanding how that contact can qualify as genuinely perceptual. But to reject both views, while acknowledging that they are, case by case, exhaustive, would commit us to concluding that we do not have perceptual access to the physical world at all. The sheet of paper on which I am currently writing is not something that I genuinely see; the pen I am holding is not something that I genuinely feel; the laughter of the students in the quadrangle below my window is not something I genuinely hear. The only kind of sense-experiential contact that human subjects can have with items in the physical world is one which, because it psychologically decomposes, does not count as genuine perception.

It goes without saying that this would be a very difficult conclusion to accept, not just because it clashes with what we ordinarily take for granted, but also because, even at the level of philosophical reflection, denying the perceptual accessibility of the world would be deeply counter-intuitive. Indeed, a natural, and not unreasonable, response of someone who has read the discussion so far would be to suppose that, if it really yields that outcome, something must have gone wrong in the reasoning. Nor does the problem end there. For, as well as being counter-intuitive in itself, there are a number of further ways in which the denial of perceptual accessibility would pose a challenge, or *prima facie* challenge, to our ordinary understanding of the world and our place within it, making it even harder to settle for the conclusion that access to the world is unavailable. Three points of challenge, in particular, are worth mentioning, though I shall not pursue any of them in detail at this stage.

The first point concerns our ordinary conception of the physical world as forming *our* world—the world which provides the framework of our lives and supplies the subject matter of our everyday thought and discourse. Part of this conception turns on the fact that we ordinarily think of the physical world as the realm in which, through our embodiment, we are materially located, and the denial of perceptual access to this realm would not directly affect this aspect of our thinking. But another element in this conception, and surely one that is equally crucial, is that we take the physical world to be our world *experientially*: we think of it as a realm in which we are not only materially but *empirically* located; and, on the face of it, a central part of this thought is that the world is something within which we have a

perceptual viewpoint and whose contents become, in appropriate conditions, perceptually accessible to us in the perspective of that viewpoint. If we had to give up our belief in the perceptual accessibility of the physical world, it seems that our conception of the world as forming *our* world would, at best, survive in a significantly attenuated form.

The second point concerns our ordinary beliefs about the sorts of item that the physical world contains. We have a rich stock of sortal concepts by which we ordinarily classify the material objects that we believe to be found in the world—concepts such as *apple*, *tree*, *river*, *table*, *house*, *pen*, and so on. On the face of it, these concepts involve thinking of the objects to which they apply as disposed to sensibly appear to us in certain characteristic ways in certain conditions of sensory encounter. So, in classifying something as an *apple* or as a *pen*, it seems that, amongst other things, we are representing it as of a type whose instances are disposed to look and feel in certain characteristically apple-like or pen-like ways when we make the right sorts of sensory contact with them. And the same seems to be true of all the other familiar concepts of this kind. But, of course, things can only have a sensible appearance in so far as they are perceived: for an item to sensibly appear to a subject in a certain way is for him to perceive it in a certain phenomenal manner. So, if we accept that the concepts have these dispositional implications, but deny that we have perceptual access to the physical world, we are forced to conclude that the concepts have no application. We are forced to conclude that the world does not contain such things as apples, trees, rivers, tables, houses, and pens, or any other of the vast range of familiar types of object that these concepts identify and whose existence we ordinarily take for granted.¹⁶

The third point of challenge concerns our assumption that sensory experience is a source of accurate information about the physical world. Lack of perceptual access to the world would not entail that this assumption was definitely false: even if they were not physically perceptive, it would still be possible that our sensory experiences normally provide a reasonably accurate representation of the current environmental situation and that the environmental information they purport to give us is predominantly correct. But the reason why we ordinarily assume that this putative information is

¹⁶ Accepting that the concepts have these dispositional implications does not, of course, involve taking the relevant dispositions to be essential properties of those objects, if there are any, to which the concepts apply.

correct is that we assume that the experiences put us into perceptual contact with the physical environment and thereby enable us to read off the relevant environmental facts from what our perceptual awareness reveals. Once this latter assumption is abandoned, the issue of whether our experiences really do provide a reliable way of finding out about the physical world becomes immediately pressing, and the history of the discussion of this issue does not encourage us to suppose that a sceptical conclusion can be avoided. Indeed, for reasons which I have elaborated elsewhere,¹⁷ and which will partly emerge in a later phase of our discussion, I am sure that it cannot. I am sure that our only chance of being able to successfully resist the sceptic, and preserve our entitlement to trust what our senses purport to tell us, is by taking our stand on the common-sense position that we have direct epistemic access to the physical world through perception.

VII

The situation is looking awkward. The arguments I have advanced seem to show that we do not have perceptual access to the physical world, and, even without the further ways in which it threatens our ordinary understanding of the world and our place within it, this would be a very difficult conclusion to accept. Indeed, as I have acknowledged, a natural response would be to assume that, if it leads to this conclusion, something must have gone astray in my reasoning. However, there is one approach to the situation which we have not yet considered, and which could be thought to offer a way out of the difficulty.

The problem of perception has arisen because, given the failure of the fundamentalist view, we cannot understand how it is possible for a subject's awareness to reach beyond the boundaries of his mind and make contact with things in the external environment. Thus, with the exclusion of that view, the only kind of psychological contact that we can think of a subject as having, through sense experience, with something in the external reality is one that breaks down into his being in some more fundamental psychological state, which is not in itself physically perceptive, and additional facts that do not add

¹⁷ See *The Nature of Perception*, 226–40.

anything further to his current psychological condition. And we cannot then see how such additional facts could turn the overall psychological situation into one in which the subject is genuinely aware of the external item. But notice that this only creates a problem, with respect to physical-item perception, because we are taking it for granted that the physical world is something *external*—something with an existence distinct from, and logically independent of, facts about human mentality. We are assuming, in other words, that our awareness has to reach to things beyond the boundaries of the mind if it is to make contact with things in the world. This assumption is certainly part of our ordinary understanding of the situation, and, following my stipulation at the outset, it has formed the framework of our discussion so far. But there is an alternative, and one that, given the problem that has arisen, deserves consideration. This—stated in its barest form—is to take the physical world to be something whose very existence is constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. On this view of the situation, which is a form of idealism, there would be no difficulty in supposing that our sensory experiences bring us into perceptual contact with physical items. For we would no longer have to think of these items as belonging to a reality that lies beyond the realm of experience. Awareness would not need to reach beyond the boundaries of the mind in order to make ontological contact with the factors by which the existence of the physical world is constitutively sustained.

It was Bishop Berkeley who first insisted that we should respond to the issue of perception in this sort of way. Conscious of the two assumptions of our ordinary thinking, that the physical world is something to which we have perceptual access and that it is also something which is ontologically independent of the mind, and convinced, on philosophical grounds, that things whose existence was independent of the mind would be beyond the reach of perception, he concluded that the only way of achieving a credible, or indeed intelligible, outcome was to retain the first assumption and discard the second. And, although he did not explicitly endorse it, and sometimes endorsed positions that were incompatible with it, much of his philosophy, as I interpret it, is pointing in the direction of the idealist position that I have just identified.¹⁸ Whether such a position is correct, or even on the right

¹⁸ This is especially true, on my interpretation, of the position he develops in his *Principles of Human Knowledge* (G. Berkeley, *Philosophical Works*, ed. M. Ayers (London: Dent, 1975)). For a discussion of Berkeley's conflicting positions on the nature of the physical world see my essay 'Berkeley on the physical world', in J. Foster and H. Robinson (eds.), *Essays on Berkeley*, 83–108.

lines, is not something that we can at present decide. If we find it difficult to accept that we have no perceptual access to the physical world, we would also find it difficult to abandon the assumption of mind independence. All I am insisting, at present, is that there is an issue here that needs to be explored.

2

The Inscrutability of Intrinsic Content

I

In our everyday thinking we take it for granted that the physical world is ontologically independent of the human mind—that it has an existence which is logically independent of facts about human mentality. But when endorsed philosophically, this common-sense assumption has been seen to create a problem. For in recognizing the world as thus mind-independent, we can no longer understand how it is possible for us to have perceptual access to it. The only kind of psychological contact that a subject can have, through sense experience, with items in such a world is one that breaks down into his being in some more fundamental psychological state that is not in itself physically perceptible, together with certain additional facts that do not add anything further to his current psychological condition. And contact of this sort does not suffice for genuine perceptual awareness: it does not involve the subject's awareness ever reaching beyond the boundaries of his own mind. Hence the problem. For it is very hard to accept that the physical world is perceptually inaccessible—that we never see such things as trees and houses, never feel such things as pens and glasses, never hear such things as birdsong and human speech. We also noted that the denial of our perceptual access to the world poses a *prima facie* challenge to our ordinary view of things in other ways. In particular, we identified such a challenge to our ordinary conception of the world as *our* world, to our ordinary understanding of the kinds of objects that the world contains, and to our ordinary assumption that sense experience is a reliable source of physical information.

Given that the assumption of mind independence creates these difficulties and potential difficulties, we have been led to wonder whether, as philosophers, we should abandon it. In particular, we have been led to wonder whether we should adopt instead a certain form of idealism—a form which represents the physical world as something whose existence is constituted

by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. On such a view, the problem of perception would be eliminated, since awareness would not need to reach beyond the boundaries of the mind to make contact with the physical world. In effect, we would have access to the physical world simply by having access to the contents of our own experiences.

As well as rejecting our ordinary assumption of mind independence, such an idealist view is offering a radically reductive account of the physical world. For, in claiming that the existence of the world is constituted in the way envisaged, it is claiming that the obtaining of any physical fact is logically due to, and involves nothing over and above, the obtaining of certain non-physical facts, and so is claiming that the whole realm of the physical is, through and through, reducible to something else. This too, of course, is in conflict with our common-sense position. For just as we ordinarily assume the physical world to be ontologically independent of the human mind, so we also, at least implicitly, assume it to be something whose existence is philosophically fundamental, in a sense that excludes any reduction of the physical to the non-physical. These two common-sense assumptions, indeed, are closely associated: there is no remotely plausible way of taking the existence of the world to be, in the relevant sense, philosophically fundamental, without also recognizing its ontological independence of the human mind; and it is hard to see what rationale there could be for questioning the fundamental status of the world once we accept its independence. Added together, the two assumptions amount to what we can appropriately speak of as the *realist* view of the physical world; and the idealist approach we have envisaged opposes this realist view at both points.

The issue between the realist and idealist views of the physical world is due to take centre stage in our future discussion. But, to prepare the ground for this, there is something else that I want to cover first. In doing so, I shall, for the time being, continue to work within the common-sense framework of realism, though, as in our earlier discussion, the conclusions that will emerge will lead us to question whether this framework is ultimately tenable.

II

We have noted how denying the perceptual accessibility of the physical world calls into question our assumption that sense experience is a reliable source

of physical information, and the questioning of this assumption would call into question our whole capacity to acquire physical knowledge or form well-grounded physical beliefs. But even if we retain the assumption of perceptual access, and so avoid this radical sceptical challenge, there turns out to be a severe restriction on what we know, or can empirically discover, about the nature of the physical world, and this restriction applies both to our knowledge of the nature of physical space and to our knowledge of the nature of the things that are located within it. It is on the nature of this restriction that I want now to focus. Since the restriction is one that we have to acknowledge even when we take ourselves to have perceptual access to the physical world, I shall, in pursuing this topic, provisionally set aside the problem of perception and work on the assumption that such access is available. What is relevant, in the present context, is the sort of information about the world that we can hope, through perception, and reflection on what perception reveals, to acquire.

As I have indicated, the restriction in question is one that applies both to our knowledge of the nature of physical space and to our knowledge of the nature of the things located in it, and I am going to begin with the case of the things in space. Specifically, I am going to focus on the issue of what we know, or are capable of empirically discovering, about the nature of material objects, which we ordinarily think of as forming, with space, the main ontological ingredients of the physical world. One thing which I shall assume that we know about such objects, and which, indeed, I take to be implicit in our very conception of them as material, is that they are three-dimensional and persisting occupants of space. The way in which they persist through time is, it must be stressed, sharply different from the way in which they are three-dimensionally extended in space. To the extent that an object is spatially extended, it divides into distinct spatial parts, coinciding with the different spatial regions it covers. But a persisting object does not divide into temporal parts. For it is not the object itself that is temporally extended, but its existence, so that the whole object exists at each moment throughout the period of its persistence. This does not preclude our speaking of the temporal stages and phases of such objects (indeed, momentary stages featured in the discussion in the last chapter), but these entities must not be thought of as parts of the objects.

For reasons that will emerge, our investigation into the kinds of thing that are known, or could be empirically discovered, about the nature of material objects will mainly focus on how things stand with respect to one particular

category of properties that such objects possess. In order to identify this category, it will be helpful to begin by identifying two other categories of properties from which it is to be distinguished, and which it does not overlap. I should mention that I am here using the term 'property' to cover relations (two-or-more-place properties) as well as properties of items taken singly.

The first of these categories covers those properties which items possess in virtue of the way in which they occur within, and are structured by, the medium of space and time. Applied to the domain of material objects, this category includes both very general properties, like those of three-dimensional extendedness and persistence, which characterize all objects at all times, and more specific properties, like those of shape, position, distance, and motion, which vary across objects and times. I shall refer to the members of this category, collectively, as *spatio-temporal* properties, though, of course, many of them will be either purely spatial or purely temporal in character. Of particular importance in our discussion will be those spatial properties of material objects that cover aspects of their *intrinsic* characters—aspects of what the objects are like in themselves, or like in themselves at a given time. At any time, these intrinsic spatial properties of an object combine to form its current *spatial structure*, and this structure will cover both its overall shape and size and the internal arrangement of its parts. The second category covers properties that are of a causal, dispositional, or nomological kind. So, applied to the domain of material objects, it covers such things as the ways in which objects causally interact, their powers and sensitivities, their behavioural propensities, and the conformity of their behaviour and modes of interaction to certain natural laws. I shall call the members of this category *functional* properties. Note that although our present concern is specifically with material objects, the categories of spatio-temporal and functional properties, as specified, are not restricted to items in this domain. Physical events, for example, have spatio-temporal location and stand in causal relations. And perhaps physical science requires us to recognize the existence of certain types of submicroscopic object which, while having spatio-temporal and functional properties, do not qualify as material in the ordinary sense.

These two categories of properties do not exhaust the basic character of the material realm. Consider the current situation of the apple on the table now in front of me. There are its intrinsic spatial properties, of shape, size, and internal arrangement, which form its current spatial structure. There is its current position in relation to other material objects—for example, its central position on the table and its distance from the light switch. There is its current

state of rest relative to the table, and to almost everything else in the room, and there are the various states of rest and motion of its components. There are its current behavioural and causal dispositions, such as its disposition to float in water and its disposition to taste sweet when orally sampled, and there are the behavioural and causal dispositions of its components. There are various internal forces that keep the arrangement of its different parts relatively stable. But these spatio-temporal and functional properties of the apple and its components, together with any other properties of these general kinds, do not cover the whole of its character at the present time. They do not even cover all that is included in its basic physical character, as a chunk of matter occupying a region of space. And this is how things stand in the case of any material object at any time. What the spatio-temporal and functional properties fail to cover, in the basic physical character of a material object at a time, are those qualitative ingredients which provide, as we might put it, elements of qualitative content for the object's spatial structure—elements that fit into the structural framework that the object's spatial properties create. So, they fail to cover such things as the nature of the forms of material stuff of which the object is composed, and any further non-functional qualities that pervade portions of its surface or internal regions. Like the various aspects of its spatial structure, these qualitative elements contribute to the object's intrinsic character at the relevant time—to what, at that time, the object is like in itself—and I shall speak of them as aspects of its *intrinsic content*. More generally, for any space-located physical item, or item at a time, I shall speak of something as an aspect of its intrinsic content if and only if (i) it is an aspect of the item's intrinsic physical character, and (ii) it is wholly additional to what is covered by the item's spatio-temporal and functional properties, and those of its components. And I shall speak of those properties of an item that cover such aspects as its properties of intrinsic content, or, for short, its *content properties*. It is on the category of the content properties of material objects, and their components, that our epistemological investigation will mainly focus.

The reason why the investigation will focus on these properties in particular is that it is in this area that the interesting epistemological issue arises. Given our assumption that we have perceptual access to the physical world, and so can use that access as a means of epistemic access, there is no reason to doubt that we are already in possession of a great deal of knowledge of the spatio-temporal and functional properties of material objects, and that we are equipped to discover a great deal more. The shape and size of

objects, the spatial arrangements of their parts, the distances between them, their spatial paths over time, the ways in which they are disposed to behave and interact, the causal processes in which they feature, the laws of nature that govern their behaviour and interaction—all these are things which, in suitable circumstances, and with suitable techniques, are open to empirical scrutiny. Where I think there is a crucial epistemological issue is over the extent of our knowledge, or potential knowledge, of the content properties of material objects. Indeed, I shall try to show that, in sharp contrast with the richness of our knowledge and potential knowledge of the spatio-temporal and functional character of the material realm, properties of intrinsic content are wholly beyond the reach of empirical discovery—or, at least, of discovery in a form that reveals what these properties are.

Before I turn to the issue of intrinsic content, there are two preliminary matters that I need to deal with.

The first of these concerns my use of the term ‘intrinsic’, which I have already employed above, and which will continue to feature prominently in the ensuing discussion. As I have indicated, I am understanding the intrinsic character of a material object (or of an object at a time) to be what (what at that time) the object is like in itself, and I am understanding the intrinsic properties of an object to be those properties that cover aspects of its intrinsic character. And this is how I understand the notion of intrinsic character and intrinsic property quite generally, whether the ontological items involved are material objects or something else. But two points need to be underlined here. The first is that, understood in this way, the properties that are intrinsic to an object do not have to be logically essential to it: they do not have to be properties which it is logically impossible for the object to exist without. For example, the precise shape of an apple at a given time is an intrinsic property of the apple at that time, since it is an aspect of what, at that time, the apple is like in itself. But it is not a logically essential property of the apple, since the apple is logically capable of existing with a different shape, and, indeed, could have had a different shape at that time. The second point is that when I speak of what an object is *like in itself*, I intend to exclude not only the various ways in which it is related to other things, but also all its dispositional properties, and those of its parts—all the various ways in which it, and its components, are disposed to behave, or to exert certain forms of causal influence, in certain conditions. So, the disposition of an apple to float in water and the disposition of its pips to sprout if sown are not, in the relevant sense, aspects of what the apple is like in itself, and do not count,

for my purposes, as ingredients of its intrinsic character. I am not, of course, denying that the behavioural and causal dispositions of material objects are typically grounded on aspects of their intrinsic characters: typically, it is aspects of what these objects are like in themselves which, in the framework of the laws of nature, dispose them to behave and influence things in the ways they do. I am not even, here, excluding the possibility of supposing that aspects of the intrinsic characters of objects sometimes *on their own* (without the assistance of contingent laws) suffice to endow these objects with certain dispositions—though this is not a supposition that I would want to make myself.¹ All I am pointing out is that, as I am using the term ‘intrinsic’, dispositional aspects of an object’s character, however intimately associated with intrinsic aspects, do not themselves count as intrinsic.

The other preliminary matter concerns the way in which I have presented the issue of intrinsic content. I have represented it as uncontroversial that material objects have forms of intrinsic content, and I have taken the relevant issue to be an epistemological one—of whether we possess, or might be able to empirically acquire, any knowledge of what those forms of content are. Now I am in no doubt that the view which I have represented as uncontroversial is, in fact, correct—or, at least, is so relative to the framework of physical realism in which we are currently working. And I am also sure that, at this stage in the discussion, it is best for us to treat it as if it were uncontroversial. Nonetheless, I should at least mention that there is a way in which the view could be challenged—a way in which the nature of the material realm could be represented as entirely spatio-temporal and functional—and that I shall be examining this challenge in due course. It is clear, from what I have just said, that I regard the challenge as misconceived, and I shall explain why I think this when I come to consider the matter. For the time being, I shall continue to ignore any controversy over the existence of intrinsic content, and take the only thing at issue to be its epistemic accessibility.

III

One way in which we might suppose that we can, in a very direct fashion, discern aspects of the intrinsic content of material objects is by attending to the ways in which they sensibly appear to us. Take, again, the case of the

¹ Thus, see my *The Divine Lawmaker* (Oxford: Oxford University Press, 2004), ch. 5, sect. V.

apple I now see on the table in front of me. The apple visually appears to me as a roughly hemispherical patch of a certain shade of green, located at a certain distance in a certain direction. The apparent shape of the apple is not, at present, relevant to our enquiry, since an object's shape is an aspect of its spatial structure, not of its intrinsic content. What we need to focus on is the apparent colour of its visible surface. Assuming that the lighting is normal, and that my colour vision is not defective, our initial inclination may be to suppose that the colour appearance of the apple reveals an aspect of its true character. If it does, this will be the revealing of an aspect of intrinsic content. For if the apple really is, colourwise, as it looks—if its surface is genuinely pervaded by the colour quality that features in its visual appearance—that will be a clear instance of its possession of a content property in the relevant sense. Analogous cases arise with respect to forms of sensible appearance in other sense realms. Thus, if I now pick up the apple to eat it, its surface will (let us assume) feel cold to me, and if the quality of coldness that features in this tactual appearance genuinely pervades the apple's surface, this will be a further instance of intrinsic content. Again, if I go on to bite into the apple, its flesh will (let us assume) taste sweet to me, and if this sweetness really characterizes the flesh, this too will be an instance of intrinsic content. In these cases too, then, we might suppose that the forms of sensible appearance reveal forms of intrinsic content that the object possesses. And, of course, if, in the case of the apple, aspects of intrinsic content can be detected through sensible appearance in the ways envisaged, similar aspects of content in other material objects will be available for detection in a similar fashion.

Should we, then, think that sensible appearance is able to reveal aspects of intrinsic content in this sort of way? One problem with doing so is that the same object, with the same intrinsic character, is liable to assume different forms of appearance in different conditions of observation, and in many cases it is hard to see how, without some degree of arbitrariness, we can decide which of the different modes of appearance is representationally correct. Thus, the colour appearance of an object may vary with the distance at which it is viewed; the apparent temperature of something held in the hand will vary with the temperature of the hand; the apparent flavour of something one is tasting can vary with the kind of food one has just been eating. It is not easy, in such cases, to think of an objective basis for selecting a certain form of appearance as what uniquely captures the object's true character in the relevant respect. Added to this difficulty is the further problem that, even in what would count, in the ordinary sense, as the same conditions of

observation, different subjects may perceive the relevant aspects of an object's character in different ways as a result of some relevant difference in their physical or psychophysical make-up, and, again, there often seems to be no objective basis for deciding whose perception is correct. Jonathan Bennett, for example, drew attention to the case of the substance phenol-thio-urea, which tastes bitter to about three-quarters of the population and is tasteless to the rest.² There is also, in this area, the quite different problem that among the differences in the ways in which different subjects perceive things there could be ones which, because they cannot manifest themselves through differences in how the subjects behave or in what they say, are empirically undetectable. Maybe my wife sees the things that we would both describe as green in the way in which I see things that we would both describe as yellow, but, because of corresponding differences in how our colour experiences align with physical items across the whole spectrum, this difference is not something that either of us, or anyone else, can detect. Even the possibility that this is the case seems to undermine my entitlement to take my own system of colour experience as a reliable basis for colour ascription.

These various points about actual and possible variations in the character of sensible appearance across conditions of observation and subjects create difficulties for the view that we can establish the precise sensible character of material objects in the relevant respects, but they do not prove that there is no basis for thinking of objects as having sensible qualities of the relevant kinds at all. So, for example, the fact that my system of colour experience may be different, in the way envisaged, from my wife's may undermine my entitlement to think that the surface of the apple in front of me has the green colour that I see it as having. But it does not show that I am not justified in taking it to have some sensible colour—some colour of the kind that can feature in the content of visual appearance. However, there is something else that does show this, and, indeed, that shows that there is no basis for thinking that any of the kinds of sensible quality we have been considering are physically realized at all. For any such basis we might suppose ourselves to possess is immediately undermined once we take account of the factors that causally underlie the relevant forms of sensible appearance.

Take, first, the case of colour appearance, and consider why, as I now view the apple on the table, its visible surface looks coloured to me in the way that

² J. Bennett, 'Substance, reality, and primary qualities', *American Philosophical Quarterly*, 2 (1965), 1–17.

it does. Is it because it is characterized by some sensible colour—the same colour as features in its visual appearance or some other colour of the same general (sensible) sort? We might initially suppose that it is: intuitively, it is as natural to think of objects as looking coloured because they are (in the relevant kind of way) coloured as it is to think of them as looking extended and shaped because they are extended and shaped. But the scientific findings prove otherwise. The reason why this surface portion looks green to me in the way it does is that the light it transmits to my eyes is of a certain wavelength composition—a composition that equips it to induce those forms of firing in my optic nerves that give rise (in my psychophysical system) to the relevant form of colour experience. And the reason why it transmits light of that sort is not that it is characterized by the relevant form of greenness, or by any other sensible colour, but that, in the framework of the laws of nature, certain aspects of its microphysical character—aspects that have nothing to do with sensible colour—dispose it to absorb certain wavelengths of light and to reflect others. The same story holds for all other cases of an object's colour appearance, except that in some cases, like that of the sun, or that of a red-hot piece of coal, the object becomes visible, and has its colour appearance, by generating light, rather than by reflecting it. But once it is accepted that the reason why material objects appear coloured has nothing to do with their possession of the sort of colour that features in colour appearance, this completely undermines the view that they do possess colour of that sort. For it is clear that the only thing we could even think of as giving us grounds for ascribing such colour to objects is that this accords with how the objects visually appear; and once we know that the reasons for their so appearing have nothing to do with their possession of such colour, we know that there are no proper grounds here at all. In the light of our scientific knowledge of the situation, the only way in which we could secure a warrant for thinking of material objects as coloured would be by equating the physical colours they possess with their dispositions to assume certain forms of colour appearance in certain conditions, or with the microphysical properties on which these dispositions are grounded, and in neither case would this involve thinking of the objects as having the sensible form of colour that features in colour appearance.

Analogous considerations apply to cases of thermal and gustatory appearance. As science makes clear, the reason why the apple feels cold to me and tastes sweet has nothing to do with its possession of the coldness and sweetness that feature in these forms of appearance, or with its possession of

any other sensible qualities of temperature and flavour, capable of featuring in the content of thermal and gustatory appearance. Rather, it is entirely to do with certain aspects of the apple's microphysical character, as specified by the sciences of physics and chemistry, and with the kinds of effect that these aspects are equipped to have on my sensory nerves, and thereby on my sensory experience, in the relevant conditions of tactual and gustatory encounter. And, of course, a similar situation obtains in cases of thermal and gustatory appearance quite generally. Once this fact is recognized, it removes any basis for supposing that material objects have qualities of temperature and flavour of those sensible kinds. As with the case of colour, the only way in which we can be justified in taking them to have qualities of temperature and flavour is in the form of phenomenal dispositions or the microphysical properties on which these dispositions are grounded (though in the case of temperature perhaps the relevant dispositions are better taken to be of a broader—not purely phenomenal—kind). Analogous facts about the factors which causally underlie instances of auditory and olfactory appearance show that, in these cases too, we are not warranted in thinking that the sensible qualities of sound and odour that feature in them genuinely characterize material objects, or occur in the physical world at all.

What we have reached here is, in effect, an endorsement of John Locke's view of the secondary qualities, as set out in Book 2 of his *Essay*.³ The qualities in question are those of colour, sound, flavour, odour, and (as tactually or somatically perceived) temperature. As we ordinarily conceive of them, these qualities are sensible qualities, which feature, or are capable of featuring, in the content of sensible appearance, and what makes them, as sensible qualities, relevant to our present discussion is that (i) we might initially suppose that, through their featuring in sensible appearance, we can, in appropriate circumstances, discern material objects to possess them, and (ii) if material objects do genuinely possess them, they qualify, as thus possessed, as forms of intrinsic content. Locke's position was that, taken in their sensible form, the secondary qualities are confined, in their occurrence, to the content of sensible appearance, and that the only way in which they can characterize material objects, or physical items of any kind, is in the form of powers (causal dispositions)—typically powers to produce the kinds of sensory experience in a percipient that will give the physical items that possess

³ J. Locke, *An Essay Concerning Human Understanding*, ed. A. Campbell Fraser (New York: Dover, 1959), Bk. 2, ch. 8.

them certain kinds of sensible appearance. Locke's way of expressing this position was sometimes misleading, and not all the arguments he advanced in its support are persuasive. But what matters here is that, for the purposes that concern us, the position is vindicated by the scientific findings about how the relevant forms of sensible appearance are causally generated. And, though the science he relied on was of a theoretically cruder kind, this was, as I understand him, Locke's own main reason for adopting it.

Admittedly, the scientific findings do not, strictly speaking, establish that the relevant categories of sensible quality are not physically realized. They only establish that we have no reason to suppose that they are. Specifically, what they establish is that we should not regard the ways in which physical items sensibly appear to us as providing any evidence that they possess such qualities (since such possession does not figure in the factors that causally underlie the relevant modes of appearance); and, since there is nothing other than their modes of sensible appearance that could give us any grounds for supposing that physical items possess such qualities, we are left with no warrant for thinking that the qualities are physically realized at all. This result leaves open the possibility that at least some of the relevant qualities are, in fact, physically realized, though in a form that has nothing to do with the ways in which things sensibly appear to us, and that is entirely beyond the reach of empirical detection. So, it leaves open the possibility that, invisibly and undetectably, the visible surface of the apple now in front of me is pervaded by a sensible colour. Indeed, it leaves open the possibility that, invisibly and undetectably, this surface is pervaded by the very colour that I seem to see pervading it. And, of course, if the surface of the apple were genuinely pervaded by a sensible colour, that would count as an instance of intrinsic content. But such possibilities do not concern us in the context of our present investigation. What we are trying to identify are examples of things that we know or could hope to discover about the intrinsic content of material objects. *Ex hypothesi*, any such undetectable physical realization of a sensible colour, or of any of the other relevant kinds of sensible quality, would not provide an example of that sort.

I should add that, although the scientific findings themselves leave open the possibility of the relevant kinds of quality being physically realized, this is something which, in my view, we can exclude for other reasons; or at least we can do so unless we are prepared to entertain the idea of the physical world being, in its own character, wholly or partly mental. The reason why we can exclude it is that we cannot ultimately make sense of such qualities

being realized outside the content of sensory experience. This, I think, is virtually self-evident, on reflection, in the case of qualities of flavour and odour. For it is surely clear, from our very concept of them, that the kinds of flavour and odour that feature in the content of gustatory and olfactory experience are essentially tied to the content of such experience, and that, barring a wholly or partly mentalistic account of the physical world, the only form in which we can make sense of physical flavours and odours is either as Lockean dispositions or as the microphysical properties on which they are grounded. In the case of the secondary qualities associated with the other sense realms the point may be less obvious, but, as I argued in *The Nature of Perception*, I think that there are considerations which should lead us to the same conclusion.⁴ Two considerations I regard as especially important, though I shall only touch on them briefly here. Let us say that, for any qualitative item Q (any quality, property, kind, and so on), a conception of Q is *transparent* if and only it reveals what, as a qualitative item, Q is; and where a conception of a qualitative item is not transparent, let us speak of it as *opaque*. So, conceiving of triangularity as *triangularity* is transparent, while conceiving of it as *the geometrical shape discussed in the fourth chapter of such and such a book* is opaque. Then the first consideration is that in the case of any secondary sensible quality, we can only form a transparent conception of it by focusing on what it is like to encounter it in the content of the relevant form of experience. So, we can only form a transparent conception of a sensible colour by focusing on what it is like to be visually aware of it, and can only form a transparent conception of a sensible pitch by focusing on what it is like to be auditorily aware of it. The second consideration is that even when, in our everyday thinking, we take some secondary sensible quality to be physically realized, we think of this realization as in some (though not a precisely defined) way essentially linked with the realization of a certain mode of objective appearance. So, in ascribing a sensible colour to an object, we think of the object's possession of this colour as in some way essentially linked with the way in which it objectively looks, and, in ascribing a sensible pitch to the chime of bell, we think of the chime's being of this pitch as in some way essentially linked with how it objectively sounds. As I see it, these two points can only be adequately explained if we accept that the qualities in question are essentially experiential and that their only form of realization is in the content of the relevant kinds of experience. So, to explain the first point, I think we

⁴ J. Foster, *The Nature of Perception* (Oxford: Oxford University Press, 2000), 130–47.

have to say that the reason why we can only transparently conceive of the qualities in terms of what it is like to encounter them experientially is simply that they are exclusively experiential qualities. And, to explain the second, I think we have to say that the way in which our ordinary thinking links their physical realization with objective modes of appearance is an implicit, though confused, acknowledgement of their experiential status. I shall not elaborate further on these claims here, though they are covered in detail in the earlier book.

In representing the secondary sensible qualities as ones whose only form of realization is experiential I should, to be precise, have merely represented them as qualities whose only *fundamental* form of realization is experiential. Obviously, there would be no problem about recognizing the non-experiential realization of these qualities if we could represent that realization as reducible to facts of a different kind. This point has no bearing on the matters we are presently discussing, but it will acquire some importance at a later stage, when the possibility of an idealist account of the world is under consideration.

IV

We are considering the question of what sorts of thing we know, or have the capacity to discover, about the nature of material objects, and we are focusing particular attention on the issue of intrinsic content—the issue of what we know, or are equipped to discover, about those ingredients of the intrinsic character of an object at a time that form elements of qualitative content within the framework of its spatial structure. All that we have established so far is that, despite what our experiences initially suggest, and what we doubtless tend to accept in our ordinary thinking, there is ultimately no basis for supposing that, taken in their sensible form—as items which feature, or are capable of featuring, in the content of sensible appearance—the secondary qualities, of colour, sound, flavour, odour, and temperature, genuinely characterize material objects, or anything else in the physical world, and that the only way in which we can justifiably think of these qualities as physically realized is in the form of phenomenal dispositions (or, in the case of temperature, perhaps dispositions of a broader kind) or the microphysical properties on which these dispositions are grounded. I shall also assume, in accordance with what I have tried to establish in detail

elsewhere, that, in their sensible form, the qualities in question are not even capable of fundamental realization outside the content of sensory experience. All this is more or less in line with the view of the secondary qualities espoused by Locke.

Where do we go from here? Well, let us stay, for the time being, with Locke. For, as well as excluding the secondary qualities from the physical world except in the form of powers, Locke himself, in his *Essay*, offered his own account of the intrinsic nature of material objects, and it will be helpful to begin by looking at what he proposed. One thing we need to bear in mind, here, if we are to understand what led him to his position, is that Locke was focusing on the issue of the nature of material objects at a time when the empirical science of matter was in its infancy, and so without the benefit of the knowledge we now possess about the different chemical elements and their internal composition. Believing that there were no physical secondary qualities in any but a dispositional form, and finding no grounds in the scientific knowledge of his day for supposing that there were different fundamental forms of matter, he found it natural to assume that matter was qualitatively homogeneous, so that the only way in which two parcels of matter could, at any time, differ in intrinsic character was by differing in their spatial structure.

Approaching the issue in these terms, what Locke concluded was that the property which makes matter, in all its instances, the sort of thing it is, and which, in my terminology, furnishes each material object with the totality of its intrinsic content, is *solidity*. We need to begin by saying how Locke understood the notion of solidity.

The term 'solid' can be used in a variety of different senses. It is used in geometry to signify the property of three-dimensional extendedness, so that a cube counts as solid and a square does not. It is used in chemistry and physics to contrast with the term 'fluid', and so to apply to such things as stone and ice, but not to such things as water and air. And it can be used in ordinary parlance to apply to things that are strong and stable, in contrast with things that are weak or flimsy. Locke's sense is different again. What Locke meant by solidity, as he explains in the fourth chapter of Book 2, is that intrinsic property of matter which makes material objects mutually impenetrable—that which gives each particle and parcel of matter, in relation to other such particles and parcels, exclusive possession of the region of space it occupies, so long as it occupies it. He took solidity in this sense to be an essential property of matter, 'inseparably inherent in body, wherever or

however modified'.⁵ And, apart from three-dimensional extendedness and persistence, he took it to form the whole intrinsic nature of matter, so that, in his system, parcels of matter become, in effect, simply mobile parcels of realized solidity. Or, at least, this is how we might represent his system if we discount his flirtation with the absurd idea of an in-itself-featureless substratum.

Now if matter is indeed characterized by this intrinsic property of solidity, this will certainly provide a form of intrinsic content for material objects. And this is a point that will continue to hold good in a framework in which we reject Locke's assumption that matter is homogeneous, and accept that there must be further forms of intrinsic content to distinguish objects composed of one sort of matter from those composed of another. But before we conclude that we are making some progress in acquiring knowledge of intrinsic content, there are two issues that we need to address.

The first issue is over whether we should recognize a property of solidity of the relevant kind at all. Let us, for the sake of argument, assume that Locke is right in thinking of bodies as mutually impenetrable. Even so, it might be argued that this impenetrability is a consequence not of the character of the stuff of which they are made, but of their principle of individuation. Thus, it might be claimed that, at any time, a given parcel of matter is numerically distinguished from all other parcels of matter by the distinct region of space that it occupies, and that the reason why it is impossible for two parcels of matter to occupy the same three-dimensional region simultaneously is that there can only be numerical distinctness of matter where there is spatial separation. If this were accepted, then there would seem to be no room for a property of the sort that Locke envisages, since this property is explicitly introduced as that which underlies, and accounts for, impenetrability; and if impenetrability is already logically ensured by the principle of individuation, there is no such role for any property to play. What Locke seems to need, for his account to work, is that impenetrability should be a causal power—the disposition of a body to prevent any other body from entering the space it currently occupies—so that solidity can be identified as that intrinsic property of bodies on which, in the framework of the laws of nature, this power is grounded. Curiously, in a later chapter Locke himself explicitly endorsed the spatial principle of individuation, though without referring back to his discussion of solidity, and without reference to the topic of solidity at all. But

⁵ *Essay* 152.

in the present context this is a complication which it will be convenient to ignore.⁶

The most straightforward way of trying to defend Locke's account of solidity, in the face of the envisaged challenge, would be by simply rejecting the putative principle. And I am inclined to think that, taken as something exceptionless, the principle should indeed be rejected. Thus, while it is plausible to regard position in space as providing the basic framework for the individuation of material objects, I am also inclined to think that we can coherently envisage cases in which different parcels of matter simultaneously occupy the same three-dimensional region, and, indeed, that we can envisage the kinds of empirical evidence (in terms of how things detectably behave in the situations in question) that would make it plausible to conclude that this had actually happened. But even if the spatial principle were accepted, someone wanting to retain the substance of the Lockean position would still have an effective line of defence. This is because while the principle represents the mutual impenetrability of bodies as a logical consequence of their mode of individuation, impenetrability of that sort would not exhaust the sense in which bodies are mutually obstructive. After all, the principle allows for the possibility of cases in which, by all empirical tests, two bits of matter *seem* to pass through each other unhindered: it simply obliges us to describe such cases as ones in which the quantity of matter diminishes as the objects seem to merge, and increases as they seem to separate. The fact that, in practice, bodies do not behave in this way—that they are in competition *causally*, and not just *formally*, for the occupancy of space—shows that they possess a power of mutual obstructiveness that is additional to their mode of individuation. So, irrespective of the status of the principle, it would be possible to appeal to this mutual obstructiveness in explaining the relevant notion of solidity. It would be possible to say that, whether or not there is the logical possibility of spatial co-occupancy, solidity is that aspect of the intrinsic nature of bodies which, in the framework of the laws of nature, makes them mutually causally obstructive. If, as I would recommend, the spatial principle is rejected, we can think of this obstructiveness as a power to resist penetration, which is how Locke himself conceived of it. With the principle in place, the obstructiveness would become, in effect, the power of

⁶ The endorsement comes in chapter 27 of Book 2 ('Of Identity and Diversity'), which Locke added to the *Essay* in its second edition. Consider in particular: 'Could two bodies be in the same place at the same time; then those two parcels of matter must be one and the same, take them great or little' (*Essay* 441).

matter to resist annihilation in situations of spatial competition. But either way the basic thrust of Locke's account can be preserved.

One way or another, then, a defender of the Lockean position can meet the challenge over the issue of individuation. And this leaves us on track for being able to think of ourselves as having, through our recognition of Lockean solidity, knowledge of an aspect of the intrinsic nature of matter, beyond its three-dimensional extendedness, and hence knowledge of a content property of material objects. But it is at this point that we encounter the second of the issues that needs to be addressed. For what we now have to consider is whether our knowledge of the solidity of material objects (assuming we have it) gives us knowledge of an aspect of their intrinsic content in any interesting sense.

It is not hard to see why this issue arises. Locke introduces solidity as what it is about material objects that makes them mutually impenetrable—that ensures that each particle and parcel of matter has, in relation to other such particles and parcels, exclusive possession of the region of space it occupies. And let us assume, for the sake of argument, that we can think of impenetrability as a causal power—a disposition to prevent penetration—and can take the solidity that Locke has in mind to be the intrinsic property on which, in the framework of the laws of nature, this power is grounded. But none of this, of course, tells us what solidity as such is. It identifies solidity by reference to its role as the ground of impenetrability, but does not reveal the qualitative essence of what it is that plays this role. The identification of the relevant property is, to use the terminology I introduced earlier, *opaque* rather than *transparent*. But if the only way in which we can identify solidity is opaque, then the knowledge that material objects have this property does not tell us anything of substance about their intrinsic characters—about what such objects are like in themselves. It only tells us that there is an aspect of what they are like in themselves that nomologically sustains the relevant power.

Locke himself seems to think that transparent knowledge of solidity is available. He stresses that the property he has in mind is one which we experientially encounter whenever we feel the resistance of other material objects to some form of tactual pressure that our own bodies exert on them. And although he introduces us to the property in an opaquely identifying way, as that which underlies impenetrability, and thereby accounts for the resistance we feel, he seems to think that we can discover its qualitative essence by attending to the content of this tactual experience. Thus, at the

end of the chapter, he writes, 'If anyone asks me, *What this solidity is*, I send him to his senses to inform him. Let him put a flint or a football between his hands, and then endeavour to join them, and he will know'.⁷ And the rest of the passage seems to indicate that he thought of the kind of experience involved as directly manifesting the real nature of solidity, in the same way that visual experience manifests to the sighted the real nature of sensible colour. But surely, in this, Locke is simply mistaken. All that the tactual experience reveals is the force of resistance that the object exerts to the pressure directed on to it. The solidity of the object is detected only in so far as one feels a barrier to the progress of one's hands—a region that one is unable to penetrate. There is no discovery of what it is about the object that empowers it to be such a barrier; indeed, the experience does not, in itself, indicate or suggest that the resistance is grounded on anything intrinsic at all. Of course, the feeling of resistance is not the only ingredient of the total tactual experience. If I press something between my hands, I will feel aspects of its shape and size, and of the texture and temperature of its surface. But these other experiential elements do not contribute to my experience of the object's solidity in the Lockean sense. My experience of the solidity is simply the experience of the obstructiveness, and the intrinsic factor that underlies this obstructiveness remains concealed.

If we cannot gain a transparent knowledge of solidity from what is revealed by our tactual experience of material objects, it is clear that we cannot gain such knowledge from anything else that is revealed by ordinary observation, whatever sense realm or realms the observation involves. The most, in this area, that we can learn by ordinary observation is that such objects are, in the relevant way, mutually obstructive (however we care to interpret this), and the only conception of solidity that this will enable us to form will be as that property on which the power of obstruction is grounded—a conception that does not reveal what the property is. Unless there is some further discovery that we can make about it, which takes us beyond the limits of what is accessible to ordinary observation, Lockean solidity, even if a genuine aspect of the intrinsic content of material objects, will not provide us with a case where we can achieve knowledge of such content in any interesting sense.

⁷ *Essay*, 156–7.

V

Even if material objects are characterized by an intrinsic property of solidity of the sort Locke had in mind, its nature is not discoverable through ordinary observation. Nor, given our conclusions about the secondary qualities, can we now think of ourselves as able, by ordinary observational means, to gain transparent knowledge of any other form of intrinsic content. Ordinary observation is equipped to reveal certain facts about the spatio-temporal properties of material objects, including facts about the intrinsic spatial structures of objects at times. It is also equipped to reveal facts about the behavioural and causal dispositions of objects, including their dispositions to assume certain forms of sensible appearance, and including the mutual obstructiveness that we experientially encounter through touch. But if tactual perception does not reveal the nature of the intrinsic property on which the obstructiveness is grounded, and if there is no basis for believing that the secondary qualities are physically realized in their sensible form, the only transparent knowledge that we can rationally hope to gain, through ordinary observation, about the intrinsic character of an object at a time will be confined to facts about its spatial structure, and will not cover the nature of any of those qualitative ingredients that provide elements of content for that structure.

We cannot hope to gain transparent knowledge of intrinsic content by ordinary observation. But we might still hope to do better by appeal to physical science. We normally think of physical science, or at least the sciences of chemistry and physics, as able to get below the level of ordinary observation and provide an account of the fundamental nature of things. So, perhaps it will be at this deeper, more penetrating level of investigation that a transparent specification of intrinsic content will emerge. To explore this possibility, let us continue with the example of the apple on my table and consider what science might be equipped to tell us about its nature.

The way in which we would think of science as equipped to provide us with a more penetrating account of the nature of the apple would be through its capacity, quite generally, to probe, below the level of what is ordinarily observable, into the internal make-up of material objects, identifying the types of imperceptible constituent of which perceptible objects are composed, and

determining the ways in which these constituents fit together, spatially and through certain kinds of causal bonding, to form the larger objects we observe. There is a series of levels through which such compositional accounts can be developed, as the scientific investigation focuses on increasingly more fundamental—mereologically simpler—types of constituent. Thus, in the case of the apple, the account might begin with a specification of its cellular make-up, as revealed by examining various portions of it under a microscope. This could be developed into a description at a more fundamental level by specifying the molecular composition of the types of cell already identified. This, in turn, might lead into a still more fundamental account of internal make-up in terms of the atomic composition of the molecules. Finally, there is the most fundamental level, or perhaps series of levels, of description available, which takes account of the internal make-up of the atoms.

Where, in all this, might we hope to gain any transparent knowledge of the content properties of the apple? Obviously not, at any level, in the specification of the spatial arrangement of the types of constituent under consideration, nor in the identification of the causal forces that bind them together. The only point where there could be any chance of intrinsic content showing up would be in the specification of the nature of the constituents themselves. But the trouble is that for any type of constituent, the most that science can tell us about its character, or its character at a given time, is concerned with either (a) its spatial properties, (b) its behavioural and causal dispositions, or (c), if it is not physically fundamental, the way in which it is composed of simpler types of constituent. Information about (a) and (b) contributes nothing to the specification of intrinsic content; and information about (c) could only do so if it included a transparent specification of, or of certain aspects of, the intrinsic content of the simpler constituents. But since (a), (b), and (c) represent the only options for scientific information at any level of investigation, exactly the same limitations will apply to the case of the simpler constituents, which means that the ultimate elements of content can never be revealed.

To illustrate, suppose we have already reached the level of description at which we know the molecular composition of the apple. To seek for knowledge of intrinsic content, we then ask our scientific informants to tell us about the nature of the molecules involved. For each type of molecule, they may well be able to specify its shape and size, and they may well be able to specify its most important dispositional properties, such as its gravitational mass and its powers of chemical interaction. But any attempt

to reveal its intrinsic content will inevitably take the form of moving to the next lower level of description and specifying the way in which molecules of this type are composed of atoms. Then, of course, in the hunt for intrinsic content, the same question will be raised with respect to the nature of each type of atom, and the only way of trying to capture the content properties of an atom will be to move to a still more fundamental level of description, in terms of subatomic particles. At some point the process of moving to more fundamental levels of description will have to stop, either because scientific investigation has reached the point where there is no further physical complexity to be unpacked or because its current techniques are unable to discern the further complexity that is there. And, at whatever point it stops, the only transparent knowledge available of the nature of the types of physical particle that we are then left with—the types that are physically fundamental or that science has provisionally to treat as such—will exclusively concern their spatial and dispositional properties: their properties of intrinsic content will remain concealed. Thus, suppose the process ends up recognizing, amongst other things, a fundamental ontology of electrons, or, at least, one that has provisionally to be treated as fundamental. We may be able to learn, transparently, of the spatial character of these particles (of their shape and size, assuming they are spatially extended). We may also be able to learn, transparently, of their behavioural and causal dispositions, such as their disposition to be attracted to protons and to repel one another. But we can learn nothing transparent of their intrinsic content—nothing, beyond their spatial properties, of what electrons are like in themselves. The scientific knowledge available will, at best, enable us to identify their intrinsic content opaquely, as that (whatever it is) which, in the framework of the laws of nature, endows this type of particle with its distinctive dispositions. The same limitation will apply to what we can learn about any other type of particle that the investigation reveals and that it treats as physically fundamental.

It turns out, then, that science is no better equipped than ordinary observation to provide us with transparent knowledge of the content properties of the apple. And, of course, what holds for the case of the apple holds for material objects quite generally. Scientific investigation can probe much more deeply than ordinary observation into the internal make-up of material objects and uncover levels of compositional complexity that such observation cannot discern. But, whatever kinds of object are at issue, and to whatever microscopic or submicroscopic level the compositional investigation is taken, science can only provide transparent information about the spatio-temporal

and functional properties of the constituent entities whose existence it reveals. It cannot provide such information about their intrinsic content.

It is not difficult to understand why a limit on the scope of scientific discovery occurs at this point. Even at a level where the constituents of objects are too small to be accessible to ordinary perception, we can understand how a scientific investigation may manage to uncover things about their shape and size, their spatial and spatio-temporal arrangement, their relationships of bonding, their behavioural causal dispositions, and the laws that govern them. For, in favourable circumstances, hypotheses about these things can be scientifically evaluated by how well they explain the empirical data. So, in the case of the apple, a scientific investigation may reveal the need to recognize certain spatio-temporal and functional facts about the chemical composition of the apple's cells in order to have an adequate explanation of the ways in which cells of this type are found to behave in a range of conditions. In contrast, even if scientists could formulate hypotheses about intrinsic content, such hypotheses would, from a scientific standpoint, be of no explanatory value. Thus, once science has determined the spatial character of the types of particle that it treats as fundamental, and determined the ways in which, under the laws of nature, they are disposed to behave and exert forms of causal influence, there is nothing further, by way of explaining any empirical data, that hypotheses about intrinsic content could add. With the number of different types of particle established, and their spatial functional characters fully specified, everything about them that could contribute to a scientific explanation of anything is already in place—irrespective of what their forms of intrinsic content happen to be. The only way of envisaging anything further of explanatory relevance would be by envisaging a further level of compositional analysis, in which some of the types of particle that had originally been treated as fundamental were represented as composed of simpler constituents.

Since neither ordinary observation nor scientific investigation is capable of revealing what forms of intrinsic content material objects possess, we are forced to the conclusion that the nature of this content is, in all cases, empirically inscrutable. Forms of content may sometimes be empirically identifiable opaquely, by reference to the detectable dispositions which they nomologically sustain; for example, we may, in this way, be able to identify the Lockean property of solidity as that which, in the framework of the relevant laws, makes bodies mutually obstructive, or identify the content property of a certain type of particle as that which, in the framework of

the relevant laws, endows it with its distinctive powers and sensitivities. But empirical investigation can never identify intrinsic content in a way that is transparent—that reveals what that content, in itself, is.

Even as it stands, this would be a far-reaching conclusion. But, in one respect, the inscrutability of intrinsic content goes even deeper. For it turns out that the limitation on the scope of empirical knowledge is matched by a coextensive limitation on what we are capable of expressing or conceiving of in physical terms—a coextensive limitation on the descriptive resources of our physical language and our system of physical concepts. It is not just that we cannot empirically discover what forms of intrinsic content material objects possess: we cannot, in physical terms, even envisage what these forms might be. At least, we cannot do so, granted that we are now accepting that, in their sensible form, the secondary qualities are incapable of fundamental realization outside the content of experience. In this sense, the limitation on what physical science can reveal, when investigating the compositional complexity of material objects, is not perceived as a practical limitation from the viewpoint of the scientist. He never finds himself wanting to evaluate hypotheses about the nature of particle content, since the possibilities for content are not scientifically specifiable. The point where the nature of the physical situation falls beyond the scope of empirical tests is the point where he runs out of vocabulary with which to formulate the options, and concepts by which to conceive of them. Nor could this deficiency ever be remedied by a further enrichment of the scientific conceptual scheme.

VI

So far, we have been taking it for granted that material objects do have properties of intrinsic content. But, given what our discussion has revealed—that the nature of any such content is empirically inscrutable, and that we cannot even, in physical terms, envisage what it might be—it might be wondered whether this assumption should be abandoned. Perhaps we can find some way of representing the natures of material objects as entirely spatial and functional, and so as in principle amenable to empirical discovery and physical specification. And if we can, perhaps the right conclusion to draw from the inscrutability of content is that this way of representing the situation is correct, and that the properties of content that we have been trying to identify do not exist.

If we were to pursue this idea, the point where we would have to look for this new representation would be in our understanding of the nature of the particles that form the fundamental constituents of material objects. Instead of thinking of these particles as possessors of intrinsic content—for example, as items of some kind of space-occupying stuff—we should have to think of them simply as mobile items of causal power, with no further space occupant to form the vehicle of the power or cluster of powers involved, nor any non-functional properties on which the power or power cluster is nomologically grounded. I shall speak of this conception of the fundamental particles as the *powers view*. Given that the relevant items of power have no intrinsic content, and so cannot be items of space-occupying stuff, it is natural to think of them as point-located entities, without spatial extension, though the point at which an item is at any time located will be the centre of an extended field, or range of fields, of potential causal influence, in accordance with the nature of the power or powers it embodies. It was in this form that the view received its classic modern exposition in a book by Rom Harré and Edward Madden.⁸ It must be stressed that the kind of causal power that features in this view of the world is very different from the kind that has featured in our previous discussion. In the previous discussion we have always taken a causal power to be a property—a disposition of something to exert a certain kind of causal influence in certain kinds of condition. But in the context of the view on which we are focusing, causal powers are present in the world in an ontological form, as concrete space-located entities, and these entities are not (at least in the first instance) possessors of power, but instances of power.

There is no denying that, if it can be made to work, the powers view has considerable appeal. There is an awkwardness in having to recognize forms of physical property whose nature is beyond the reach of empirical discovery or physical specification, especially when those properties are assigned such a fundamental role in the make-up of the world. If we could find an acceptable way of eliminating them, that would be an advantage. Moreover, there is a sense in which the powers view is distinctively congenial to the outlook of physical science. For since any content properties of the types of particle it postulates are empirically inscrutable, and not even transparently

⁸ R. Harré and E. Madden, *Causal Powers* (Oxford: Blackwell, 1975). As well as expounding and defending their own version of the powers view, the authors see it as the culmination of the ideas of a number of earlier philosophers and scientists, who include Boscovich, Kant, Faraday, and Priestley.

conceivable in physical terms, it is inevitable that science treats these particles, for all practical and theoretical purposes, as if they were merely mobile items of causal power. It treats them as if there was nothing more to their natures than the spatial and dynamic properties that the powers view accords them. What we still have to decide is whether the powers view can be made to work, and, indeed, whether it is even intelligible. One *prima facie* problem for it, as we shall now see, concerns an issue about the *content* of the forms of power it envisages—about what these powers are, supposedly, powers to do. This notion of the content of a power—*dynamic* content—must not, of course, be confused with that of the intrinsic content of an object.

It is not difficult to see how, under the powers view, a *prima facie* problem of dynamic content arises. If all that was envisaged was a situation in which *some* of the fundamental types of particle were items of power, leaving others in possession of intrinsic content, there would be no suggestion of anything amiss. For, as potential targets of causal influence, the intrinsically contentful types of particle could form a source of content for the relevant particle powers, as things whose behaviour the powers were, directly, or indirectly, powers to affect. But with all the fundamental particles construed as merely items of power, it seems that all the powers involved will have to be powers to affect the behaviour of items of power, and the difficulty then will be in understanding how there could be anything to give any of the powers any content at all. It seems that any putative possession of dynamic content will presuppose some further possession of such content being already in place. Let us start with a simple example. Suppose that there are just two kinds of fundamental particle, K1 and K2, and that their only power is one of mutual attraction, an attraction which, let us assume, varies inversely in strength with the distance between them. Then, under the powers view, K1 particles will get construed as mobile point-located instances of the relevant power to attract K2 particles, and K2 particles will get construed as mobile point-located instances of the relevant power to attract K1 particles. But it is obvious that these construals create a vicious circle with respect to our understanding of the content of the powers involved. For the content of the K1 power is specified by reference to the behaviour, or potential behaviour, of instances of the K2 power, whose content is in turn specified by reference to the behaviour, or potential behaviour, of instances of the K1 power. So there is, in the end, nothing to give any content to either power. Nor, it seems, can we improve the situation by giving these particles further powers

or by envisaging the existence of further particles. For example, we might, as a starter, take K1 particles to be mutually repulsive. But then how is the content of this power of repulsion to be understood? If we try to understand it as, or as in part, the power to repel particles that are, amongst other things, instances of itself, we get an immediate vicious circle. But if we try to understand it as the power to repel particles that are, amongst other things, instances of the K1 power of attraction—the power to attract K2 particles—we still get a circle when we take account of the construal of K2 particles. For such particles are instances of the power to attract K1 particles, and so have to be construed either as instances of the power to attract particles that are, amongst other things, instances of the K1 power of attraction, or as instances of the power to attract particles that are, amongst other things, instances of the K1 power of repulsion, or as instances of the power to attract particles that are instances of both these powers. Again, we might introduce an additional kind of fundamental particle, K3, such that K2 and K3 particles are mutually repulsive and such that K3 particles have the power to destroy K1 particles that come within a certain distance. But this will still leave us unable to assign content to any of the powers involved without going in a circle. For example, the K2 power to repel K3 particles will have to be understood either (i) as, or as in part, the power to repel particles that are, amongst other things, instances of the power to repel particles that are, amongst other things, instances of the K2 power of repulsion, which will be immediately circular, or (ii) as, or as in part, the power to repel particles that are, amongst other things, instances of the power to repel particles that are, amongst other things, instances of the K2 power of attraction, which will yield a circle when we try to spell out the content of this latter power in terms of its effects on K1 particles, or (iii) as, or as in part, the power to repel particles that are, amongst other things, instances of the K3 power of destruction, which will again yield a circle when we focus on the K1 particles that are the target of the destruction. Quite generally, it is clear that by multiplying types of particle and the power relations in which they stand to one another we cannot escape the basic problem that, if each power is simply a power to affect the behaviour of items of power, there is nothing, in the end, that will give any power genuine content. The situation will be like that of a supposed language in which the only source of meaning lies in the meaning relations in which terms and expressions stand to one another. Just as a language can only have meaning if some of its terms or expressions make direct semantic contact with things that lie outside it, so the total system of

physical powers can only have content if at least some of its components are powers to affect something other than the behaviour of items of power.

There are three ways in which a defender of the powers view could try to deal with this problem, though only one of them offers any prospect of success. The first way would be to say that, while all the actual fundamental particles are to be construed as concrete items of power, there are certain uninstantiated types of particle with properties of intrinsic content, and that the powers of the actual particles possess their dynamic content by being, directly or indirectly, powers to affect how particles of these uninstantiated types would behave, if they existed. This response would clearly be perverse and wholly against the spirit of the powers view. Any reason for supposing that intrinsic content features in the story in that way would be a reason for ascribing intrinsic content to the actual particles themselves. The second way of trying to deal with the problem would be to say that certain particle powers are powers to affect human experience and that these experience-affecting powers are what form the source of content for the others. But, for reasons I shall not here elaborate, this approach cannot, in practice, be satisfactorily developed;⁹ and, in any case, it would take us out of the framework of physical realism that we are currently assuming, since, if all particle powers were, directly or indirectly, targeted on to human experience, the existence of the physical world would become logically dependent on facts about human mentality. The third way of trying to handle the problem would be to take some particle powers to be powers to affect the geometrical behaviour of space. This, it seems to me, is the only line of defence which deserves to be taken seriously. We already have one scientifically accredited example of a power of this sort. For in the general theory of relativity the power of gravitational attraction is represented as a power which is exercised not by two objects directly attracting each other, but by two objects directly affecting the geometrical structure of the surrounding space and by the resulting geometrical structure affecting how the objects are disposed to move. And if we are prepared to accept this case, there can be nothing in principle against supposing that there are other ways in which the occupants of space have the power to directly affect its geometry. So one way in which we could try to develop the powers view, without incurring the relevant problem, would be by taking the basic set of particle powers to be powers to affect the geometrical

⁹ For a specification of these reasons see my *The Case for Idealism* (London: Routledge & Kegan Paul, 1982), 70–1.

structure of space, and then trying to define all other powers in terms of how their content directly or indirectly relates to powers in this set. Whether this approach could be made to square with our scientific knowledge of the world, I am not in a position to say; in any case, science has yet to deliver its final theory of the world. But, ignoring any possible scientific objections, I can see no reason why, as a way of dealing with the problem of dynamic content, the approach should not be successful.

However, even if it turns out to have some satisfactory way of assigning content to the particle powers, the powers view is vulnerable to a more fundamental objection. For, as we shall now see, the whole idea of construing a particle as a concrete item of power is incoherent.

I need to begin by dealing with a preliminary matter. Particles are persisting entities: they move through space and preserve their identities through time. For this reason, the powers view, as I have represented it, takes particles to be items of power that are similarly persisting and mobile. But this seems to create a problem. Items of power, in the relevant sense, are not objects which possess power—objects which have their own non-functional qualitative characters and which are disposed to exert certain kinds of causal influence on other things. They are simply concrete instances of power, whose whole natures are spatial and dynamic. Otherwise, we would still be left with intrinsic content. But, even if we can accept an ontology of items of power in this sense, there seems to be a special difficulty in thinking of such items as equipped to persist—to preserve their identities—through time. It seems that the only way in which we could countenance talk of persisting items of power would be as a shorthand for describing situations in which there was a spatio-temporally and causally continuous series of momentary power items of the same power type. And if this is so, it seems that we can only exclude intrinsic content from particles by denying their existence altogether. However, I think that, on this point, the defender of the powers view has an effective response. It is true, I think, that he would do best to restrict his *fundamental* ontology of power items to ones that are necessarily momentary—each such item being tied to a particular point of space at a particular time. But all this means is that in recognizing persisting particles, he has to think of them as entities that are ontologically derivative—entities whose existence is constituted by more fundamental facts about momentary items of power. Specifically, for each particle P, he would have to think of P's existence as constituted by the existence of a spatio-temporally and causally continuous series of momentary power items of the same power

type, together with the obtaining of whatever natural laws it is that make the spatio-temporal and causal continuities involved non-accidental.

The fundamental problem for the powers view is that, irrespective of whether we think of them as momentary or persisting entities, we cannot make sense of there being spatially located items of power at all in the form that the view requires. To see why not, let us focus again on the example of the two kinds of fundamental particle, K1 and K2, that are mutually attractive, the force of the attraction between any two particles varying inversely with the distance between them. But this time, to avoid any problem over the specification of the content of the relevant powers, let us assume, though it is contrary to what the powers view itself requires, that K2 particles have intrinsic content, so that it is only K1 particles that are candidates for the powers-view treatment, and the K1 power of attraction can be fully specified, without circularity, by reference to the potential behaviour of K2 particles. So, the relevant powers-view claim will be that K1 particles are mobile point-located instances of the power to attract K2 particles in the relevant way; or, if we follow what was suggested above, it will be that each K1 particle is something whose existence is constituted by a spatio-temporally and causally continuous series of momentary instances of such a power, together with some appropriate framework of laws. Consider, then, the particular case in which it is claimed that an instance of the relevant power is located at a certain point p at a certain time t . Can we make sense of this claim? And, more specifically, can we make sense of it in the terms required by the powers view, where there is nothing else about that point at that time on which the presence of the power is grounded? Well, we can certainly make sense of the claim that, at t , any K2 particle in existence is subject to the relevant form of dispositional attraction towards p ; and if there is any instance of power which is responsible for K2 particles being disposed in that way, p will centrally feature in its content. But what it is hard to understand is what it would be for p to serve, additionally, as the *locus* of the power. If it were really possible for an instance of the power to be located there, at the point to which the power is a power to attract, it would be equally possible for the same type of power (the power to attract K2 particles to p) to be instantiated elsewhere, for example at a point a mile away from p . Nor would the locus of the power have to be a point: it would be equally possible to envisage an instance of the power extending over a region of space, or over an aggregate of separated regions, or indeed over the whole of space. But as soon as we turn our attention to putative locations other than the point which forms the focus of

the attraction, it is obvious that we cannot make sense of what is envisaged unless we think of the location as characterized by something that accounts for the presence of the power. And what this shows is that we cannot make sense of what is envisaged in the special case where the supposed location of the power is the point which centrally features in its content. We can think of the relevant power as located at p at t if we also suppose this to be the location of a particle with intrinsic content, whose character, combined with the laws of nature, requires K2 particles to be attracted to that point at that time. Or again, we can think of the power as located at p at t if we suppose p itself to instantiate some property at t which, in the framework of the laws, plays a similar—attraction-sustaining—role. But, without anything about p at t to form a ground for the power—something about p at t which in some way accounts for the presence of the power—we cannot get any purchase on the claim that the power is located there. We cannot move beyond the claim that p is the point to which, at that time, K2 particles are dispositionally attracted.

It turns out, then, that, even if we would like to be able to adopt it, the powers view cannot be made to work, because we cannot make sense of the ontology of power items it requires. This means that we cannot avoid accepting that the fundamental constituents of material objects have some form of intrinsic content, even though the nature of this content cannot be empirically discovered, and even though we cannot, within our system of physical concepts, form any conception of what it might be. At least, we cannot avoid this conclusion within the framework of the physical-realist position that we are currently taking for granted.

VII

We are addressing the issue of what we know, or might be able to empirically discover, about the nature of the physical world, and we have so far confined our attention to the nature of its material ingredients. In particular, we have focused on the question of how things stand, epistemologically, with respect to their properties of intrinsic content. These properties are intrinsic properties of the material objects they characterize, pertaining to what the objects are like in themselves, as opposed to their relations to other things and their behavioural and causal dispositions; and, for a given object at a given time, they cover those aspects of the object's intrinsic character that

provide elements of qualitative content for its spatial structure—aspects such as the kinds of stuff of which the object is composed, and any non-functional qualities that pervade portions of its surface or internal regions. The conclusion we have reached is that these content properties are ones of which we have, and can empirically acquire, no transparent knowledge. We can, at best, only identify them opaquely, as those properties, whatever they are, which endow the objects that have them with certain behavioural and causal dispositions. Moreover, we have seen that, as well as their being beyond the reach of transparent knowledge, it is not even possible, in physical terms, to describe, or conceive of, what these properties might be. It is also clear that the considerations which oblige us to accept these forms of inscrutability in the case of the content properties of material objects would lead to analogous conclusions in the case of any other type of space-located physical entity. So, if there are certain kinds of submicroscopic physical particle that it is hard to think of as particles of matter (for example, because they lack mass), or if there is any other form of spatially located ontological item which does not qualify as material in the ordinary sense, we shall equally have to accept the empirical and conceptual inscrutability of their properties of intrinsic content—of what, beyond their intrinsic spatial (or spatio-temporal) properties, these things are like in themselves.

But material objects, and whatever other space-located items there may be, are not the only ontological ingredients of the physical world. There is also the space in which these items occur. So, having identified a crucial limitation on what we know, or might be able to discover, about the nature of the things located in physical space, I want now to turn to the question of what we know, or might be able to discover, about the intrinsic nature of the spatial medium itself. In one respect, of course, this question overlaps the question of the sorts of things we can know about the nature of material objects. For, at any time, the intrinsic character of a material object includes the properties of spatial structure which it derives from the specific form of its space occupancy. So, in investigating the scope of our actual and potential knowledge about the intrinsic nature of space, we are still investigating, in a specific respect, the scope of our knowledge about the nature of such objects.

One thing which it seems that we know, or could come to empirically discover, about the intrinsic nature of physical space is its geometrical structure. It used to be thought that, taking for granted its three-dimensionality, and leaving aside the question of whether it is bounded, the geometrical structure of physical space could be determined a priori, by the application of

Euclid's axioms. It is now generally agreed that this view is incorrect. It has been shown that there are other sets of axioms that are internally consistent, but incompatible with Euclid's set at certain points, and it is hard to see how, other than by empirical tests, we could establish the appropriate set for physical space. Thus, if our empirical measurements of distances in space consistently seemed to indicate its possession of some form of internal curvature (in the way, for example, that the distances on the surface of a sphere indicate its curvature), and if there were no reason, other than adherence to the Euclidean axioms, for thinking that our measurements were systematically erroneous, there would be a strong case for concluding that physical space was non-Euclidean. Moreover, once we accept the possibility that space is non-Euclidean, we must accept the further possibility that its geometrical structure is not, as we intuitively suppose, homogeneous and static, but varies in detail from place to place and from time to time. And this, in effect, is what modern physics, in the form of the general theory of relativity, holds to be the case. For, in giving a purely geometrical account of gravitational fields, the general theory postulates a four-dimensional space–time continuum whose curvature varies from region to region, and indeed from point to point, with the varying distributions of matter. (We have already noted the potential relevance of this to the development of the powers view.)

Even if it goes against current scientific theory, it will be best, for the purposes of our discussion, if we retain the traditional view of physical space as internally uniform and Euclidean, and as having that structure at all times. The philosophical points that I want to make do not require this view: they apply, in the same way, whatever our account of the geometry of space. But they are easier to make within the Euclidean framework. Without the requisite mathematics, it is hard to come to terms with non-Euclidean geometry, and even harder to come to terms with the rippling curvatures of the space–time continuum. For most of us the Euclidean view remains the one that is familiar and accessible, and we should see how things work out in this framework before we contemplate taking on anything more complicated or more obscure.

For ease of discussion, then, I shall work on the assumption that physical space is three-dimensional and Euclidean, and I shall assume that this is something that we have established empirically, by making the appropriate measurements, and, more generally, by observing how things in space behave. But a knowledge of the geometrical structure of space does not amount to a full knowledge of its intrinsic nature. For it does not tell us anything about

the kind of thing that *has* this geometrical structure. In order for the space to be three-dimensional and Euclidean, its constituent points must collectively instantiate a network of distance relations that satisfies the relevant geometrical requirements, and these requirements can be specified in a purely formal and mathematical way. The simplest way of doing this is by saying that, relative to a suitable coordinate system, the distances between points conform to the Pythagorean principle; in other words, for some one-to-one correlation between points and ordered triples of real numbers (intuitively, the triples form the identifying coordinates for the points), and for any pair of points p_1 and p_2 , if the ordered triple correlated with p_1 is $\langle x_1, y_1, z_1 \rangle$ and the ordered triple correlated with p_2 is $\langle x_2, y_2, z_2 \rangle$, then the distance between p_1 and p_2 is equal to the square root of $\{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2\}$. But knowing that the network of distance relations meets these requirements does not tell us the nature of the spatial thing to which the points that stand in these relations belong. One thing which serves to underline this point is that we can think of other types of complex entity that have, or can be represented as having, the relevant form of geometrical structure. To take a simple example, we can represent the set of all ordered triples of real numbers as itself forming a kind of abstract space, with the triples serving as points, and with the distance relations defined by the Pythagorean formula itself (so that the distance between point $\langle x_1, y_1, z_1 \rangle$ and point $\langle x_2, y_2, z_2 \rangle$ is fixed as the square root of $\{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2\}$). Trivially, this will have a three-dimensional Euclidean structure, but, as an abstract mathematical space, its nature is not, in any other way, remotely like that of the concrete space which forms an ontological ingredient of the physical world. Clearly, there must be something further about the intrinsic nature of physical space that its geometrical specification does not capture. It is also clear that this additional factor cannot merely consist in the concreteness of physical space. There must be something further about the nature of the concrete thing that has the geometrical structure—something that covers, as it were, the nature of the spatial material, or fabric, in which the geometrical structure is realized.

Specifying the geometrical structure of physical space does not specify its full intrinsic nature. It does not reveal the nature of the spatial thing which exemplifies the geometrical structure. The questions now are: Do we have any further relevant knowledge which we can add to the geometrical specification and which will make up the deficiency? And if not, can we think of a way in which such knowledge might be empirically acquired? One item

of knowledge that will not serve to shed much light on the situation is that the nature of physical space equips it to be a medium for the existence of material objects. No doubt this knowledge does, in a sense, suffice to uniquely identify the nature of the space, by telling us something about this nature which distinguishes it from anything else. But what it does not provide is a *transparent* identification, which reveals what that nature is. This is not just because it is indirect—because it identifies the nature of physical space by reference to its role in relation to material objects. That, in itself, might not matter if we already had an adequate grasp of the forms of intrinsic content of such objects. For such a grasp, combined with the knowledge that such objects are occupants of the space, might be enough to reveal the space's own qualitative character. But without any transparent knowledge of the intrinsic content of material objects, the most we can learn from the knowledge that the nature of physical space equips it to be a medium for their existence is that its nature equips it to accommodate occupants that are three-dimensional and persisting, and this still leaves us entirely in the dark as to what, beyond its three-dimensional geometry, the medium is like in itself.

In our ordinary, pre-philosophical, thinking, we tend to conceive of spatial properties in visual terms—or, at least, we do so if we are sighted. So, if someone asks me to think of a right-angled triangle, my normal response will be to visualize such a figure in my mind's eye. Or if I am asked to think of two snooker balls in spatial contact, I will normally, again, form a visual picture of the two balls in the appropriate contiguous relationship. And this, I am sure, would be the normal practice of sighted subjects in general. There is obviously something distinctive about this visual way of representing spatial properties. So, one suggestion might be that we should think of it as revealing something distinctive about the intrinsic nature of physical space—something additional to facts about its geometrical structure. In other words, it might be suggested that we should accept two propositions: first, that the visual representation of spatial properties characterizes them in terms that go beyond what is covered by their geometrical specification; second, that the additional (geometry-transcending) elements of this characterization reveal something about the real nature of the properties in their physical form, and so about the intrinsic nature of physical space itself.

The first of these propositions, that the visual representation of spatial properties characterizes them in terms that transcend their geometrical specification, is, in my view, correct. Thus, I do not think that what is distinctive about the visual representation of these properties is simply a

matter of the distinctively visual character of the qualities of sensible colour, with which, in the context of such representation, they combine. I think that, in addition to what is distinctively visual about sensible colour, there is, in the context of visual representation, a distinctively visual way in which elements of colour are spatially arranged—a way that is exactly suited to what is distinctively visual about sensible colour itself. So, I think that when I visualize a triangular patch of colour there is, in addition to the colour, something distinctively visual about the way in which my representation characterizes the extension and shape of the patch. And I think that when I visualize the contact between two snooker balls there is, in addition to the colours in my image of the balls and the table on which they rest, something distinctively visual about the way in which my representation characterizes the contiguity. And, if I am right about this, it is also clear that, by being distinctively visual, these characterizations go beyond what is covered by a mere geometrical specification of the relevant forms of spatial arrangement. But the suggestion we are considering also claims that in the respects in which the visual representation of spatial properties goes beyond their geometrical specification, it reveals something about their real nature in their physical form, and so reveals something, additional to geometrical structure, about the intrinsic nature of physical space. And this, it seems to me, is where the suggestion goes astray. For, on reflection, I think we can see that in so far as spatial properties are accorded a distinctively visual character, there is no rational basis for taking them to be physically realized.

The points that are relevant here involve themes that are familiar from our previous discussion. To begin with, although we may find it natural to ascribe a distinctively visual character to physical space, and to the spatial properties of the things for which it forms a medium, we can see, on reflection, that this is not empirically warranted, since it has no explanatory value with respect to our empirical data. It does not help to explain why the occupants of space behave, with respect to it, in the ways they do: a geometrical specification of space and physical spatial properties provides all the explanatory information of a spatial kind that is available or needed. Nor does it even help to explain why physical spatial arrangement has the forms of visual appearance it does. For the only way in which such arrangement affects the character of our visual experiences is by affecting the transmission of light to our eyes, and here again, as with physical behaviour quite generally, it is only the geometrical aspects of the arrangement that have any explanatory relevance. In short, there is no more empirical justification for ascribing a distinctive visual character to the

spatial aspects of the physical world than for supposing that material objects are characterized by the distinctively visual qualities of sensible colour, or for supposing that any of the categories of secondary quality are physically realized in their sensible form. Moreover, although I accept that there are distinctively visual forms of spatial property—the forms of extension and patterning that feature, with sensible colour, in the content of visual experience—it seems to me that we can only make sense of their distinctively visual character by taking them to be essentially experiential properties, which are incapable of realization, or at least fundamental realization, outside the content of visual experience. In other words, I take the situation here to be the same as the situation I recognized in the case of sensible colour, and, with variation in the sense realm involved, in the case of the other forms of sensible quality on the secondary list. My reasons for taking this to be the situation are also the same—reasons which turn on the two considerations I mentioned in the earlier context, and which I have elaborated in detail in *The Nature of Perception*.¹⁰

There is no rational basis for supposing that physical space has a distinctively visual character, conforming to the distinctive way in which physical spatial properties are visually represented. But, of course, the visual realm is not the only sense realm by which we perceptually register the spatial aspects of the physical world. So, although we cannot look to what is distinctive in the visual representation of spatial properties to reveal something relevant about the intrinsic nature of physical space—something that transcends its geometrical specification—there is still the possibility that we can look to something distinctive in some other sensory mode of representation to provide such a revelation. For example, it might be suggested that our tactual experience of material objects gives rise to a distinctive way of conceiving of their spatial properties—a way that represents these properties in terms that go beyond their geometrical character—and that we can then look to this distinctively tactual mode of conception to give us some transparent insight into what these properties are and what the space in which the objects are located is like in itself. But, even if we can find some other sense realm which offers a distinctive representation of physical spatiality—and the tactual realm would be the obvious candidate—the considerations which show that there is no rational basis for ascribing a distinctively visual character to physical space

¹⁰ *The Nature of Perception*, 130–47.

also show that there is no rational basis for ascribing a distinctively sensible character to it of any sort. Such an ascription would have no empirical warrant, since it would not help to explain the empirical data: whether we are wanting to account for how things behave within the physical world or for how things in the world sensibly appear to us, it would add nothing of explanatory value to the spatial information supplied by a purely geometrical specification of space and the arrangement of things within it. Moreover, if my reasons are well founded for thinking that distinctively visual forms of spatial property are incapable of fundamental realization outside the content of visual experience, they would also serve to show that any distinctively sensible form of spatial property was confined, in its fundamental realization, to the content of sensory experiences of the relevant kind—though this is not something I can appropriately pursue here.

If there is no basis for ascribing a distinctively sensible character to physical space—a character that conforms to the distinctive way in which spatial properties are represented in a particular sense realm—it is clear that we cannot look to ordinary observation to give us any transparent knowledge of the intrinsic nature of space beyond its geometrical structure. This conclusion is parallel to the conclusion we reached earlier, in the case of material objects, that ordinary observation cannot give us transparent knowledge of properties of intrinsic content. The next question, then, and again in parallel with the earlier case, is whether we can hope to do better by appeal to physical science. As I remarked in that earlier context, we normally think of the investigative techniques of science as able to get below the level of ordinary observation and provide an account of the fundamental nature of things. So, should we think of it as equipped to reveal something, beyond a geometrical specification, about the fundamental nature of space?

It does not take long to see that, as in the case of intrinsic content, science has nothing to offer in this area. The techniques of science certainly equip it to probe more deeply than ordinary observation into the geometrical character of space, as well as into the character of the objects and events that occur within it. This is why science has been able to establish that space is not, as ordinary observation would lead us to suppose, Euclidean—though, as I explained, this is something which, for ease of discussion, I have decided to ignore. But the techniques of science do not equip it to uncover anything more about the nature of the thing that has the relevant geometrical character

and forms the medium for the relevant types of object and event. What, in this case, restricts the scope of scientific discovery is the same as what restricts its scope in the case of the nature of material objects, or of any other kind of space-located physical item, and the same as what excludes any empirical warrant for thinking of space as having a distinctively sensible character. The point is that any empirical warrant that the scientist could acquire for ascribing some intrinsic property to physical space would have to derive from the contribution that such an ascription made to the provision of an explanation of his empirical data, and, whatever empirical data his investigation might make available, there is nothing intrinsic that he could think of ascribing, apart from geometrical properties, that would make such a contribution. Even if they could be formulated, hypotheses about the intrinsic nature of space, beyond its geometrical structure, would be as idle for the purposes of scientific explanation as hypotheses about the intrinsic content of physical particles.

It turns out, then, that, like the content properties of material objects, the intrinsic nature of physical space, beyond its geometrical structure, is empirically inscrutable. We can identify this nature opaquely, as that, whatever it is, which, together with its geometrical structure, equips space to form a medium for its material occupants. But what this nature is, is not something that can be empirically revealed, either by the methods of ordinary observation or by the investigative techniques of science. Moreover, as in the case of intrinsic content, the inscrutability has a conceptual aspect too. For, as well as not being able to discover the intrinsic nature of physical space, beyond a knowledge of its geometrical structure, we cannot, in physical terms, even envisage what this nature might be. Any options lie permanently outside the scope of the physical conceptual scheme. Or, at least, they do so, granted that, like the secondary qualities in their sensible form, any distinctively sensible forms of spatial property are incapable of fundamental realization outside the content of sensory experience.

VIII

We have been addressing the question of what sorts of knowledge we possess, or are equipped to acquire, about the nature of the physical world, and have focused our enquiry, in turn, on the case of knowledge about the properties of the things located in physical space—in particular, about the properties of

material objects—and the case of knowledge about the nature of space itself. The situation that has emerged is this:

1. With respect to things located in space, we can acquire knowledge of their spatio-temporal and functional properties, but not of their properties of intrinsic content. So, in the case of material objects, we can acquire knowledge of such things as their shape and size, the internal arrangement of their parts, their positions at a time, and their spatial paths through time. Equally, we can acquire knowledge of their behavioural and causal dispositions, the causal processes in which they feature, and the laws to which their behaviour and modes of interaction are required to conform. But for a given object, or a given object at a given time, we can acquire no transparent knowledge of those ingredients of its intrinsic character that form the elements of qualitative content for its spatial structure. The only way we can hope to empirically identify these elements is opaquely, by reference to the behavioural and causal dispositions which they nomologically sustain.
2. With respect to physical space, we can acquire empirical knowledge of its geometrical structure—a structure which, for simplicity, we have been assuming to be (and to be at all times and in all places) that of a three-dimensional Euclidean continuum. But, beyond a specification of this structure, we can acquire no transparent knowledge of its intrinsic nature—of what the thing which has the relevant structure is like in itself. We can, at best, identify this nature opaquely, as whatever it is that equips physical space to be a medium for material objects. This limitation on what we can know about the nature of physical space means that even with respect to the spatial properties of the things located in it, our knowledge of them is restricted to what can be captured by a geometrical specification, and does not fully reveal what these properties are.
3. With respect to both physical space and the things located in it, those qualitative aspects that are empirically inscrutable are also, in a certain sense, conceptually inscrutable, in that we cannot, in physical terms, conceive of what they might be. At least, we cannot do so, granted that neither the secondary qualities in their sensible form nor any distinctively sensible forms of spatial property have any fundamental realization outside the content of sensory experience.

So far I have reserved the notion of intrinsic content for the inscrutable aspects of material objects and of other things that are located in physical

space. But there is a close analogy between what is inscrutable about the nature of the things in space and what is inscrutable about the nature of space itself, and it will be both convenient and appropriate to extend the notion to cover the latter case as well. So, just as the notion already covers those ingredients of the intrinsic character of a material object at a time that form the elements of qualitative content for its spatial structure, it will be convenient and appropriate to extend it to cover, in addition, that aspect of the intrinsic nature of physical space that transcends its geometrical structure—that aspect that forms the qualitative nature of the thing that possesses the geometrical structure. With the notion thus extended, and allowing the notion of the organization of the world to cover all its various functional aspects, we can express the overall situation succinctly thus: what, in the physical world, is open to empirical investigation are the aspects of its structure and organization; what is empirically inscrutable, and even beyond the reach of physical specification, are its forms of intrinsic content.

There can be no denying that this inscrutability of intrinsic content constitutes a severe limitation on the scope of our epistemic and conceptual capacities with respect to the physical world. It means that, in one obvious sense, we do not know what sort of thing the physical world is, nor even can have any notion, at least in the terms that are ordinarily designed for the characterization of that world, of what it might be. Apart from the topic-neutral facts of structure and organization—facts that do not reveal the nature of the ontological domain to which they belong—the world has turned out to be, as we might put it, a know-not-what. This outcome is likely to prompt surprise, and it may also engender unease. What we must next consider is how we should respond to it.

3

Realism and Phenomenalistic Idealism

I

Even when we set aside the problem of perception, and assume that our sensory experiences give us perceptual, and thereby epistemic, access to the physical world, we have to recognize a severe limitation on the scope of what empirical investigation into the nature of the world is capable of revealing. This limitation affects both our knowledge of physical space and our knowledge of the things located in it. In the case of space, empirical investigation is equipped to reveal facts about its geometrical structure and the ways in which this structure affects what takes place within it. What it cannot reveal, or even help to reveal, is the nature of the thing which has this structure—the nature, as it were, of the spatial material, or fabric, in which the geometrical structure is realized. In the case of the things located in space, empirical investigation is equipped to reveal facts about the ways in which they are structured in space and time—though, with respect to the spatial aspects of this structuring, any revelation will (in line with what can be revealed about space) be confined to what can be covered by a geometrical specification. It is also equipped to reveal facts about their functional properties—about their behavioural and causal dispositions, the causal processes in which they feature, and the laws which govern them. What it cannot reveal, or help to reveal, with respect to any space-located item at a time, are those ingredients of its intrinsic character that provide what we can think of as elements of qualitative content for its spatial structure. So, in the case of material objects, which, from the standpoint of our ordinary conceptual scheme, form the most important category of space-located item, it cannot reveal, beyond a specification of their spatial and dispositional character, the nature of the forms of stuff of which such objects are composed, or of any non-functional qualities that pervade portions of their surfaces or internal regions. I have introduced the notion of *intrinsic content* to cover that aspect of the intrinsic

nature of space that transcends its geometrical character and those aspects of the intrinsic physical characters of space-located items that transcend their spatio-temporal and functional properties. The epistemological situation can then be succinctly expressed by saying that, while the structure and organization of the physical world is amenable to empirical investigation, its forms of intrinsic content—with respect to both physical space and the items that exist within it—are empirically inscrutable.

I have described this limitation on what we can empirically discover as severe. What makes it so is that it leaves our knowledge of the world as, in a certain sense, topic-neutral. However much information we acquire about the structure and organization of the world, the inscrutability of its intrinsic content means that we remain wholly ignorant of the nature of the ontological domain in which this structure and organization are realized, and so, in one clear sense, are left not knowing what sort of thing the world is. Nor, as we saw, can we, in physical terms, even conceive of what sort of thing it might be. For the nature of the relevant forms of content is not only beyond the reach of empirical knowledge, but also beyond the scope of what can be specified within the physical conceptual scheme. Within the system of physical thought, the limits on possible knowledge and possible conception fall at the same point.

II

There is no denying that this limitation on what we can come to know, or even conceive of, about the nature of the physical world is at radical variance with our ordinary, pre-reflective view of the situation. This, at least in large part, is because we ordinarily assume that we can learn much about the intrinsic content of the world from the ways in which it is disposed to sensibly appear to us. There are three aspects to this, which link up with different phases in our discussion in the last chapter. In the first place, we ordinarily assume that material objects are characterized by secondary qualities in their sensible form—the form in which they feature in the content of sensible appearance—and that, in suitable conditions, we can detect the presence of these qualities through ordinary perception. In particular, we ordinarily assume that such objects are typically, in both their surfaces and internal regions, pervaded by various forms of sensible colour, and that we can detect the physical presence of these colours by viewing the

objects, or their relevant parts, at a suitable distance, and in the right kind of illumination. Second, I think we ordinarily tend to assume, with Locke, that our tactual experience of the resistance of material objects to the pressure of our bodies, as when we hold an object in our hands, or push against a barrier, or simply run our hands over a material surface, reveals a fundamental aspect of the intrinsic nature of matter—that aspect which equips it to be a space-filling stuff—and we then take our supposedly transparent grasp of this aspect to give us a basic understanding of the sorts of thing, beyond their three-dimensional extendedness, that material objects are. Third, and most crucially, we ordinarily assume that our visual and tactual experience of the spatial properties of material objects reveals aspects of the nature of those properties beyond what is covered by a geometrical specification, and, in this way, reveals aspects of what, beyond its geometrical structure, physical space is like in itself. So, in feeling the surface of a table, or in seeing the shape of the moon and the arrangement of stars in the sky, we take ourselves to gain insight, beyond mere geometry, into the nature of physical extension, shape, and arrangement, and thereby insight into the intrinsic nature of the spatial medium in which these objects are located and structured. What makes this third point the most crucial is that our conception of physical space, and the spatiality of the things that exist and occur within it, is, along with our conception of time, the most important contributor to our overall view of the world.

The radical difference between our ordinary view of the epistemological situation and what we are obliged to accept by the inscrutability thesis raises a crucial question. Can we accept the new limits on the extent of our physical knowledge without having to abandon our belief in a physical world altogether? Take an analogy. Suppose a biblical scholar argues—it does not matter on what basis—that, apart from his infantile adventure in the bulrushes, everything that the Bible attributes to Moses is pure invention: the things he was said to have experienced or done were, as far as we know, not experienced or done by him or by anyone else. Conceivably, we may find his argument persuasive. But if we do, should we conclude, as the argument itself ostensibly suggests, that there is very little we know about Moses? Or should we simply abandon our belief that Moses really existed? The point of entertaining this second conclusion is that, arguably, with so little left of what we had previously taken ourselves to know about Moses, we are no longer entitled to think that there is any real individual who satisfies what our identifying conception of Moses requires. Likewise, in the case of the physical

world, what we have to decide is whether, as we have been supposing, the inscrutability of intrinsic content means that our knowledge of the character of the world is restricted to facts of structure and organization, or whether it means that we no longer have enough putative physical knowledge to warrant believing in a physical world at all. Should we continue to think, as before, of the inscrutability of intrinsic *physical* content? Or is the limitation on what we can know so severe, in relation to our initial understanding of the situation, that we are no longer entitled to think of the external reality whose content is inscrutable as satisfying what is conceptually required of a physical world?

What seems to provide strong support for the second position is that it is hard to think of our topic-neutral knowledge of this reality as providing all that is needed for an adequate conception of physical space and its material occupants—a conception that measures up to those concepts of space and material object that are central to our physical conceptual scheme and indispensable if we are to hold any physical beliefs at all. The issue about space is the crucial one. If the topic-neutral knowledge provides an adequate conception of physical space, we can use this, I think, to achieve an adequate conception of its material occupants: it will suffice to think of these latter objects precisely as mobile space occupants of unknown intrinsic content, with the spatial properties that relate to their space occupancy, and with certain behavioural and causal dispositions. This conception will not, of course, measure up to our ordinary understanding of the nature of such objects, since that understanding reflects our ordinary assumption that we can gain knowledge of aspects of their intrinsic content from their modes of sensible appearance. But so long as our conception of physical space is adequate, the conception of material objects in terms of it—as mobile three-dimensional occupants with certain behavioural and causal dispositions—is, I think, all that is needed to cover their status as material, and so preserve our entitlement to believe in a physical world. It is our conception of physical space itself that poses the problem. Our topic-neutral knowledge includes, we are assuming, a knowledge of the geometrical structure of the external space, mathematically specified. It is clear that this geometrical knowledge does not, on its own, suffice for an adequate conception of physical space. It does not even cover our conception of physical space as a genuine space, since, as mathematically specified, the geometrical structure of physical space can be exemplified by things that we would not regard as genuine spaces, in the ordinary sense, at all. Thus, as we saw earlier, if we assume that physical

space is uniformly Euclidean, we can represent the set of all ordered triples of real numbers as exemplifying its geometrical structure, by treating the triples as points, and defining distances between them in the appropriate (Pythagorean) way. But apart from its geometrical structure, all we are supposedly left knowing about the character of physical space is that it is something concrete and is equipped to form a medium for the existence of material objects, and, since all we know about the intrinsic nature of these objects is that they are mobile space occupants, with certain specific spatial properties, this does not take us much further. In effect, the only additional insight we gain is that physical space is equipped to form a medium for a category of mobile three-dimensional occupants, the three-dimensional character of these occupants being understood, like the three-dimensional character of the space itself, in purely geometrical terms. This insight may be enough to explain why what we are referring to as physical space qualifies as a genuine space, but it still seems inadequate to sustain that concept of space which is central to our physical conceptual scheme and which we need to be able to apply if we are to retain a belief in the physical world at all.

Even the insight that physical space is equipped to form a medium for a category of mobile occupants needs, I think, to be taken with a certain qualification. For, although it is part of our basic conception of the physical world that it contains material objects, and such objects are, by definition, persisting space occupants, with a capacity for motion, our empirical evidence, and, indeed, any further evidence that we could hope to obtain, is, as I see it, open to the interpretation that the fundamental physical ontology is that of space and instances of certain region-characterizing properties, and that the mobile occupants that feature in our ordinary physical beliefs and in our scientific theories are ontologically derivative entities, whose existence is ultimately constituted by facts about these properties. Thus, suppose that within the domain of mobile occupants, K is the class of types of such occupant that we are required by the scientific evidence to treat as relatively fundamental (so that facts about other types can be represented as ultimately constituted by facts about them), and, for convenience, let us assume that the occupants of these types are known to be spherical particles, of unvarying size and dispositions. Then, as I see it, there is nothing either in our concept of a physical particle or in the empirical evidence to exclude, or count against, the hypothesis that, for each K -type T , there is a one-place property P , which is non-geometric and non-functional, and whose domain of instantiation is restricted to T -sized spherical regions at times, such that the spatio-temporal

distribution and functional organization of P exactly simulates the situation of there being mobile items of P instantiation with the spatial and dispositional character of T particles, and such that the existence of the T particles is precisely constituted by this distribution and organization. If this is right, then although it does not undermine our assumption that physical space forms a medium for certain kinds of mobile occupant, we have to leave room for the possibility that such occupants are not ingredients of the fundamental physical reality, and this means that the knowledge that space is equipped to accommodate them tells us less about its nature than we might have supposed. It is even arguable that we are obliged to think of the mobile occupants of space as ontologically derivative entities, in the sort of way envisaged, in order to make sense of their existence at all. But this is not a point that I shall pursue.

It might now be said that, although our topic-neutral knowledge of the external reality does not, on its own, furnish us with an adequate conception of physical space and its material occupants—the kind of conception that we need if we are to be entitled to retain our belief in the physical world—this point is of no great consequence. For we should not expect an adequate conception of these things to come solely from a knowledge of their intrinsic nature and the way in which they relate to each other. This is because our identifying conception of the physical world, and of its spatial and material components, is partly in terms of its relationship to *us*. For all we know, there may be other concrete realities, distinct from the physical world, but of the same general intrinsic nature—realities with the same kinds of space, space occupant, and organization. Given any such reality, what disqualifies it from having any claim to be the physical world is not that it lacks the appropriate character, but that it does not have the right kind of connection with human mentality to form a world for us. In particular, it is not functionally linked with human mentality in ways that equip us to have epistemic access to it through our senses or practical access to it through our wills. So, it might be said that what provides an adequate conception of physical space and its occupants is a combination of our topic-neutral knowledge of their nature and relationship and a representation of them as appropriately related to human mentality—related in ways that make them the space and space occupants of our world.

I certainly do not want to deny that it is implicit in our identifying conception of the physical world that it is related to our mentality in a way that makes it our world, or a world for us. Indeed, the recognition that the

world has to be a world for us, and the implications of this for our ultimate understanding of its metaphysical status, will be central themes of my own approach. But I do not think that this point takes care of what seems to be the distinctive inadequacy of our topic-neutral conception of physical space and its material occupants. For the apparent inadequacy of this conception is an inadequacy in what it tells us about the *character* of the space and these occupants. In effect, the problem is that, by its topic neutrality, this conception fails to characterize these items in distinctively physical terms at all, and, even when we add in the relevant connections with human mentality, this makes it difficult to think that there is enough here to sustain those concepts of space and material object that are central to our physical conceptual scheme, and without which we cannot preserve our belief in a physical world.

One thing that brings out this problem very clearly is the fact that, by concealing its forms of intrinsic content, our topic-neutral knowledge of the external reality leaves us free to entertain the hypothesis that this reality is, in substance and character, purely mental—that the ontological domain in which the empirically discernible forms of structure and organization are realized is a domain of minds and the things that exist and occur within them. Of course, there could be independent reasons why this mentalistic hypothesis should be excluded. The idea of such a mental reality might be in conflict with our philosophy of mind: it would not, for example be available to someone who accepts a reductive or physicalistic account of the mind. And, even if the hypothesis is compatible with our philosophy of mind, we might still think that in practice the mental realm does not have the ontological and qualitative resources to provide a reality with the relevant structure and organization. We might think, for example, that there is no type of mental item, or complex of items, that could serve as a three-dimensional space. But, while important or of interest in themselves, these issues about the nature of the mind and the resources of the mental realm are not crucial in the present context. All that here matters is that, whatever further considerations may affect its ultimate availability, the mentalistic hypothesis is not excluded by the content of our topic-neutral knowledge. And the reason why this matters is that even when we accept that there could be a mental external reality of the relevant structural and organizational sort, it is still very difficult to accept that the physical world itself could be such a reality. And so, irrespective of whether there could be a reality of that sort, it is hard to accept that the conception of physical space and material objects provided by our topic-neutral knowledge is adequate in the relevant sense.

To illustrate the point, let us, for the sake of argument, suppose that it is logically possible for there to be a sense field with the geometrical structure of physical space—a sense field whose positions are defined by triples of values along three continuous dimensions, and whose positions, thus defined, stand in the right relations of field topology and distance to realize the relevant geometrical structure. And let us suppose that, by envisaging the existence of a non-human mind in which such a sense field concretely exists and persists, and by envisaging, in respect of this field, a suitable field–time distribution and organization of region-pervading sense qualities, we can come to envisage a situation in which there is a mental reality with the structure and organization of the physical world. This will involve—in line, as I have argued, with what the empirical evidence allows—thinking of the fundamental physical reality as one in which persisting space occupants have been replaced by region-characterizing properties that are spatio-temporally distributed and functionally organized in an appropriate way. So far, so good. But can we go on to envisage that, by having the appropriate functional links with human mentality—in particular, by being disposed to causally affect human sensory experience and to causally respond to human volition in all the appropriate ways—this external reality could actually turn out to be the physical world? Would it make sense to suppose that when we move from place to place in the physical world, we (our bodies) are moving around in some alien being’s sense field, and that when we encounter familiar material objects and observe their properties, what we are fundamentally encountering and observing are elements and aspects of this being’s experiential condition? It is hard to think that it would. Even when there are assumed to be no problems on the mental side with envisaging a being whose experiential life is of the relevant kind, it is hard to think that we have the genuine option of supposing that the world that forms the target of our ordinary physical beliefs—the world to whose entities we are trying to refer, and whose states of affairs we are trying to record, in our ordinary physical assertions—is a component of such a life. But if we do not have this option—given the assumptions—then there must be some inadequacy in the topic-neutral conception of the physical world, since this conception allows for that option; and presumably this inadequacy must pertain to the topic-neutral conception of physical space and its occupants.

It might be replied that the only reason why we find difficulty in entertaining the idea of the physical world’s being something mental is that, even in our reflective consideration of the issue, we have not properly

detached ourselves from the perspective of our ordinary ways of thinking. As I have stressed, our ordinary understanding of the character of the world is heavily conditioned by the ways in which types of physical situation sensibly appear to us, and this conditioning affects our beliefs about intrinsic content, as well as those about structure and organization. In particular, we ordinarily conceive of the spatial character of the world not just in the geometrical terms that feature in its topic-neutral specification, but also in terms of what we take to be additionally revealed by the distinctive ways in which physical spatial properties sensibly appear to us in the context of visual and tactual perception. Given its dominant role in our ordinary thinking, it is hardly surprising that this phenomenally conditioned conception of the world should influence our initial judgement when we try to evaluate the adequacy of the topic-neutral conception and the intelligibility of a mentalistic account. It is hardly surprising that we should initially feel that the topic-neutral characterization of physical space and its occupants, which preserves no trace of the phenomenal perspective, does not measure up to what our actual concepts of these items require, and hardly surprising that we should initially find it hard to make sense of the suggestion that the ultimately correct way of understanding the nature of the physical world might be by thinking of it as contained within the experiential life of some alien mind. The remedy, it will be said, is simply to remind ourselves that only the structural and organizational aspects of the external reality are empirically accessible, and that the phenomenal perspective of our ordinary thinking needs to be discounted. Once we are clear about the epistemological situation, and consciously distance ourselves from the misleading aspects of our ordinary outlook, the doubts over the adequacy of the topic-neutral conception of the world and the intelligibility of a mentalistic construal will disappear.

But this reply misses the point, because it fails to get to grips with the real issue. There is no denying that the reason why the topic-neutral conception of the physical world seems *prima facie* inadequate, and why the mentalistic construal seems *prima facie* unintelligible, is that they are radically out of line with our ordinary understanding of the situation—an understanding which is crucially shaped by the ways in which things sensibly appear to us. And it is also true that, if the reasoning in the previous chapter is sound, there is no warrant for accepting the veracity of such sensible appearance beyond what it purports to reveal about structure and organization. But the crucial issue remains of whether, when we detach ourselves from the perspective of our ordinary thinking, and accept that only facts of external structure

and organization are empirically accessible, we can continue to think of the external reality as satisfying the conceptual requirements of a physical world. And it is here that the real challenge to our belief in a physical world arises. It is true that we can, in certain respects, detach ourselves from this ordinary perspective without calling in question the status of the external reality as a physical world. We have no difficulty, for example, in accepting a Lockean view of the secondary qualities, without this affecting our ability or entitlement to think of the items whose possession of these qualities is at issue as genuinely physical. So, we can move from the naive position, which ascribes sensible colours to external objects, to the scientifically enlightened position, which only credits them with dispositions to look sensibly coloured, without jeopardizing the status of these objects as material ones. Similarly, we have no difficulty in retaining our belief in a physical world when we accept that what pass, under ordinary observational scrutiny, as fully materially occupied volumes of the external space turn out, under microstructural scientific analysis, to be largely empty. We can even entertain the idea of all persisting space occupants being ontologically derivative entities, whose existence is constituted by the spatio-temporal distribution and functional organization of certain region-characterizing properties. These departures from our ordinary ways of thinking do not threaten our belief in the physical world, because they do not threaten our belief in physical space, and they do not threaten our belief in physical space because, while revising our ordinary assumptions about the sorts of thing that exist and occur within it, they do not require us to abandon our ordinary conception of this space itself. Admittedly, there are further scientific considerations that seem to indicate that we should revise the traditional view of physical space as uniformly Euclidean—considerations which seem to indicate that the space is subject to internal forms of curvature that vary from place to place and from time to time. But even this revision would not affect our ordinary understanding of the basic nature of physical space, and, in particular, would not affect those aspects of that understanding which derive from the distinctive ways in which physical spatial properties sensibly appear to us in the context of visual and tactual perception. The trouble with the topic-neutral conception of the physical world is that, except in respect of geometrical structure, it discards our ordinary understanding of the nature of physical space altogether, and in doing so it loses contact with anything that we can straightforwardly recognize as physical at all. It may well be that there is no empirical warrant for supposing that the external space conforms

to that understanding; certainly, that is how things seem to stand as the result of our previous discussion. But that does not remove the difficulty of thinking that, without conformity to that understanding, the space will meet the conceptual requirements of a physical space, and the external reality will qualify as a physical world.

III

If intrinsic content is inscrutable, the only knowledge we can have of the external reality is topic-neutral: it is knowledge confined to facts about the structure and organization of this reality, and about its links with the realm of human mentality; it is knowledge that does not reveal anything about the nature of the ontological domain in which this structure and organization are realized. This has been seen to create a problem. For it seems that knowledge of that topic-neutral sort would not provide what is needed for an adequate conception of physical space and its material occupants. It seems that it would not provide a conception that measured up to those concepts of space and material occupant which are central to our physical conceptual scheme and which we need to be able to apply if we are to retain a belief in a physical world at all. The problem has presented itself in two ways. In the first place, a topic-neutral conception of physical space and its occupants strikes us as deficient in its own terms, just by the very paucity of what it covers about the nature of these items. This ostensible deficiency presumably stems from the fact that the conception omits all the distinctively phenomenal aspects of how we ordinarily conceive of them—all that we ordinarily suppose ourselves to know about their properties of intrinsic content from the ways in which things sensibly appear to us. Most crucially, the conception omits all that we ordinarily take ourselves to know about the nature of the spatial medium from the ways in which properties of physical extension and arrangement sensibly appear to us in the context of visual and tactual perception. Second, one consequence of the fact that the conception does not cover any aspects of intrinsic physical content is that, given suitable assumptions about the nature of the mind and the resources of the mental realm, it leaves us free to entertain the hypothesis of the world's being, in substance and character, purely mental. But, even granted the relevant assumptions, we find it hard to accept that such a mentalistic account of the world is a genuine option.

How we should respond to this problem is not, I think, at this stage clear. Certainly, there is no denying the difficulty over thinking of the topic-neutral conception of physical space and its occupants as adequate for the retention of a belief in a physical world. At the same time, it is hard to accept the conclusion that our belief in such a world needs to be abandoned. But before we consider this issue any further, we need to notice that the inscrutability thesis creates an additional problem, and one which is, in a sense, more far-reaching.

So far we have been assuming that, whatever its intrinsic content, the external reality is a spatial one—a reality whose ontological ingredients are a three-dimensional space and the things that exist and occur within it—and that this is something that the empirical evidence reveals. The issue over space has simply been whether our topic-neutral knowledge of the reality yields a conception of space that measures up to what is needed for belief in a physical world. But when we reflect on the matter, I think we can see that, if the empirical evidence does not reveal anything about the intrinsic content of the putative external space—about the intrinsic nature of the spatial medium in which the relevant geometrical structure is supposedly realized—it does not even reveal that there is an external space at all. What it reveals is neutral between the hypothesis that the external reality contains a genuine space and the hypothesis that, while not genuinely spatial, its ingredients are organized in a way that functionally simulates their location in a space.

The point I have in mind will apply whatever precise form we suppose the empirical evidence to take, but it will be helpful if we focus on a concrete case. Let us suppose that, within the framework of our ordinary belief in a physical world, and our ordinary recognition of an ontology of persisting space occupants, the empirical evidence, properly scientifically evaluated, suggests that, at its fundamental level, the physical reality consists of a three-dimensional space, certain types of mobile spherical particle distributed over space and time, causal relations between elements in this distribution, and certain laws of nature governing the behaviour and causal activities of the particles. We have already seen how, in this kind of case, the empirical evidence is equally amenable to the interpretation that the fundamental physical ontology is simply that of space and instances of region-characterizing properties, and that corresponding to each type of particle there is a region-characterizing property whose spatio-temporal distribution and functional organization simulates the situation of there being mobile items of property instantiation of the relevant particle type. What we now have to recognize is that the empirical evidence is also amenable to what we might think of as the mirror-image interpretation,

which rejects the existence of an external space and takes the fundamental external ontology to consist of a category of persisting items that are organized, with respect to the states that occur within them and their modes of interaction, as if they were mobile occupants of a space. Let us speak of the framework of our ordinary belief in a physical world, and our ordinary recognition of an ontology of persisting space occupants, as the *standard doxastic framework*. Then, set out in detail, the interpretation I have in mind is this:

1. There is no concrete space, at least none at the level of what is fundamental. Instead, the external reality, at this fundamental level, is ontologically composed of a stock of persisting entities—I shall refer to them as *P-entities*—grouped into a range of different intrinsic types. The number of types is the same as the number of types of fundamental physical particle that the empirical evidence suggests within the standard doxastic framework.
2. There are three continuous qualitative dimensions X, Y, and Z, and there is a certain class K of complex states (type states) such that each K state is defined by a unique triple of values drawn, respectively, from X, Y, and Z, and such that each such triple defines a unique K state. It is logically necessary (logically guaranteed by the nature of P-entities and the nature of K states) that, at any time, any P-entity that exists at that time is in one and only one K state.
3. Because they are defined by triples of values along three continuous dimensions, we can represent K states as the points of an abstract three-dimensional space. Then, relative to such a representation, and to a one-to-one correlation of types of P-entity with types of particle, there is an exact match between what, in the standard doxastic framework, the empirical evidence suggests about the physical particles in respect of their positions (the positions of their centres) in physical space at times and what holds in the case of P-entities in respect of the realization of K states at times. This exact match covers not only what the evidence suggests about the actual spatio-temporal distribution of the particles, but also what it suggests about the functional facts that relate to such distribution—about the causal relations between elements in the distribution, about the laws that govern the character of the distribution, and about the behavioural and causal dispositions that the particles have in the framework of these laws.

In this way we can envisage a form of external reality that is entirely consonant with the empirical evidence, but in which the ontological spatial medium

that we would ordinarily think of the evidence as systematically reflecting is replaced by the quasi-spatial organization of the relevant persisting entities. It is also clear, I think, that the method I have employed, in constructing this non-spatial interpretation of the evidence, could be adapted to provide analogous forms of non-spatial interpretation for other assumptions about the kind of fundamental physical reality that the evidence, taken in the standard doxastic framework, suggests.

I said that we could think of this non-spatial way of interpreting the evidence as the mirror image of the interpretation which takes the fundamental physical ontology to be that of space and instances of region-characterizing properties, and it is not difficult to see what I had in mind. The standard interpretation of the evidence, exemplified in our actual physical beliefs, is of an external reality that is ontologically composed of both a three-dimensional space and a stock of persisting space occupants. The two alternative interpretations postulate external realities that deviate from this ontological composition in opposite ways—the one retaining an external space, but replacing its persisting occupants by the quasi-occupant organization of certain region-characterizing properties, the other retaining a domain of persisting objects, but putting the quasi-spatial organization of these objects in place of the ontological space. But in one crucial respect the two non-standard interpretations are not analogous. The first one, at least as it was initially conceived, purports to offer an alternative—and equally empirically accredited—account of the fundamental nature of the *physical world*. Thus, the external space that it recognizes is supposed to be the same space—physical space—as features in the standard interpretation, and this, in turn, allows it to represent the material occupants that feature in that interpretation as ontologically derivative entities, whose existence is constitutively sustained by the spatio-temporal distribution and functional organization of the relevant properties. But the second interpretation does not purport to offer an account of the physical world. It does not represent the P-entities that it takes to be subject to the quasi-spatial organization as the same as the physical particles that feature in the standard interpretation. It does not represent these entities as physical at all. And this is hardly surprising. For there cannot be a physical world without a physical space, and the interpretation explicitly takes the fundamental external reality to be non-spatial.

It might be suggested that we could eliminate this disanalogy between the two non-standard interpretations by developing the second one to a further

stage. In the first interpretation, which excludes material objects from the *fundamental* physical ontology, these objects are not excluded from the physical ontology altogether, but represented as ontologically derivative entities, whose existence is constitutively sustained by the distribution and organization of the relevant properties in the way explained. So, the suggestion might be that, under the second interpretation, physical space could be represented as something ontologically derivative, whose existence is constituted by the quasi-spatial organization of the P-entities, and that these entities themselves could then be identified with the physical particles that feature in the standard interpretation. So, just as on the first interpretation we are left with a fundamental physical space and a derivative stock of physical particles to be its occupants, so, on the expanded version of the second, we would be left with a fundamental stock of physical particles and a derivative space for them to occupy.

But the interpretation cannot be developed in this way. For, even if we could legitimately think of physical space as deriving its existence from the quasi-spatial organization of the P-entities, we could not legitimately identify the P-entities with physical particles. Physical particles, like all physical space occupants, are, in their very being, occupants of physical space and logically incapable of existing in any other form. If the space in which they are located is taken to be ontologically derivative, these particle occupants have to be taken to be ontologically derivative too, deriving their existence from the same fundamental factors that supposedly sustain the existence of the space. And once we recognize the particles as ontologically derivative, we obviously cannot identify them with the P-entities whose existence is stipulated to be fundamental. What also needs to be stressed is that even taking physical space to be something ontologically derivative is not an available option in the framework of physical realism within which we are currently working. For if physical space were ontologically derivative, the physical world itself would be something whose existence and character were constituted by factors of a different kind, and this would be in direct conflict with the realist claim that the world has an existence which is philosophically fundamental.

We can see now why the inscrutability of intrinsic content creates a further problem. The problem earlier identified was that, by not revealing anything about the nature of the ontological domain in which the external structure and organization are realized, a topic-neutral knowledge of the external reality does not seem to yield an adequate conception of physical space and its material occupants—a conception that measures up to those concepts of

space and occupant that are central to our physical conceptual scheme and that we have to be able to apply if we are to retain a belief in a physical world at all. And I highlighted this problem by drawing attention to the fact that such topic-neutral knowledge is compatible with the hypothesis that the external reality is, in substance and character, purely mental. The new problem that has emerged is that, if intrinsic content is empirically concealed, we cannot even think of the empirical evidence as revealing that the external reality has the kind of ontological structure that is required of a physical world. For we cannot think of it as revealing whether the external reality is genuinely spatial or merely quasi-spatially organized. We can also see why, in a certain sense, this new problem is more far-reaching. For it seems to mean that even if we could bring ourselves to accept that, despite its thinness, the topic-neutral conception of physical space and its occupants is adequate for our purposes—that it is all that belief in a physical world requires—we have no way of telling whether the external reality is structured and organized in a way that satisfies that conception, or even comes close to satisfying it.

It might be objected that, in the form in which it is gathered and has to be evaluated, the evidence that I am claiming to be neutral between a spatial and a non-spatial account of the external reality is explicitly about how things stand in the physical (spatial) world, and so—at least in the framework of physical realism—presupposes that a spatial account is correct. But we have to be careful here. It is certainly true that, as it is initially conceived, and, indeed, as it has to be initially conceived, the subject matter of the relevant evidence is physical. So, in the case on which we have focused, the evidence which supports the acceptance of the relevant particle theory when taken in the standard doxastic framework has also been acquired through scientific investigation conducted within that framework, and consists in items of putative knowledge of the properties and behaviour of things in the physical world. But this does not mean that when we consider what the evidence, understood in those physical terms, may indicate, we are barred from entertaining the hypothesis that what ultimately lies behind it is a non-spatial reality of the kind envisaged, and that the evidence itself needs to be ultimately understood in a correspondingly different way—as aspects of the spatial fashion in which this non-spatial (but quasi-spatially organized) reality empirically appears to us in the context of our ordinary conceptual scheme and system of physical beliefs. This kind of reinterpretation of the evidence is something that is already needed within the context of physical science itself, where the putative observational information which provides science

with its ultimate source of empirical evidence represents the world as having, even with respect to its spatial structure, a very different character from that ascribed to it by the microphysical theories that that evidence is taken to support, and where, in consequence, things which, in the initial evidential context, are taken to be aspects of the actual character of the world have to be ultimately construed as aspects of how a very differently characterized world observationally appears. And the same kind of reinterpretation of the evidence would also be needed in the case where, compatibly with physical science, but going beyond it, we envisage a fundamental physical ontology of space and region-characterizing properties. In these cases, of course, the change in the understanding of the nature of the evidence does not affect its physical status: the evidence continues to be represented as aspects of the way in which a physical (spatial) reality empirically presents itself to us in the context of our ordinary system of beliefs. But the fact that the reconstruction of the evidence in the case of the non-spatial hypothesis is, in that respect, more radical does not show that it cannot be entertained.

It might still be said that, whatever ways of reconstructing the evidence are theoretically available, the existence of the physical world is something already clear to us, which we are entitled to take for granted. Well, in the end I shall not disagree with this, as we shall see. But the point remains that, in the realist framework in which we are currently working, and with the inscrutability of intrinsic content established, it is hard to see how that claim of entitlement can be defended. Once we have conceded that we have no knowledge of what, beyond its geometrical structure, the putative external space is like in itself, there seems to be no way of avoiding the conclusion that, for all we know, the external reality might be one that gives rise to the appearance of spatiality at our viewpoint through its quasi-spatial organization, rather than by containing a genuine space. And a reality that did not contain a genuine space—a genuine three-dimensional space—would not meet the requirements of a physical world.

IV

Given the inscrutability of intrinsic content, it is hard to see how we can preserve our entitlement to believe in a physical world. On the one hand, a topic-neutral conception of the world—a conception that entirely leaves open the nature of the ontological domain in which the physical structure

and organization are realized—does not seem to measure up to all that is conceptually required of such a world. In particular, it does not seem to measure up to all that is conceptually required of a physical space and its material occupants. On the other hand, even if we were to allow the adequacy of a topic-neutral conception of the world, it is hard to see how, without any knowledge of intrinsic content, we could be entitled to believe that the ontological structure of the external reality meets the requirements of a physical world. For it is hard to see how if we cannot detect anything about the intrinsic content of physical space—anything about the nature of the spatial medium in which the geometrical structure is realized—we could be equipped to tell whether the appearance of a spatial world at our viewpoint is the reflection of a reality that is genuinely spatial or merely quasi-spatially organized.

Should we, then, conclude that the belief in a physical world is no longer philosophically tenable? Well, before we contemplate such an unpalatable outcome, there is one other approach that needs to be explored.

The problems that have arisen, have arisen within the framework of a realist conception of the world—a conception that takes the world to be ontologically independent of the human mind and to be something whose existence is philosophically fundamental. It was in this framework that we pursued the epistemological enquiry of the previous chapter—the enquiry that yielded the inscrutability thesis—and it is this framework that obliges us to regard respects in which we are ignorant of the nature of the external reality as respects in which we are ignorant of the nature of the physical world. One possibility we need to consider, then, given the difficulties we now face, is that it is the framework of realism itself that is at fault. Perhaps the right way of responding to the problems of inscrutability is not to give up our belief in a physical world, but to give up our realist view of it, and develop, in its place, a view of the world for which these problems do not arise.

This suggestion brings us back to an issue that has already surfaced in our earlier discussion, but not yet been resolved. We saw, in Chapter 1, that within the framework of physical realism we cannot achieve a satisfactory account of physical-item perception, since such a framework does not allow us to understand how the physical world, which it takes to be ontologically independent of the human mind, can fall within the scope of our perceptual awareness. A suggestion that was then put forward for consideration, but not at that stage evaluated, was that instead of conceding that the physical world is perceptually inaccessible, we should replace our realist view of it by a certain

form of idealism. The form of idealism in question was one that takes the physical world to be something whose existence is constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. This would enable us to make awareness contact with the physical world through sensory experience, since it would be by facts about such experience, either on their own or as part of a richer complex, that the world was constitutively created. If we are thinking of abandoning realism in response to our present difficulties, we shall obviously need to consider whether this same form of idealism offers a possible solution, and, if it does, whether this is the right position to adopt.

We are now moving into the central area of our whole discussion, which focuses on the issue between physical realism and the relevant form of idealism. To prepare the ground for a proper investigation of that issue I need to begin by setting out, in more detail, precisely what these rival views of the world involve.

V

I have represented physical realism as the conjunction of two claims. The first claim is that the physical world is ontologically independent of the human mind, and I shall refer to this as the *independence claim*. The second claim is that the physical world is something whose existence is philosophically fundamental, and I shall refer to this as the *fundamentalist claim*. (This latter claim must not, of course, be confused with what I earlier labelled the *fundamentalist view* in my discussion of the topic of perception.) To understand the realist position, then, we need to get clear about the meanings of these two claims and what they do and do not imply.

In claiming that the physical world is ontologically independent of the human mind, the realist is claiming that its existence is logically independent of facts about human mentality, and this is equivalent to claiming that such facts do not logically contribute to its existence. By 'facts about human mentality', I do not just mean facts that cover actual instances of mentality—facts such as that a certain subject had a certain experience at a certain time, or that a certain subject held a certain belief over a certain period. I also mean to include facts about how the realm of human mentality is organized, and about what is naturally necessary or probable or possible in that realm as a result. So, if someone thought that the physical world logically

depended, for its existence, on certain organizational aspects of the realm of human sensory experience, then, even if he took that to be the only respect in which human mentality logically contributed to the world's existence, his position would be in conflict with the realist's independence claim. It should also be noted that, in taking the world's existence to be logically independent of facts about *human* mentality, the realist is also, in effect, committed to taking its existence to be logically independent of facts about the mentality of any subject or group of subjects that are physically embodied. So, if certain types of non-human animal are endowed with minds, the realist is committed to saying that facts about their mentality do not logically contribute to the world's existence. This said, it is only the application of the independence claim to human mentality that is of any interest, since it is only in this area that there could be any serious issue about whether the claim is correct.

In taking the world to be something whose existence is logically independent of facts about human mentality, the realist is, with one qualification, thinking of this ontological independence as applying to the world with respect to every part of its spatio-temporal spread. So, if someone accepted that, in general, the world's existence is logically independent of facts about human mentality, but still insisted that such facts logically contribute to its existence over a certain period in its history or in a certain portion of its spatial extent, then, subject to the relevant qualification, his position would be in conflict with the independence claim as the realist intends this to be understood. The need for a qualification arises as follows. Realism allows for the possibility of mental events causing, or contributing to the causation of, physical events, and, although this may be contrary to the actual laws of nature, it allows for the possibility of a human mental event causing, or being part of what causes, an increase in the spatial extent of the universe or its persistence beyond a certain time. Now a *causal* contribution to the occurrence of an event is not, as such, a *logical* contribution. But it could still be thought that, where an event is an effect (as it normally is), its identity logically depends on the identity of its cause (so that that particular event logically could not have occurred without that particular cause), and that, where an effect is that of the coming into existence of a certain item, the identity of that item logically depends on the identity of the cause (so that that particular item logically could not have existed without the event of its coming into existence having that particular cause). This would mean that, if a human mental event were to be responsible, or part of what was responsible, for causally generating a certain spatial portion or temporal phase of the

universe, it would logically contribute to the identity of that portion or phase, and, in that respect, logically contribute to the existence of this portion or phase. Since we do not want to define physical realism in a way that excludes this view about identity, we need to attach an appropriate qualification to the scope of the independence claim. We need to take the claim to be merely: the physical world is something whose existence, throughout its spatio-temporal spread, is logically independent of facts about human mentality, apart from the respects, if any, in which the causal contribution of human mentality to the existence of certain portions or phases of the universe logically affects the identities of these portions or phases. Having made this point, I shall, for convenience, normally omit the qualification in future formulations of the claim.

In insisting that facts about human mentality make no logical contribution to the *existence* of the physical world—or to the existence of any spatial portion or temporal phase—the realist is not committed to saying that they make no logical contribution to its *character*. He is free to acknowledge, what is surely uncontroversial, that there are many physical properties and types of physical object whose instantiation logically depends on a contribution from human mentality. For example, he is free to acknowledge that the colour of a material object consists in, or logically involves, its disposition to look coloured in a certain way to the normal human percipient who views it in standard conditions, and that something's being a clock logically involves its having the potential to indicate the time to a human subject who understands its chronometric significance; and, clearly, there are a host of further examples of these and similar kinds. All that the realist is required to say, by his insistence on ontological independence, is that facts about human mentality do not logically contribute to the character of the world *in its primary core*. To express this requirement more precisely, let us say that a fact is *physically fundamental* if and only if it is purely physical and is constitutively basic relative to the physical realm—in other words, if and only if it is purely physical and is not wholly or partly constituted by other physical facts. Let us also say that a physical fact is a *core* physical fact if and only if it is either physically fundamental or wholly constituted by facts that are physically fundamental. Finally, let us say that, within the total physical reality (of all that physically exists or obtains), the *core physical reality* is that portion which is covered by the totality of core physical facts. Then what the realist is required to say is that human mental facts do not logically contribute to any aspect of the character of the core physical reality. There is room for

different views about the kinds of ingredient that feature in the core physical reality. I shall assume—I hope uncontroversially—that this reality at least includes: space and time; the location, within space and time, of various kinds of material object and object-involving event; causal relations between different object-involving events; and certain laws of nature that govern the occurrence of these events and the causal relations between them.

There are two final points to be made about this. First, in presenting the realist as committed to saying that facts about human mentality do not logically contribute to the character of the core physical reality, I am taking the *character* of this reality to be something that is neutral with respect to the identities of the physical particulars that feature in it. So, for example, if the reality includes the fact, or state of affairs, of certain particles being spatially arranged in a certain way at a certain time, the relevant aspect of its character will cover the natures of these particles and the nature of their arrangement, but will not cover their identities. This means that, unlike the earlier case, I do not need to attach a qualification to the relevant claim of no logical contribution in order to take care of any complication that may arise from the occurrence of mind-to-world causation. Thus, in his commitment to saying that facts about human mentality do not logically contribute to the character of the core physical reality, the realist is still free to suppose that human mentality causally contributes to the existence or occurrence of some of the particulars involved in that reality, and thereby logically contributes to their identities.

Second, although the point about the identities of physical particulars is taken care of, there is still one way in which I need slightly to amend my account of the realist's position. Thus, instead of representing the realist as committed to claiming that facts about human mentality make no logical contribution to the character of the core physical reality, I should, to be strictly accurate, represent him as committed to the weaker claim that such facts make no such contribution to the character of this reality *in their capacity as mental facts*. The reason why this qualification is needed is that realism leaves open the option of adopting a physicalistic account of psychological properties—an account which represents such properties as ones that are to be ultimately understood as physical properties. If the realist were to adopt such an account—for example, by embracing analytical behaviourism or the type version of the psychophysical identity theory—he would inevitably have to acknowledge that certain facts about human mentality logically contribute to the character of the core physical reality simply by being themselves aspects

of that character. What he could still insist, and what, as a realist, he would have to insist, is that any human mental facts that logically contribute to the character of the core physical reality in that way do so only in their capacity as *physical* facts—as facts understood in physical terms—and not in their capacity as aspects of human psychology. Hence the point of the qualification. I might add that, in my view, any physicalistic account of psychological properties is hopelessly implausible, as I have tried to show elsewhere.¹ But that does not mean that I am entitled to define the realist's position in a way that excludes it.

This said, the qualification will have no relevance to the issues that I want to discuss, and, as with the earlier qualification, I shall tend to omit it in future formulations of the realist's position.

So much, then, for the realist's independence claim. Let us now turn to the second of his claims, that the physical world is something whose existence is philosophically fundamental. How is this to be interpreted?

I have already indicated that this fundamentalist claim is to be understood as excluding the reduction of the physical realm to something else. But I need to begin by noting that there are two general types of potential reduction at issue. Thus, on the one hand, there is the possibility of someone's taking the physical realm to be amenable to *conceptual* (or *analytical*) reduction. This would involve claiming that all physical statements can be transformed, by conceptual analysis, into non-physical statements with the same factual content. In other words, it would involve claiming that, for any statement, S1, that is explicitly about the physical world, there is a statement, S2, which is devoid of any physical ontology or physical concepts, such that an analysis of the meaning of S1 reveals that what it claims is expressed more perspicuously by S2. On the other hand, there is the possibility of someone's taking the physical realm to be amenable to *constitutive* (or *metaphysical*) reduction. This would involve claiming that all physical facts are ultimately constituted by non-physical facts. In other words, it would involve claiming that, for any fact, F1, about the physical world, there is either some fact, F2, which is devoid of physical entities and physical properties, or some set of facts, S, all of whose members are devoid of physical entities and physical properties, such that the obtaining of F1 is logically due to, and involves nothing over and above, the obtaining of F2 (the obtaining of the members of S). In claiming that the physical world is something whose existence is

¹ John Foster, *The Immaterial Self*, (London: Routledge, 1991), chs. 2 and 4.

philosophically fundamental, the realist is excluding both types of reduction. He is denying that the physical realm is reducible to something non-physical in either a conceptual or a constitutive way.

The fundamentalist claim excludes both types of physical reduction. But to do justice to the full force of the realist's position, we should not understand it as merely equivalent to the claim that both types of reduction fail. Rather, we should take it to be equivalent to the claim that both types of reduction fail *through and through*. The point is that, whichever type of reduction is involved, the reductive thesis makes a claim about the whole physical realm: it asserts that *all* physical statements are relevantly transformable into non-physical statements, or that *all* physical facts are ultimately constituted by non-physical facts. So, a denial that the physical realm is amenable to either type of reduction is saying something relatively weak: it is saying that at least one physical statement is not thus transformable and that at least one physical fact is not thus constituted. Clearly, the realist wants to say something much stronger: he wants to say that the physical realm is not amenable to either type of reduction *at any point*. So, we must construe his fundamentalist claim as equivalent to the claim that *no* physical statement is analytically transformable into a non-physical statement, and *no* physical fact is constituted by non-physical facts. This claim does not, of course, exclude forms of reduction *within* the physical realm, where one kind of physical subject matter is conceptually or constitutively reducible to another. Nor does it even exclude forms of reduction that stay, with respect to each reduced statement or fact, *partly* within the physical realm. What it excludes is the reduction of something physical to something wholly non-physical—to something where the physical subject matter, whether at the level of concept or at the level of fact, disappears altogether.

Although the fundamentalist claim excludes both types of reduction, and excludes them through and through, it is only the constitutive type that will be of any relevance in our future discussion. This is because it is only with respect to its exclusion of constitutive reduction that, as I see it, the claim is open to serious challenge. It is true that, in the heyday of logical positivism, some philosophers, invoking a verificationist theory of meaning, did seriously suggest that, in so far as they have factual meaning at all, statements about the physical world must be analytically transformable into statements about sense data.² But no remotely plausible method of transformation was ever

² This, for example, was the view of A. J. Ayer in his *Language, Truth, and Logic* (2nd edn., London: Gollancz, 1946).

devised, and it is now, I think, universally acknowledged that the kind of verificationism they espoused was misconceived. Nor, as far as I can see, is there any prospect of developing some other form of conceptual reduction of the physical to the non-physical. So, although I shall be exploring the possibility of a constitutively reductive account of the world, which excludes physical facts from the realm of what is metaphysically fundamental, I shall, from now on, simply take it for granted that the realist is right to insist that the physical language is descriptively irreplaceable—that what we express by means of our physical ontology and our system of physical concepts cannot be re-expressed, preserving factual content, in any other way.

VI

As I have indicated, the form of idealism that concerns us takes the physical world to be something whose existence is constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. I shall speak of this form as *phenomenalistic idealism*. It is immediately clear that phenomenalistic idealism opposes realism in both its claims. Thus, by taking the world's existence to be constituted by non-physical facts, it opposes the fundamentalist claim, and, by taking the constitutive non-physical facts to be, or to centrally involve, facts about human sensory experience, it opposes the independence claim. Moreover, it is clear that it opposes these realist claims through and through. It takes all physical facts to be constitutively reducible to non-physical facts and all aspects of the character of the physical world to be logically dependent on facts about human mentality.

Thus defined, phenomenalistic idealism is a generic position, which can be developed in a number of ways. To begin with, there is the question of what kinds of fact about human sensory experience the idealist takes to be constitutively involved in the creation of the world. In addition, there is the question of whether he takes the relevant experiential facts—whatever they are—to be constitutively responsible for the existence of the world on their own or as part of a richer complex. And, if he takes the latter view, there is the further question of what additional facts he thinks of as featuring in this complex. Now I do not want, at this stage, to try to decide, in all its details, precisely what position it would be best for the idealist to adopt. But if there is to be any serious discussion of phenomenalistic idealism at all, I

do need to settle on an idealist position which is, in certain respects, more specific. The trouble with the position as defined is that it is not easy to see what (if I may put it colloquially) the idealist is *getting at*: the idealist's claim, as formulated, seems hardly more intelligible in real terms than the claim that the existence of the world is constituted by facts about the singing of angels or about the dancing of fairies. Before phenomenalist idealism can become a serious option, I need to present it in a form which gives it some kind of rationale—a form which enables us to understand how the idealist is viewing things—and I can only do this by presenting a position which is more specific. In particular, I need to present a position which is more specific about the kinds of sense-experiential fact constitutively involved in the creation of the world and about the nature of the constitutive role they play.

Human sensory experience tends to be orderly in ways that empirically invite physical interpretation—ways that empirically suggest to us that we are mobile percipients in a three-dimensional world and that our sensory experiences are, typically, perceptions of portions or ingredients of that world in the perspective of our position within it. I shall speak of this orderly character of sensory experience as its *world-suggestive* orderliness. Such orderliness characterizes sensory experience in its own psychological domain and independently of the physical circumstances of its occurrence. It is exemplified, in the first instance, by the relevantly orderly character of typical sequences of experiences, belonging to a single sense realm, and occurring within a single stream of experiential awareness. It is also exemplified by the harmonious way in which, within the experiential life of a single subject, the forms of world-suggestive orderliness of different sequences of experiences fit together to suggest a richer, but unitary, physical account. And, analogously, it is exemplified by the way in which these separate subjective accounts fit together to suggest a still richer, but still unitary, physical story. It is obvious that, whatever other factors may be involved, the world-suggestive orderliness of our sensory experiences plays a major role in the formation and sustainment of our system of physical beliefs. Put at its simplest, the orderliness makes it empirically appear, at the human viewpoint, that we are denizens of a certain kind of three-dimensional world, and, for the most part, our physical beliefs are formed, directly or indirectly, to accord with that appearance. In detail, of course, the situation is more complicated in a number of ways. One obvious complication is that although we can legitimately speak of a shared human viewpoint at which the world-suggestive appearance occurs, and from which

we make a shared doxastic response, everything about this appearance and this response has to be ultimately understood in terms of what obtains and occurs in individual human minds. Another complication is that, although it all ultimately derives from the relevant orderliness of sensory experience, the relevant empirical evidence, and our doxastic response to it, comes at different theoretical levels, where evidence of a distinctively higher-level kind is only available once the account of the world suggested by the lower-level evidence is already in place. Thus, evidence relevant to the testing of theories about the microphysical character of the world is only available in the framework of our common-sense understanding of the world, which preserves the perspective of ordinary observation. For the moment, however, it is only the overall nature of the situation, not its details, that need concern us.

I have represented sensory experience as inviting physical interpretation through its *orderliness*. But there are two points about this that need to be understood.

First, when I speak of the world-suggestive orderliness of sensory experience, I am, as I have made clear, thinking of something that characterizes experiences collectively, rather than individually—an orderliness in the way in which complex groups of experiential items occur and fit together. But what we also need to recognize is that, at least in subjects of sufficient maturity, individual experiences also, typically, though in a quite different way, invite physical interpretation, simply because, by their very content, they phenomenologically pose as presentational perceptions of items in the subject's environment. So an individual visual experience will typically, by its very content, phenomenologically seem to the subject to be the presentational awareness of a three-dimensional arrangement of coloured surfaces and regions in his spatial environment, viewed in the perspective of his own position within it. Likewise, an individual tactual experience will standardly, by its very content, phenomenologically seem to the subject to be the presentational awareness of a certain form of spatial contact between a part of his body and some other object in his environment. And, with variations in detail, the same situation holds in all perceptual sense realms; indeed, it is only when sensory experiences have this phenomenological character that they qualify as perceptual experiences in the ordinary sense. This means that, at least in the case of subjects of relevant maturity, the invitations to physical interpretation that stem from the orderliness of groups of sensory experiences will tend to incorporate the invitations to physical interpretation carried by the content of the experiences themselves: in effect, it will be the orderly way

in which these individual invitations occur and fit together that empirically invites a unitary collective interpretation. There is still the question of what it is about the content of the individual experiences that enables them to pose as presentational perceptions of items in the environment. This is not something that I want to discuss here in any detail. As I indicated in Chapter 1,³ my own view of the matter, which I elaborated in *The Nature of Perception*, is that, when sensory experiences have this phenomenological character, and so qualify as properly perceptual, they combine two elements, one of which is the presentation of a sense datum, or sense quale, that is internal to the mind, and the other of which is an element of conceptual interpretation, which is directed on to the item presented, and which represents it as something environmental.⁴ But, whether or not this is the right account, it is clear, I think, that, in some way, these experiences owe their phenomenological character, wholly or partly, to their conceptual content. In recognizing that the typical experiences of sufficiently mature subjects have the relevant form of conceptual content, we could still suppose, and on the face of it with some plausibility, that in very earliest infancy a subject's sensory experiences are of a more primitive and non-conceptual kind, and do not represent themselves as physically perceptive. And, in the context of such a supposition, we might then conjecture, again with *prima facie* plausibility, that both the beginnings of our physical-belief system, and our facility for having experiences of the ordinary, properly perceptual kind, in some way develop in response to the orderliness of these more primitive experiences. But any hypothesis about the experiential life and psychological development of the neonate subject is bound to be highly speculative, even when such behavioural evidence as can be marshalled is taken into account. All that is certain is that the kinds of sensory experience that we have as adults, and have had as far back into infancy as we can clearly remember, are, typically, of the properly perceptual kind—experiences which, by their very content, make it phenomenologically seem to the subject that he is perceiving an item in his spatial environment from the perspective of his position within it—and, in the context of the present discussion, it is the world-suggestive orderliness of these experiences that is at issue.

The second point concerns the notion of orderliness itself. Normally, when we describe some complex of items as orderly, we mean that, in their arrangement or functioning, these items exhibit some kind of regularity—a

³ Ch. 1, Sect. V.

⁴ *The Nature of Perception*, 147–64.

regularity that can be understood purely in terms of the properties of the items themselves and the domain of reality to which they essentially belong. So the colour pattern on a chess board is orderly in this sense by being uniformly composed of alternating squares of contrasting colours, and the motion of the planets around the sun is orderly in this sense by being uniformly elliptical. It might be thought, then, that when I speak of human sensory experience as relevantly orderly, I have in mind an orderliness of this ordinary kind—regularities of occurrence that can be defined in terms of factors that pertain to the experiential realm itself. But this is not what I have in mind at all. No doubt the realm of experience does exhibit certain aspects of orderliness of this ordinary—regularity-involving—kind, if only in the form of certain statistical tendencies, rather than exceptionless uniformities. But the orderliness I have in mind, as what empirically invites physical interpretation, simply consists in the fact that, by and large, our sensory experiences occur in ways which make them systematically amenable to physical interpretation—to a unitary interpretation across different periods and subjects—and to the extent that experiential regularities contribute to this orderliness, they do so by contributing to this amenability. Admittedly, there is still a sense in which this kind of orderliness is intimately associated with regularity. For although the experiential realm is not especially regular in its own terms, it comes to acquire a much greater degree of regularity in the context of the notional world of its physical interpretation. In other words, this interpretation confers regularity on the realm by representing it as part of a larger, psychophysical reality, whose workings, even with respect to the occurrence of sensory experience, are conspicuously orderly in the ordinary—regularity-involving—sense. It is, indeed, partly for this reason that the *amenability* of our sensory experiences to unitary physical interpretation becomes an *empirical invitation* to such interpretation. As Hume saw, human nature includes a doxastic bias in favour of regularity—a propensity to form beliefs that represent things as working in a uniform and inductively predictable way—and part of the basic reason why our experiences combine to make it empirically appear to us that we are mobile percipients in a three-dimensional world is that the supposition that we are is precisely what would be needed to give their occurrence the regularity and predictability that our minds doxastically seek. I say that this is part of the *basic* reason for the relevant empirical appearance, because it is something that is operative independently of the system of physical beliefs that our experiences lead us to acquire. I say that it is *part* of the basic reason, because,

even at this basic level, it would be wrong to assume that this is the only factor involved. We should not, for example, exclude the possibility that, in addition to the bias in favour of regularity, we have a natural propensity to try to make sense of the course of our sensory experiences in specifically physical (unitary physical) terms. And, more obviously, whatever contribution is made by the natural endowments of our minds, the fact that our sensory experiences tend to be collectively amenable to unitary physical interpretation would not have the empirical significance it does have if the domain of experience were insufficiently rich. It is because the amenability is systematic over such a vast range of experiential items that it is able to create such a clear-cut invitation to us to accept the interpretation and take ourselves to be perceptive inhabitants of the relevant kind of world.

The fact that, over such a rich domain, sensory experience tends to be systematically amenable to a unitary physical interpretation cannot be credibly thought of as merely accidental. We can reasonably assume that there is something that systematically controls the course of experience, making provision for its richness and disposing it to conform to the relevant world-suggestive pattern. I shall speak of this presumed system of control, with its relevant experiential provisions and constraints, as the *sensory organization*. We should not think of this organization as sufficing, on its own, to determine exactly what sensory experiences occur, or even to determine what experience occurs in any particular subject on any particular occasion. But, in the context of the relevant endowments of the human mind, the organization ensures that, whatever its specific details, the course of experience has both the richness and the orderliness needed to make it empirically appear to us that we are perceptive inhabitants of a certain kind of world. It is this that provides the key idea for the idealist position on which I want to focus. In the framework of physical realism, we would think of the sensory organization as stemming from the presence of the physical world, realistically construed, and the ways in which, under the control of certain psychophysical laws, it is disposed to affect human sensory experience: it would be the physical world itself that controlled the course of experience and imprinted, as it were, its own character on it. The idealist position I have in mind sees the lines of dependence as running in the other direction—though the envisaged dependence is of a quite different kind. Thus, rather than taking the physical world to be what is responsible for organizing human sensory experience, it takes the sensory organization to be the central component of that by which the existence and character of the physical world are constituted. Its

central thesis is that, whether on its own or as part of a richer complex, the sensory organization, in the context of the relevant endowments of the human mind, constitutively creates a physical world by disposing things to appear systematically worldwise at the human empirical viewpoint, and it creates all the details of the world's specific character—or, at least, the character of the core physical reality—by disposing things to appear systematically worldwise in the relevantly specific ways. The underlying idea would be that, in the context of the other relevant factors, the sensory organization suffices for the existence of a physical world because it ensures that there is something that will serve as a world *from our standpoint*—something that will serve as world *for us*. As I see it, this position is the only version of phenomenalistic idealism with any prospect of success, and it is on this idealist position that I shall focus in the subsequent discussion.

If the underlying idea of this position is as I have represented it, it might be wondered why I have assigned a constitutive role to the sensory organization at all. Could I not have settled for the theoretically simpler position that what, in the sense-experiential realm, contributes to the constitutive creation of the world, is the specific world-relevant orderliness of the actual course of sensory experience? After all, it is this orderliness, in the context of the relevant endowments of the human mind, that makes things appear relevantly worldwise at the human empirical viewpoint. So would not this orderliness, without the help of the organization that generates it, suffice to provide something that would serve as a world from our standpoint?

There are two reasons why this theoretically simpler position would not serve the idealist's purpose, and why the central creative role has to be assigned to the sensory organization.

The first, and most important, reason is that, although the relevant orderliness of sensory experience, combined with the relevant endowments of the human mind, suffices to make things appear systematically worldwise at the human empirical viewpoint, there are aspects of the total physical reality, and indeed of the core physical reality, that lie outside what is covered by the content of that appearance—aspects for which there is no actual empirical evidence. I am not thinking here of the inscrutability of properties of intrinsic content: for reasons that will shortly emerge, the fact that the factors he envisages as sufficing for the creation of the physical world fail to cover such properties would not trouble the idealist. Rather, I am thinking of the commonplace point that, even within the domain of what is in principle open to empirical investigation, by subjects who are suitably positioned and

equipped, the physical universe abounds with aspects that are not, in practice, empirically manifested, or even empirically suggested by what is. Indeed, given the vastness of the universe, and the inner richness of its perceptible objects, the amount of its empirically accessible character which comes, in practice, to be empirically manifested or empirically suggested must be comparatively small. It is clear, then, that an idealism which took the orderliness of the actual course of sensory experience to be the only sense-experiential factor that constitutively contributed to the existence and character of the world would be hopelessly inadequate. In contrast, an idealist who assigns the central creative role to the sensory organization can hope to make provision for everything about the physical universe that is empirically accessible. For because this organization *disposes* sensory experience to conform to the relevant world-suggestive pattern, and thereby *disposes* things to appear relevantly worldwise at the human empirical viewpoint, it has implications for the outcomes of counterfactual cases, as well as actual, and so, where an aspect of the world that is in principle open to empirical manifestation is not in fact manifested, it is able, by imposing the appropriate dispositions on the sense-experiential system, to secure the potential for its manifestation in the appropriate counterfactual circumstances. Such organizationally sustained potentials are not, of course, limited to cases in which the relevant aspects of the world are unmanifested: for each aspect of the world, the sensory organization creates a range of experiential dispositions, with respect to that aspect, which, within the framework of the total way in which the organization disposes things to appear worldwise at the human empirical viewpoint, secure potentials for its empirical manifestation in the relevant types of condition. The point of stressing these potentials in the unmanifested cases is that this is an area where the theoretically simpler idealist position would conspicuously fail. There is still the question of how the idealist should understand the form of the relevant dispositions and the potentials they sustain. If they are to contribute to the constitutive creation of the physical world, their content must be specifiable in non-physical terms: the idealist cannot afford to settle for such facts as that a subject has the potential to have a certain kind of experience if he finds himself in a certain type of physical situation. I shall return to this point presently.

The second reason why the idealist needs to assign the central creative role to the sensory organization is that it is only by doing so that he can ensure that what gets idealistically created qualifies as a single world for all subjects—a world in which all subjects are experientially located and

to which all have the same kind of empirical access. It is true, as I have stressed, that the world-suggestive orderliness of human sensory experience is harmonious across different minds, thereby making the experiences of different subjects collectively amenable to a unitary physical interpretation. But it is impossible to see how this intersubjective harmony on its own would lead to the emergence of a common intersubjective world, rather than (at best) a collection of separate subjective worlds that qualitatively cohere. In order for phenomenalistic idealism to yield a genuinely common world, the qualitative harmony between the experiences of different subjects has to reflect the presence of some integrated system of control, which ensures the harmony, and does so non-accidentally, and the idealist has to take this system, and its non-accidental ensuring of harmony, to be part of what constitutively creates the world. And this is what he can do by assigning the creative role I have envisaged to the sensory organization. At least, he can do this, provided he conceives of the organization as integrated in the appropriate way, so that it is essential to its method of control, and not just accidental, that the harmony is ensured.

To have any prospect of being able to accommodate the full qualitative richness of the world and the fact that it is the same world for all subjects, the phenomenalistic idealist needs, then, to assign the central creative role to the sensory organization. He needs to say that, either on its own or as part of a richer complex, this organization, together with the relevant endowments of the human mind, constitutively creates the physical world, and determines all the aspects of its core character, and that the way in which it does this is by disposing things to appear systematically worldwise, in the relevantly specific ways, at the human empirical viewpoint—a disposing that ensures the potential for the empirical manifestation of all aspects of the world, not just those that are covered by how things empirically appear in practice. What is still quite unclear—and there might be thought to be a problem here—is how the idealist should understand the *substance* of the organization. The account of the organization that we presently have is a purely functional one: the organization is a system of provisions and constraints that control the course of sensory experience and dispose it to conform to the appropriate world-suggestive pattern. And, indeed, this conception is all that we need in order to understand the envisaged role of the organization in the idealistic creation of the physical world. But if his view of the world is to have any prospect of credence, the idealist has to be able to offer some account of how the relevant system of control is concretely realized, or, at least, of some form, or range of forms, of realization

that he could plausibly take to be possibilities. What we know, for sure, is that he cannot, like the realist, take the organizing agent to be the physical world.

The issue of the concrete realization of the sensory organization is one that I shall consider in more detail later; it is, as we shall see, intimately linked with the even more fundamental issue of whether the idealist can endow the idealistically created world with the requisite objectivity. But one point I want to stress, at this stage, is that, despite initial appearances, the fact that the idealist is debarred from taking the organizing agent to be the physical world is not a genuine restriction on his explanatory options. This is because adopting a phenomenalist idealist account of the physical world does not prevent him from accepting the existence of an external—mind-independent—reality of the same structural and organizational kind as the world of the realist, and taking this reality to be what controls the course of sensory experience. In other words, whatever account of the sensory organization the realist can offer by postulating a certain form of external physical reality, within which human subjects are embodied, and which controls their sensory experiences through the various forms of stimulus input to their sense organs and sensory nerves, the idealist can offer the same account, but with the modification that the postulated reality is not the physical world. It might seem that an idealist who adopted this position would be a physical realist in all but name—that he would be endorsing the realist view of the ultimate nature of reality, both within the realm of human mentality and outside it, but refusing to allow the external component of this reality its conventional title. But this would be a mistake. The idealist's refusal to recognize the external reality as the physical world is not just the withholding of a title. It is a consequence of his insistence that it is a quite different reality that all our physical beliefs are about. It is a consequence of his insistence that what determines the truth-values of these beliefs is not how things stand in the external reality, but how they stand in a reality that is constituted by the relevant facts about human sensory experience, or by some richer complex in which such facts centrally feature. And this makes his position crucially different from that of the realist, however close it comes to the realist's view in other respects.

Not only is it open to the idealist to assign the control of sensory experience to an external reality that replicates the structure and organization of the realist's physical world, but there are three ways in which he could find such an approach attractive.

In the first place, and most obviously, it would provide him with a particularly straightforward explanation of why our sensory experiences

conform to a world-suggestive pattern. It would allow him to say that they conform to such a pattern because it is a relevantly world-like reality that controls their occurrence and constrains them to reflect its own character.

Second—and here I take up an issue raised earlier—by postulating an external controlling reality of the envisaged kind, the idealist would be able to offer a very straightforward account of how the sensory organization is equipped to impose the right kind of dispositions on the experiential system. These dispositions have to cover all aspects of the world's character, whether or not the aspects happen to be empirically manifested, and this means that, for each such aspect, they have to sustain a range of ways in which there is a potential for its manifestation in suitable conditions. Outside the context of idealism, we would think of these potentials in partly physical terms—in terms of the sorts of experience that subjects would have, or would be liable to have, in certain types of physical circumstance. Thus, we would think of such facts as that if I were now to open the drawer of my desk and inspect its contents, I would have, or be liable to have, a certain kind of (contents-revealing) visual experience, and that if I had, five minutes ago, run my finger over the surface of the desk, I would have had, or been liable to have had, a certain kind of (texture-revealing) tactual experience. This ordinary way of understanding the potentials is not available to the idealist—or, at least, not available as his *ultimate* way of understanding them—since he wants them to be constitutively involved in the idealistic creation of the world: he wants them to be aspects of the way in which, by disposing things to appear systematically worldwise at the human empirical viewpoint, the sensory organization determines the full character of the world. But by taking the relevant experiential dispositions to stem from the controlling presence of an external reality that replicates the structure and organization of the world, the idealist could ensure that an exactly analogous understanding was available in terms of the presence of this reality and our relationship to it, and the potentials, thus understood, could then feature in the idealistic creation in the required way. Thus, in the case of the potential that we would ordinarily understand as the fact that, if I were now to open my desk drawer and look inside, I would have, or be liable to have, a certain kind of visual experience, he could ensure the obtaining of an exactly analogous conditional fact concerning the kind of experience that would occur if the external correlate of my physical body were now to operate in the relevantly analogous way in relation to the external correlate of the desk, and he could then take this latter fact to constitutively contribute, in its own small way,

to how things physically stand in the relevant region of the world at the relevant time.

The third point is not directly to do with the provision of an idealistic account of the world. A realist will not think that it is just in its control of human sensory experience that the physical world is involved in the functioning of the human mind. Rather, he will think that, through the functional links between human mind and human brain, it is involved in the functioning of the mind over a very broad area. To take just two of a whole range of cases, he will almost certainly think that our possession of beliefs is causally underpinned by neural structures in the brain that in some way encode their content, and he will think that it is because our attempts to perform physical actions are able to bring about certain motor-neuronal responses that these attempts can be successful. In contrast with the realist, the idealist cannot think of the human brain, or anything else in the physical world, as having any ultimate involvement in the functioning of the human mind, since he excludes the physical from the realm of what is metaphysically fundamental.⁵ Nonetheless, he might still think that, in order to have a plausible explanation of the functioning of the mind, it is reasonable to suppose that, in the areas where the realist recognizes the involvement of the human brain, there are factors external to the mind that support its functioning in an appropriate way. And, if he does think this, postulating an external reality of the same structural and organizational kind as the realist's world, complete with objects that replicate human organisms and human brains, would be the most straightforward option.

In mentioning these ways in which an idealist could find it attractive to suppose that there is an external controlling reality of the envisaged kind, I am not, I should stress, trying to settle the issue at this stage. Indeed, I shall eventually identify a quite different approach that the idealist could take to the question of what underlies the sensory organization, and this approach, too, will merit serious consideration.

VII

I have tried to explain what is involved in the positions of physical realism and phenomenistic idealism, as I understand them, and have identified the

⁵ The discussion in Chapter 5 will reveal that some stress needs to be put on the term 'ultimate' here, though the point is best not pursued in the present context.

more specific form of phenomenalistic idealism on which I want to focus. As I have indicated, the issue between these rival views of the physical world will form the central topic of our subsequent discussion. I want to round off the present discussion by considering how the idealist view I have identified stands in relation to the problems we encountered earlier—the problems that stemmed from the inscrutability thesis established in the framework of realism.

There were two problems, both of which concerned ways in which, in the framework of realism, the inscrutability of intrinsic content seems to undermine the credentials of our belief in a physical world. The first was that a topic-neutral knowledge of the external reality—a knowledge confined to facts about its structure and organization, and about its links with human mentality—does not seem to provide all that is required for an adequate conception of physical space and its material occupants. It does not seem to suffice for a conception that measures up to those concepts of space and occupant that are central to our physical conceptual scheme, and which we have to be able to apply if a belief in a physical world is to be preserved. The second was that once we accept that our empirical evidence does not reveal anything about the intrinsic content of the putative external space—about the intrinsic nature of the spatial medium in which the relevant geometrical structure is supposedly realized—we cannot avoid the conclusion that it does not reveal whether there is an external space at all, rather than merely a quasi-spatial organization. And since a physical world has to be genuinely spatial, this would mean that we cannot, in the framework of realism, avoid the conclusion that the empirical evidence does not reveal whether such a world exists.

We can immediately see that the second of these problems would not arise in the framework of the relevant form of idealism. This idealist position does not identify the physical world with an external reality—a reality that is ontologically independent of the human mind. It takes the world to be something whose existence and character are constitutively sustained by facts of a quite different kind, in which facts about human sensory experience centrally feature. So, from an idealist standpoint, the fact that the empirical evidence does not reveal whether the external reality (if there is one) is spatial or non-spatial has no immediate bearing on the credentials of belief in a physical world. Admittedly, it could still turn out that, for some reason yet to emerge, the idealist needs to envisage a spatial external reality in order to ensure the success of his enterprise. (What kind of external reality, if any,

the idealist needs to envisage is something that I shall consider, in detail, in Chapter 6.) But that does not alter the point that the problem we identified in the framework of realism presupposes the equation of the physical world with the external reality, and does not arise once that equation has been dropped.

What, then, of the first problem? As we saw, this problem has a double aspect. In the first place, a topic-neutral conception of physical space and its material occupants strikes us as deficient in its own terms, just by the very paucity of what it covers about the nature of these items. This ostensible deficiency, as I said, presumably stems from the fact that such a conception of space and its occupants omits all the distinctively phenomenal aspects of how we ordinarily conceive of them. Second, the fact that the topic-neutral conception does not cover any aspects of intrinsic content means that, given certain assumptions about the nature of the mind and the resources of the mental realm, it leaves us free to entertain the hypothesis of the world's being, in substance and character, purely mental. And, even when we grant the relevant assumptions, and acknowledge the possibility of there being an external reality of the relevant structural and organizational kind, we find it hard to accept that such a reality could be identified with the physical world.

Here, again, I think we can see that the idealist position I have identified avoids the problem, and avoids it in both its aspects.

It avoids the first aspect because, unlike realism, it does not require us to discard all the distinctively phenomenal aspects of our ordinary conception of physical space and its occupants. In the framework of realism, as we saw, we have to give up thinking of the world as having a distinctively sensible character conforming to the ways in which it standardly sensibly appears to us. Thus, we are not entitled to think of material objects as having secondary qualities, like colour and flavour, except in the form of powers to affect human sensory experience, or the microphysical properties on which such powers are grounded. And, more crucially, we are not entitled to suppose that, beyond what is covered by their formal geometrical structure, physical figure and extension have any intrinsic resemblance to the ways in which they are visually and tactually represented. In both cases, we are not entitled to ascribe the relevant sensible features to the physical world, because such ascription would not help to explain the facts of sensible appearance, or help to explain anything else about how things empirically appear at the human viewpoint. The situation for the idealist is quite different. On his view, it is the sensory organization, in the context of certain other factors, that logically

determines the character of the physical world, and it does this by disposing things to appear systematically worldwise at the human empirical viewpoint. But the way in which things are disposed to *sensibly* appear at the human viewpoint is part of the whole complex of ways in which they are disposed to *empirically* appear, and so the idealist can think of it as making its own distinctive contribution to the full character of the world. He can endorse our common-sense view that the world is, in character, as its standard sensible appearance represents it, because he can take its character to derive, in part, from this appearance. It remains true, of course, that crediting the world with a distinctively sensible character does not help to explain the facts of sensible appearance or any other aspect of the empirical evidence. But in the idealist system this point is irrelevant, since it is not the world itself, but something else (whether world-like or not) that is responsible for organizing sensory experience and determining what empirical evidence obtains.

It is also clear that the idealist position avoids the second aspect of the problem. The idealist's physical world cannot have any aspects of character beyond what are assigned to it by the factors that constitutively create it, and these factors confine its character, or at least the character of its primary core, to what is covered by how things are disposed to appear at the human empirical viewpoint. So, there is no place in the idealist's world for anything empirically inscrutable, and, in particular, no place for any inscrutable intrinsic content that could turn out to be mental in nature. Even if he holds the appropriate views about the nature of the mind and the resources of the mental realm, the idealist will not have to take seriously the possibility that the familiar space in which we move and the familiar objects that we see and handle may be items in the psychological life of some alien mind.

In the framework of realism, denying the existence of inscrutable physical content is not an option. We cannot avoid thinking that the physically fundamental particles have intrinsic natures beyond the spatial and dispositional properties that empirical investigation is in principle capable of revealing, nor avoid thinking that physical space itself, as a concrete medium for space-occupying objects, has an undetectable inner nature in which its geometrical structure is realized. We cannot avoid these conclusions because the realist's world is something that we have to be able to make metaphysical sense of in its own terms, and, considered in its own terms, a world that lacked the inscrutable forms of particle and spatial content would be seen as metaphysically incomplete: its qualitative ingredients would not be enough to enable us to understand how there could be a genuine space and space

occupants at all. The situation in the framework of the relevant form of idealism is quite different. In this framework, we can accept that our ultimate way of making metaphysical sense of the world is not by focusing on it in its own terms, but by considering it in the perspective of its idealistic creation, and, in that perspective, where we explicitly recognize the world as something whose existence is nothing over and above the factors that idealistically create it, there is no difficulty in understanding how it can turn out to lack forms of qualitative ingredient that it would need to have if its existence were metaphysically fundamental. There is no difficulty in understanding how, as the constitutive creation of factors of a different kind, it is limited to the empirically accessible nature that these factors give it.

The fact that the idealist position avoids the problems that beset physical realism does not mean that it does not face problems of its own. And, indeed, at this stage it is still likely to seem to us that the whole idea of the physical world's being constitutively created in the way envisaged is an absurdity and does not come even close to satisfying what our basic understanding of such a world requires. In my view this is not so, and I shall be devoting the remainder of the discussion to trying to show this, and to trying to show that it is only by embracing this idealist view of the world that we can do justice to our deepest intuitions about the ultimate determinants of physical truth and falsity.

4

The Refutation of Realism

I

Physical realism makes two claims. The first is that the physical world is ontologically independent of the human mind—something whose existence is logically independent of facts about human mentality. I have labelled this its *independence* claim. The second is that the physical world is something whose existence is philosophically fundamental. I have labelled this its *fundamentalist* claim. I have tried to set out, in some detail, how these claims are to be understood, and what they do and do not commit the realist to accepting. In particular, I have made it clear that the independence claim is to be understood as applying to the world with respect to every portion of its spatio-temporal spread. And I have made it clear that the fundamentalist claim excludes any kind of reduction of something physical to something non-physical—whether the conceptual (analytical) reduction of certain physical statements to non-physical statements or the constitutive (metaphysical) reduction of certain physical facts to non-physical facts.

Standing in radical opposition to physical realism, and in opposition to both its claims, is the position I have labelled *phenomenalistic idealism*. This asserts that the physical world is something whose existence is constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. The version of phenomenalistic idealism which concerns us—the only version, as I see it, which has any prospect of acceptability—is one that assigns the central constitutive role to what I have termed the *sensory organization*. This organization is the unified system of provisions and constraints that controls the course of human sensory experience and disposes it to conform to its world-suggestive pattern. The central thesis of the relevant form of idealism is that, whether on its own or as part of a richer complex, and in the context of

certain endowments of the human mind—endowments that render the mind empirically receptive to the orderly character of its sensory experiences—the sensory organization secures the constitutive creation of the physical world by disposing things to appear systematically worldwise at the human empirical viewpoint, and it logically determines the detailed character of the world by disposing things to appear systematically worldwise in the relevantly specific ways. I shall refer to this version of phenomenalistic idealism as *canonical* idealism (the label indicating that it is the version which I think that the idealist needs to adopt).

We have already seen how physical realism encounters three problems, and ones that the idealist position I have identified avoids. In the first place, by taking the world to be an external reality, which is ontologically independent of the human mind, realism does not allow the development of a satisfactory account of physical-item perception, since there is no way of understanding how perceptual awareness can reach to objects in such a reality. Canonical idealism avoids this problem, since, with the central role assigned to the sensory organization in the constitutive creation of the world, our awareness does not need to reach beyond the boundaries of the mind to make ontological contact with things in the world. Second, under realism our knowledge of the nature of the physical world would be confined to facts about its structure and organization, leaving its properties of intrinsic content empirically inscrutable, and such knowledge—topic-neutral knowledge—does not seem to provide an adequate conception of physical space and its material occupants: it does not seem to provide a conception that measures up to those concepts of space and occupant that are central to our physical conceptual scheme and that we have to be able to apply to retain our belief in a physical world at all. Canonical idealism avoids this problem by both confining the character of the world to what is capable of revealing itself empirically and by allowing us to retain our ordinary view that we can gain knowledge of the intrinsic content of the world, and, in particular, of what, beyond its geometrical structure, physical space is like in itself, from the ways in which things sensibly appear to us. Third, once we have accepted that the empirical evidence does not reveal anything, beyond facts of geometrical structure, about the intrinsic nature of the putative external space, it seems we have no way of telling whether there is an external space at all, rather than merely a quasi-spatial organization. And, of course, unless the external reality is spatial, it cannot, as realism requires, form the physical world. This problem does not arise under canonical idealism, since even if the idealist

accepts the existence of an external reality, he does not identify it with the physical world.

These problems for realism are serious, and I do not think that the realist has any satisfactory way of dealing with them. Even so, I would be reluctant to rely on them as the sole basis for rejecting the realist position. This is because, although the problems bring out respects in which the realist view of the world has awkward consequences, they do not reveal a fundamental flaw in the view itself, and without the identification of such a flaw it is still likely to seem to us that a realist understanding of the world is the only one that can be seriously entertained. If realism is to be finally discredited, it will have to be by means of an argument that attacks it more directly. Such an argument is what I shall try to develop in the present chapter. Specifically, I shall try to show that realism provides a mistaken account of the ultimate truth-conditions of propositions about the physical world—a wrong account of the factors that ultimately determine what physical facts obtain. If successful, this argument will suffice to undermine physical realism entirely. It will also undermine it in a way that will serve to establish the truth of canonical idealism, or, at least, to do so on the assumption of the existence of a physical world.

In the last chapter I introduced the notion of the core physical reality (or the physical world in its primary core). I defined this core as that portion of the total physical reality which is covered by facts that are either physically fundamental or wholly constituted by facts that are physically fundamental, where a fact is physically fundamental just in case it is purely physical and is not wholly or partly constituted by other physical facts. The point of introducing this notion was to mark a certain limit on the implications of the realist's independence claim with respect to the *character* of the physical world. In taking the world to be ontologically independent of the human mind, the realist is free to accept that there are physical properties and types of physical object whose instantiation logically depends on facts about human mentality. So, he is free to accept that the colour of a material object consists in, or logically involves, its disposition to look coloured in a certain way to the normal human percipient who views it in standard conditions, and that something's being a clock logically involves its having the potential to indicate the time to a human subject who understands its chronometric significance. All he is required to say is that facts about human mentality make no logical contribution to the character of the *core* physical reality. This means that, in trying to evaluate the realist position, it is only the situation of the core physical reality that need concern us.

Although realism is committed to taking the character of the core physical reality to be, in all respects, logically independent of facts about human mentality, there is, as I mentioned, room for differences of view, within the realist camp, about the kinds of ingredient it contains. The argument against realism that I am about to develop can be adapted to any account of these ingredients that the realist cares to adopt; at least, it can be adapted to any account with any shred of plausibility. But, in order to have a concrete point of focus, and one which is reasonably simple, I shall assume that, at the level of what is physically fundamental, the realist takes the physical reality to be composed of:

- (1) time;
- (2) a three-dimensional and uniformly Euclidean space;
- (3) a stock of mobile material particles, of various intrinsic types;
- (4) for each particle, the fact of its being of a certain intrinsic type;
- (5) the precise arrangement of these particles in space and time;
- (6) the obtaining of certain forms of causal relation between events that pertain to this arrangement;
- (7) certain laws of nature that govern particle arrangement and the relevant forms of causal relation.

The realist's core physical reality will then consist of this physically fundamental reality, thus demarcated, together with whatever is additionally covered by physical facts that are wholly constituted by the physically fundamental facts (in other words, by the facts that feature in (4)–(7)).

For convenience, I shall assume an old-fashioned Newtonian (absolutist) view of physical space, as something whose positions have identities that are logically independent both of the things located in it and of time (though, of course, it will only be by reference to space-located items and times that these positions can be identified). I say that this is *for convenience*, because the assumption is not something that my argument requires, as a premiss or a presupposition, but merely something that will help to make its formulation, and the discussion of the issues it raises, simpler.

II

Let us speak of a concrete reality as a *relevant external reality* if and only if it satisfies three conditions: first, its character is, through and through,

logically independent of facts about human mentality; second, the facts that feature in it are in no instance constitutively reducible to facts that lie outside it; and third, whether on its own, or in conjunction with laws that link it with human mentality, it is responsible for the systematic control of human sensory experience. In developing my argument against physical realism I shall provisionally work on the assumption that there is a relevant external reality in this sense. In making this assumption I am not begging any questions against the realist, since he is committed to taking the core physical reality to be itself a relevant external reality. Thus, his realist position, as I have specified it, explicitly commits him to taking the core physical reality to satisfy the first and second of the relevant conditions; and anyone who takes this reality to satisfy the first two conditions cannot avoid thinking of it as satisfying the third as well, since there is no other remotely plausible view. The idealist, too, is likely to accept the existence of a relevant external reality, as a way of accounting for the sensory organization. If he does recognize such a reality, he will not, of course, take it to be something physical, since he takes the physical reality to be, through and through, both constitutively reducible to non-physical facts and logically dependent on facts about human sensory experience. Even if he postulates an external reality which is world-like in its character—an option which, as we have seen, his position permits—he will think of this reality not as what *forms* the physical world, but as something that *underlies* it.

However we should think of the relationship between the relevant external reality and the physical world, we can envisage the possibility that the structure and organization of this reality are at variance with what our empirical evidence suggests and what, in response to this evidence, we ascribe to the world in our ordinary and scientific thinking. Now the anti-realist argument that I want to present is concerned with the question of how we ought to interpret things physically if this possibility obtains. More precisely, it is concerned with this question for a specific range of possible cases, where certain further conditions are satisfied. In order to identify these conditions, it will be best to start by focusing on a particular example of the sort of thing I have in mind. It is an example that I have employed, in various forms, elsewhere, though my discussion of it here will be rather more elaborate.¹

¹ Thus, see my 'The succinct case for idealism', in H. Robinson (ed.) *Objections to Physicalism* (Oxford: Oxford University Press, 1993), 293–313, 'In defence of phenomenalist idealism', *Philosophy and Phenomenological Research*, 54 (1994), 509–29, and *The Nature of Perception*, pt. 5.

To prepare the way for this example I shall assume that the relevant external reality is world-like in its character and that its fundamental ingredients are of the same kind as we are taking the realist to assign to the core physical reality. In other words, I shall assume that, at its fundamental level, the external reality is, like its envisaged physical counterpart, composed of:

- (1) time;
- (2) a three-dimensional space;
- (3) a stock of particle-like mobile space occupants, of various intrinsic types;
- (4) for each such occupant, the fact of its being of a certain intrinsic type;
- (5) the precise arrangement of these occupants in space and time;
- (6) the obtaining of certain forms of causal relation between events that pertain to this arrangement;
- (7) certain laws of nature that govern occupant arrangement and the relevant forms of causal relation.

The total external reality will then, like the envisaged core physical reality, consist of what is included in these seven items, together with whatever is additionally covered by facts that are wholly constituted by the facts that feature in items (4)–(7). I shall refer to this external reality as *E*, to its space as *S*, to the occupants of *S* as *E-particles*, and to the relevant laws as *E-laws*. The point of modelling the fundamental ingredients of the external reality on those that the realist assigns to the core physical reality is that, on the realist view, the external reality has itself to form the core physical reality; and since I am intending to use the example as a basis of an argument against the realist, I do not want to exclude his position by the very terms in which the example is constructed. For the same reason, I shall, in line with what I assumed in the case of physical space, take the identities of the positions in *S* to be logically independent both of the things located in it and of time.

In addition to the composition of *E*, there is the matter of its relationship to us, and, specifically, its role as the systematic controller of our sensory experiences. To provide for this, I shall assume that, in addition to the *E-laws*, there is a system of what I shall speak of as *link laws*, which make provision for, and regulate, forms of causal interaction between *E* and the realm of human mentality, and, in particular, for the direct causal influence of *E* on human sensory experience. The details of this need not concern us, but a plausible general account, and one that keeps things in line with what would be needed by the realist, would be to suppose that there is a certain class

of complex persisting objects in E—the E-equivalents of human biological organisms—and a one-to-one correlation between these objects and human subjects, such that, for each correlated pair, there is a set of laws prescribing ways in which the condition of the E-object is empowered directly to affect the mentality of the subject, and ways in which the mentality of the subject is empowered directly to affect the condition of the object. (For a realist, of course, these E-objects, which are the E-equivalents of human biological organisms, would actually be such organisms.) Unlike the E-laws, which govern events and causal processes within E, the link laws are not part of the external reality—part of what is covered by E—since their obtaining is not logically independent of the realm of human mentality. Nonetheless, it will often be convenient to think of the E-laws and the link laws as forming a single nomological package, and I shall speak of the members of this package as the *E-relevant* laws.

With these preliminary assumptions in place—ones that have been deliberately chosen to leave open a realist construal of the situation—let us now turn to the example itself.

Let us suppose that, with one crucial exception, the E-relevant laws impose the same constraints on events across the whole of S and time. The exception is this. Within S there are two wholly separate spherical regions, R1 and R2, of the same size, and a one-to-one correlation C between R1 points and R2 points, such that the distances between the points in each region are the same as those between their C-correlates in the other, and everything is nomologically organized, both with respect to what takes place within E and with respect to the causal traffic between E and human mentality, exactly as if—by the standards of what would be required for organizational uniformity—R1 and R2 were C-wise interchanged (with each point transferred to the position, relative to the space that lies outside the regions, of its C-correlate). To illustrate, suppose there is a certain kind, K, of process in S, such that, away from the boundaries of R1 and R2, any instance of K involves the spatially continuous motion of a certain type of E-particle. Then, whenever, in the context of a K process, a particle of this type comes, in the normal (spatially continuous) way, to some point on the boundary of either R1 or R2, it instantaneously changes its location to its C-correlate on the boundary of the other region, and continues in the appropriate spatially continuous motion from there. Likewise, suppose there is a certain type, T, of momentary particle arrangement such that, away from the boundaries of R1 and R2, the instantiation of T is nomologically sufficient, under the link

laws, for the causing of a certain kind of human mental effect. Then, where particle arrangements touch or extend over the boundaries of R1 or R2, what will determine whether things are nomologically right for the production of that kind of effect will not be whether there is in fact an instantiation of T, but whether there would be such an instantiation if the total particle arrangements within R1 and R2 were C-suitably transposed. Quite generally, by the standards of how, in the rest of the space, things behave, interact, and interact with our minds, everything is organized, with respect to the boundaries of the two regions, as if each region had (C-specifically) the other's location. One way we can represent the situation is by drawing a distinction between the *intrinsic* geometry of S and its *functional* geometry. Thus, we can say that each of the two regions is functionally located (located with respect to S's functional geometry) where the other region is intrinsically located (located with respect to S's intrinsic geometry). It is the intrinsic geometry of S, of course, that forms its real geometry—its geometry in the ordinary sense. Its functional geometry is merely that non-actual intrinsic geometry that S would need to have, with its organization held constant, for the achievement of organizational uniformity—the achievement of a situation in which the same constraints on behaviour and causation were in force across all its regions.


Everything in E is organized as if, by the standards of uniformity, R1 and R2 were interchanged. And, as I have made clear, this relates not just to the constraints on behaviour and causation within S, but also to the modes of interaction between what occurs in S and human mentality. In particular, then, it covers the ways in which situations in S affect human sensory experience. Crucially, this means that the organizational anomaly in the external reality is wholly concealed at the level of empirical appearance. At the human empirical viewpoint, everything seems to indicate that, within our own physical world, things behave in a completely uniform way across the whole of space. This, in turn, of course, shapes the character of our ordinary physical beliefs, which are directly responsive to what the empirical evidence suggests. So, the physical world of our empirical beliefs—both those at the level of our common-sense thinking and those which emerge through science—comes to mirror not the external reality as it is, but the reality that would obtain if the relevant regions were interchanged. In the world we empirically construct, the region corresponding to R1 gets located in the surroundings corresponding to the S-surroundings of R2, and the region corresponding to R2 gets located in the surroundings corresponding to the S-surroundings of R1. In other words, with respect to the relevant

regions, the topology which features in the content of our physical beliefs coincides with the functional topology of *S* rather than with its intrinsic topology. For simplicity, let us assume this to be the only point of conflict between the structure and organization of *E* and those of the world we empirically construct—the world which the empirical evidence suggests and which features in the physical beliefs we form on that basis.


To bring the situation into sharper focus, let us make the example more concrete. Let us suppose that *R1* and *R2* respectively correspond to what, in the empirically constructed world, qualify as the regions of Oxford and Cambridge, or, more precisely, as the smallest, or as among the smallest, equisized spherical regions that respectively contain Oxford and Cambridge. (I say ‘or as among the smallest’ because, given that the two cities differ in shape and size, there may well be—indeed, I imagine there will be—more than one pair of equisized spherical regions which contain them and which are as small as any other such pair.) Strictly speaking, of course, Oxford and Cambridge do not define fixed spatial regions within the physical universe, since, as parts of the rotating and orbiting planet, they are in constant motion. And this means that we can only correctly represent them as corresponding to *R1* and *R2* if we reconstrue the latter as spatio-temporally continuous sequences of regions-at-times. Such a reconstrual would be perfectly feasible, and would not affect the course of the argument. But, rather than getting involved in such a complication, I shall simply pretend that our scientific beliefs represent the earth as stationary, with each of the two cities occupying the same region of physical space at all times.

We are supposing that *R1* and *R2* in *S* respectively correspond to what, in the empirically constructed world, qualify as the regions of Oxford and Cambridge. With this supposition in place, the situation specified in abstract terms above becomes, more concretely, this: all the empirical evidence suggests to us that the Oxford region is in what qualifies as Oxfordshire, and the Cambridge region is in what qualifies as Cambridgeshire, but, because of the organizational anomaly in *E*, the situation with respect to *R1* and *R2* is reversed, with the *S*-surroundings of *R1* corresponding to the relevant portion of Cambridgeshire (that is, to what remains of Cambridgeshire when the Cambridge region is subtracted), and the *S*-surroundings of *R2* corresponding to the relevant portion of Oxfordshire (that is, to what remains of Oxfordshire when the Oxford region is subtracted). To bring out what this involves, consider the case of someone who drives from Oxford to Cambridge, and note the striking difference in character between the route of

his journey in S and the route of the journey which is empirically apparent. The external process begins in R1, the region which corresponds to Oxford, but which is surrounded by what corresponds to Cambridgeshire. But as soon as the driver moves out of the Oxford region, the process instantaneously changes its S-location to an area just outside R2, corresponding to an area in Oxfordshire normally thought to be just outside Oxford. The process then continues in a way that coincides with the empirically apparent route until it reaches the boundary of R1, when, as the empirical journey finds the driver entering Cambridge, it once again instantaneously changes its location and becomes a process moving through R2. So, what seems to the subject to be a spatially continuous journey in physical space corresponds to a process in S which becomes dramatically discontinuous whenever it reaches (whether from the inside or the outside) the boundaries of R1 or R2. (See the diagram below.) And, of course, it is not just that the discontinuities



Journey in the empirically constructed world



Process in S

in the external process are not apparent to the *subject*. Rather, the E-relevant laws, which control both what takes place in E and the effects of E on human experience, ensure that they would remain completely undetected by anyone, at any time, by any empirical means. All the empirical evidence, past and future, actual and potential, from the most casual observation to the most searching experimental test, favours the conclusion that the external process is (in line with how things experientially appear to the subject) spatially continuous.

One final point. I have represented the situation in the external reality as one which, whenever an object reaches the boundary of one of the relevant regions, it is required to undergo an instantaneous shift of position to the corresponding place on the boundary of the other, and to continue its journey from there. But it might be thought that the very notion of such instantaneous relocation across a spatial distance is unintelligible—that we cannot make sense of an occupant of space persisting through time except in a spatially continuous form. (Staying in the same place, of course, counts as a limiting case of spatial continuity.) If this is so, the example on which I am wanting us to focus is not in fact available.

I have two things to say in response to this. The first is that I am not inclined to accept the principle that occupants of a space can only persist in a spatially continuous form. What is certainly true, and guaranteed by the logic of the situation, is the weaker principle, that, provided it remains an occupant of the relevant space throughout its history, any period through which a space occupant persists must either be a period of spatially continuous existence or exhaustively divide into periods of spatially continuous existence. This principle is guaranteed by the logic of the situation, because it is only by being spatially continuous that a change in spatial position could use up time: a shift from one place to a spatially separated place that did not pass through a spatially continuous series of places, or through a sequence of such series, would have to be instantaneous—it would not be able to pass through a continuous series of times. But this weaker principle explicitly allows for cases in which successive phases in the history of an occupant are not spatially continuous—cases in which, on some occasion, an occupant follows a spatially continuous path up to a certain time, and then continues on a spatially continuous path from that same time, but from a different (spatially separated) place. And I am inclined to think that this is a possibility that we should recognize. Moreover, it seems to me that we can envisage circumstances where it would be natural to conclude that such a case had

actually occurred. Envisage, for example, a situation within the physical world in which, for a certain time t :

- (1) x is the spatially continuous phase in the history of a certain space occupant up to t , and y is the spatially continuous phase of the history of a certain space occupant from t .
- (2) The terminal spatial position of x is spatially separated from the initial spatial position of y .
- (3) Apart from this spatial separation, the relationship between x and y has all the features that are characteristic of the relationship between successive phases in the history of a single occupant. In particular, it exemplifies all the relevant forms of qualitative and causal continuity.
- (4) There is nothing other than y that has any claim to count as the successor to x in the history of the x occupant, and nothing other than x that has any claim to count as the predecessor to y in the history of the y occupant.

If such a situation were to occur, I think it would be natural to conclude that there is only one persisting space occupant involved and that x and y are successive phases in its history.

My second point is that, even if I am wrong in thinking that we can make sense of a situation in which successive phases in the history of a single space occupant are spatially discontinuous, there is a simple remedy with respect to my example of R1 and R2. For I could have avoided the problem altogether if I had assigned a different form of ontological composition to the external reality. As things stand, this reality is ontologically composed, at its fundamental level, of a three-dimensional space and a stock of particle-like space occupants. But the inclusion of these particle-like occupants, which persist through time and move through space, was not essential to the purposes of the example. I could, instead, have taken the fundamental external reality to be composed of a space and certain region-characterizing properties whose spatio-temporal distribution and functional organization simulated the situation of there being particle-like occupants. Had I done so, the issue over the possibility of spatial discontinuities in the history of a persisting occupant would not have arisen. We have already seen how, in the framework of realism, this sort of ontological composition can, without doing violence to the empirical evidence, be envisaged in the case of the fundamental physical reality—with all persisting physical space occupants being represented as ontologically derivative entities, whose existence is

constituted by the distribution and organization of the relevant properties.² And if it is envisageable under realism in the case of the physical reality, it must be envisageable in the case of the relevant external reality, whether we think of this reality as physical or not.

Although I could avoid the issue over discontinuity by excluding persisting occupants from the fundamental external reality, I shall, for simplicity, assume that I am right in supposing that we can make sense of the relevant kinds of spatial discontinuity, and continue with the example in the form in which I have presented it.

III

I think that the nature of the case we are envisaging is by now clear. What still needs to be decided is what we should make of it. And one thing which must be conceded, at the outset, is that, taken as a suggestion as to how things may actually be, what is envisaged is not at all plausible. It is not just that, *ex hypothesi*, it runs counter to all the empirical evidence, interpreted in the most straightforward way. It is also that there would be something inherently puzzling about a reality that was organized in the way envisaged. If there is nothing which qualitatively distinguishes R1 and R2 from other regions in S—and there has been no suggestion of anything of that sort—it is surely very strange that the laws should treat them in such a distinctive fashion, and even stranger that they should do so in a way which contrives to prevent any empirical trace of what is happening. None of this, however, will affect the role of the hypothetical case in our present discussion. All that matters to that discussion is that what is being envisaged—however epistemically improbable—is coherently conceivable. And this is surely so. However strange such a situation would be—however strong our reasons for supposing that it does not obtain—there is surely nothing which logically excludes the suggestion that the external reality is organized in this anomalous way.

Given that what is being envisaged is conceivable, the crucial question is how we should interpret it *physically*. How does the peculiar organization with respect to R1 and R2 affect the situation of Oxford and Cambridge in the physical world? It is essential, of course, that we should address this

² Ch. 3, Sect. II.

question from an appropriately detached standpoint. As I have stressed, the physical beliefs which we form in the context of our ordinary empirical lives are automatically shaped to fit the empirical evidence. But the question which we now need to consider requires us to stand back from our empirically formed beliefs, and try to reach a verdict about their correctness when we take into account the external factors that are envisaged as lying behind them.

So, supposing ourselves to know that the relevant external reality is structured and organized in the relevant way, and that it is linked with the world of our empirical beliefs in the ways indicated, what conclusion should we reach about the physical situation? One drastic response would be to say that, with that as the external situation, there is no physical world at all. This response can be instantly dismissed. Maybe there are independent reasons for being sceptical about the existence of the physical world, but it is obvious that, with no other points of difference between the external reality and the empirically constructed world, this particular organizational anomaly leaves our belief in the existence of the physical world unscathed. Another response that can be instantly dismissed would be to say that while there is a physical world, it contains gaps—holes in its spatial fabric—at the points where we ordinarily locate Oxford and Cambridge. Given that it is only at the boundaries of R1 and R2 that anything odd happens in S (for behaviour *within* the regions is constrained in the same ways as behaviour *outside* them), and given that there is nothing to distinguish our own links with what takes place within these regions from our links with what takes place outside them, this proposal, too, is a clear non-starter. There are, in fact, only two possible responses that deserve serious consideration. The first would be to say that there is a physical world, and one which is replete with the Oxford and Cambridge regions, and that its structure and organization coincide with those of the external reality. This would involve saying that, contrary to what we ordinarily believe, and to what all the actual and potential empirical evidence suggests, Oxford is really in Cambridgeshire, and Cambridge is really in Oxfordshire. The second option would be to say that there is a physical world, and again one replete with the Oxford and Cambridge regions, but that its structure and organization are as they are empirically represented at the human viewpoint. This would involve saying that, despite the positions of R1 and R2 in the external reality, Oxford and Cambridge are, physically, where we ordinarily take them to be—Oxford in Oxfordshire, Cambridge in Cambridgeshire. In short, we have the topology of the external reality itself and we have the topology which, by its distinctive organization, this reality

projects on to the human empirical viewpoint, and what we have to decide is whether the topology of the physical world matches the one or the other.

It seems to me that, given our ordinary understanding of what is involved in physical spatial arrangement, the second of these options is the one which we need to embrace: we need to think of the physical topology as conforming to the empirically projected topology, rather than to the empirically hidden topology of the external reality. The basic point, as I see it, is that our understanding of the nature of physical spatial arrangement cannot be separated from our ordinary ways of empirically assessing what forms of arrangement obtain. This does not mean that these ordinary methods of assessment are infallible—that they never deliver a mistaken verdict. It is undeniable that there are all sorts of ways in which, in particular cases, the evidence we possess with respect to some issue of arrangement may be misleading, and its evaluation by our standard methods support a false conclusion—cases of illusion, faulty measuring instruments, and erroneous reports are obvious examples. But what it does mean, I think, is that our ordinary understanding of physical spatial arrangement commits us to recognizing the overall objective soundness of these methods, so that we are only entitled to hold back from an unqualified acceptance of the verdicts they deliver in so far as either (i) these verdicts are themselves delivered with qualifications, or (ii) we possess independent evidence that counts significantly against them, or (iii) we have to take seriously the possibility of there being the potential for such contrary evidence. And, crucially, it means that, at the theoretical limit, where some issue of arrangement is thoroughly open to empirical investigation, and where we are assuming ourselves to know that all sufficiently searching empirical tests—actual and potential—would, by those methods of assessment, unequivocally deliver a certain verdict, we are obliged to recognize that verdict as logically decisive. We are obliged to recognize that verdict as decisive, because if we were to leave room for the possibility of its being mistaken—allow that how things are, in the relevant respect, may be through and through belied by how things are disposed to present themselves empirically—we would lose our grip on what physical spatial arrangement is. The application of this to the particular case under consideration is immediate. What is here at issue are the positions of Oxford and Cambridge in the geographical layout of England—something which is thoroughly open to empirical exploration—and we are supposing ourselves to know that everything is organized in a way that ensures that Oxford is disposed to pass all the relevant empirical tests for being in

Oxfordshire, and Cambridge to pass all the relevant empirical tests for being in Cambridgeshire. So, our ordinary understanding of what is involved in physical spatial arrangement—an understanding that implicitly recognizes the overall soundness of our ordinary ways of empirically gauging such arrangement—obliges us to accept those empirically assigned locations as the correct ones. It obliges us to accept those locations as correct, even though we know that they are out of line with how the corresponding regions are located in the postulated external reality.

In claiming that our ordinary understanding of the nature of physical spatial arrangement is tied, in this way, to our ordinary methods of empirically assessing what forms of arrangement obtain, I am not, I should stress, endorsing a general verificationist account of meaning and understanding. I am not, for example, tempted to suppose, like the followers of the later Wittgenstein, that our understanding of what it is for a subject to be in a certain mental state is tied to our standard ways of gauging the mental condition of a subject in third-person perspective, nor tempted to suppose that our understanding of what it is for something to have happened in the past is tied to our standard ways of gauging the character of the past from present evidence. But there is a special reason why our understanding of physical spatial arrangement has this verificationist aspect. The physical world, to qualify as *the* physical world (as the world that forms the target of our ordinary physical beliefs) has to be *our* world, and it can only be our world in the relevant sense, if it is ours *empirically*—if it is a world that is, as we might put it, *empirically immanent*. It is arguable that being ours empirically, in the sense required, involves being ours perceptually, and we have already noted how a denial of perceptual access to the physical world would put the conception of the world as our world under pressure.³ But whether the physical world has to be ours perceptually, there can be no denying that, to be our world in the relevant sense, it has to be a world to which we have empirical access through our senses, and of whose structure and organization we are systematically equipped, through that access, to gain knowledge. It is this that explains the verificationist aspect of our understanding of physical spatial arrangement. For if the world has to be empirically immanent in that sort of way, it is hardly surprising that our understanding of the nature of physical spatial arrangement should turn out to assign the envisaged authority to our ordinary methods of empirically gauging such arrangement, and that,

³ See Ch. 1, Sect. VI.

at the theoretical limit, when nothing impedes a thorough investigation of some issue of arrangement, and when all sufficiently searching applications of these methods would unequivocally support a certain conclusion, this understanding should leave no logical room for an alternative conclusion. Indeed, our expectation must be that, quite generally, our understanding of what is involved in something's being physically the case will turn out to be linked in this way with the methods of empirically gauging what is physically the case. But for the time being I shall continue to confine my attention to the specific case of physical spatial arrangement, which forms the primary structural ingredient of the world, and continue to focus, in particular, on my claim that, in a case like that of R1 and R2, our understanding of what is involved in such an arrangement obliges us to accept that the relevant regions are arranged in accordance with what is projected on to the human empirical viewpoint, rather than in a way that matches the concealed arrangement of their counterparts in the external reality.

I can envisage three possible objections to the position I am taking, and I shall consider them in turn.

In the first place, it might be thought that, in taking it to be uncontroversial that the physical world *has to be* our world empirically, and so *has to be* a world to which we have empirical access, I am already assuming the falsity of physical realism, and so assuming the very thing that I am seeking to establish. For realism takes the world to be ontologically independent of the human mind, and, if the world has that kind of independence, it must presumably be capable of existing without standing in any relationship to human mentality, and so without being something to which we have empirical access. There are, admittedly, as we saw, respects in which the realist's independence claim needs to be qualified if it is to allow him the full range of options to which he is entitled—qualifications concerning the possible contribution of mental causation to the identities of physical particulars and the possibility of the realist's adopting a physicalistic account of psychological properties. But, even with these qualifications, it presumably remains true that a realist will think of the world, or at least the world in its primary core, as something which is logically capable of existing without being related to human mentality in the specific ways required for empirical access. So, it might still seem that, in assuming that the physical world has to be empirically accessible to us, I am begging the question against him.

Any such thought, however, involves a confusion between two quite different ways in which there could be an issue over whether empirical

accessibility is logically required for the existence of the physical world—a confusion, in effect, between a conceptual and a metaphysical issue. On the one hand, there is the question of whether, as an aspect of how things actually are, something could qualify as the physical world without our having empirical access to it. And, given the way in which we conceptually identify the physical world—as, in part, our world empirically—the answer to this question is uncontroversially negative and something that even a realist will be happy to accept. Even a realist will acknowledge that if, contrary to what we ordinarily assume, there is no world to which we have empirical access, then there is no physical world at all. On the other hand, given that there is something that qualifies as the physical world, there is the question of whether that thing, or that thing in its primary core, is logically capable of existing without being empirically accessible to us. And to this question the realist will return an affirmative answer, as a consequence of his claim of ontological independence. This answer is one which, in the development of my anti-realist argument, I am set to oppose, but it is not something whose falsity I am at present assuming. All I am currently assuming, on the issue of empirical access, is the uncontroversial point, accepted by the realist and non-realist alike, that something will not satisfy our identifying conception of the physical world if such access is lacking.

Second, it might be objected that, in claiming that our understanding of physical spatial arrangement commits us to recognizing the overall soundness of our ordinary methods of empirically gauging such arrangement, I am ruling out, as ultimately unintelligible, certain forms of radical scepticism, and that, in the present dialectical context, this is unwarranted. Consider, for example, the scenario envisaged by Descartes in which there is no physical world, and in which all our sensory experiences are directly caused by a powerful malicious demon, who is intent on creating the systematic illusion of a world.⁴ If such a scenario obtained, all the empirical tests that we conducted, and all that we could conduct within the constraints of the demon's deceptive intent, would bear out our belief in the existence of a physical world, and a world of a specific kind, but, *ex hypothesi*, such a belief would be false. So, if envisaging the scenario is coherent, my claims about the overall soundness of our methods of gauging physical spatial arrangement, and the logically decisive character of the verdicts they deliver at the theoretical limit, seem

⁴ Descartes envisages this scenario in his First Meditation. See his *Meditations on the First Philosophy*, in *The Philosophical Writings of Descartes*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch, ii (Cambridge: Cambridge University Press, 1984).

to be mistaken: it seems that all the relevant actual and potential empirical evidence could unequivocally support certain conclusions about how things are physically spatially arranged, but these conclusions be false. As for the coherence of what Descartes envisages, there could, I suppose, be an issue over the intelligibility of the notion of a demon. But even if this notion turns out to be problematic, I do not see how, without begging the question, I can deny the intelligibility of supposing that there is *some form* of external agency that causally controls and organizes our experiences in the relevant way, but without the real existence of a physical world. The only way, I think, in which I could find a basis for such a denial would be by embracing a suitable form of canonical idealism—a form in which the mere presence of an external agency that imposes the appropriate organization on human sensory experience would, set in the context of the relevant endowments of the human mind, suffice for the existence of a physical world. But, even if I were happy with this form of idealism—and I shall indicate later that I am not—I would not be entitled to assume its truth, or indeed the truth of any idealist position, at this stage in the discussion, where I am still in the process of trying to develop an argument against the realist.

This objection, too, involves a basic misunderstanding of the position I have taken. The overall soundness that I am claiming that we have to recognize in our ordinary methods of empirical assessment has nothing to do with the evidential situation that would arise if Descartes's scenario, or anything else of that ilk, obtained. For when I speak of these methods, and the spatial conclusions they support, I am taking for granted the existence of the physical world and its accessibility to empirical investigation, and I am thinking exclusively of methods of assessment that pertain to that accessibility—methods that involve genuine sense-experiential contact with the physical world, of the sort that we ordinarily take to be perceptual. So, in claiming that the recognition of the overall soundness of these methods is implicit in our very understanding of the nature of physical spatial arrangement, and that the verdicts they deliver become, in consequence, logically unchallengeable at the point where there is no further room for an empirical challenge, I am not implying that the scenarios of the radical sceptic are unintelligible; I am not, at this point, engaging with the claims of the radical sceptic at all.

Finally, it might be objected that, while the physical world has, in the relevant sense, to be our world empirically, and this, in turn, gets reflected in our understanding of what is involved in physical spatial arrangement,

the nature of this reflection is more subtle than I have allowed. Certainly, for the world to be ours empirically, there has to be, in general, agreement between how its spatial structure is empirically represented at the human viewpoint and how it really is, and this will inevitably link our conception of such structure with the empirical methods by which we gauge what forms of structure obtain. But the presence of a *general* agreement leaves room for exceptions; and, in consequence, it might be thought that, so long as they are sufficiently rare, our understanding of physical spatiality even allows for cases in which some aspect of physical spatial arrangement is totally belied by how things are disposed to appear at the human viewpoint—cases in which some issue of arrangement is thoroughly open to empirical investigation, but in which all sufficiently searching empirical tests would unequivocally support a false conclusion. If this were right—and it has a certain initial plausibility—then what I have claimed about the overall soundness of our empirical methods of spatial assessment, and the logically decisive status of the verdicts they deliver at the point where the potential for any significant contrary evidence has been excluded, would be too strong. In the particular case on which we are focusing, our understanding of the nature of physical spatial arrangement, though conditioned by the need for the physical world, as *our* world, to have a basic empirical immanence, would allow us to entertain the supposition that the locations of Oxford and Cambridge conform to the external topology, rather than to what gets projected on to the human viewpoint.

One way in which I could comfortably deal with this objection would be by changing the terms of the relevant example in a way that avoids it, or at least draws its sting. Thus, instead of envisaging just one pair of S-regions where everything is organized as if they were interchanged, I could suppose such cases to be a commonplace feature of E. And, for added security, I could even suppose that, despite the organizational anomaly involved in each case, there was a general pattern to which the various cases conformed—for example, that the pairs of spherical regions were of constant size, that there was the same distance between the regions in each pair, and that the cases were distributed over S in a uniform fashion. Developing the example in this way would not, it seems to me, affect our judgement that there was still a physical world, and one without gaps, and, since all the cases would be, at the level of what fundamentally obtained, of exactly the same type, they would have to be physically interpreted in the same way. But, even if the objection had some force in the case of the original example, it would become

impossible to suppose that, on such a scale and in such a regular way, things were organized so as to ensure that the topology of what qualified as our world was systematically belied by how things were disposed to appear at our viewpoint. Granted that we accept—as does the envisaged objector—that the physical world has to be our world empirically, and that this conditions our understanding of the nature of physical spatial arrangement, we could not, with any plausibility, refuse to accept, on the new version of the example, that the physical topology coincided, across the whole range of cases, with the functional—organizationally simulated—topology of *S*, in line with what was empirically projected.

If necessary, then, I would be happy to revise the example in the way indicated in order to make my interpretation of it secure. As it is, I think that the objection is misconceived and that the example will serve my purposes in its original form. Such initial plausibility as the objection has comes, I think, from a failure to draw a sufficiently sharp distinction between two ways in which it could be thought that the general conformity of how things physically spatially are to how they are empirically represented admits of exceptions. On the one hand, there is the uncontroversial claim that the actual empirical evidence we possess with respect to some issue of physical arrangement can sometimes be misleading and support an erroneous conclusion. I have already conceded that there are all sorts of familiar ways in which this can occur. On the other hand, there is the much stronger claim that features in the objection, that, even at what I have been referring to as the theoretical limit, where an issue of arrangement is thoroughly open to empirical investigation, and where all sufficiently searching tests (actual and potential) would unequivocally deliver the same verdict, this verdict may be in error. The objector is, in effect, supposing the type of exception envisaged by the second claim to be just an extreme version of the ordinary types of exception envisaged by the first, and something that our understanding of physical spatiality can ultimately accommodate on the same basis. And this, in my view, is his fundamental mistake. For, although the first claim allows us to envisage a spectrum of types of case of empirical misrepresentation, in which the evidence we possess in support of some erroneous spatial conclusion varies in strength, we confront something of a crucially different kind at the theoretical limit. Thus, whereas in the ordinary cases, however strong the misleading evidence, it is left theoretically possible that the sum total of actual and potential relevant evidence will, overall, support a different conclusion, such a possibility is explicitly excluded by the sort of case that the objector

is trying to envisage at the theoretical limit. It is this exclusion, it seems to me, that puts the case beyond the limits of what our empirically conditioned understanding of physical spatiality allows. Thus, it seems to me that the sense in which the physical world has to be our world empirically would be undermined if some aspect of its spatial structure could, in this way, be totally belied by how things are disposed to appear at the human empirical viewpoint, and that this is reflected in our understanding of what such spatial structure involves.

IV

Let us speak of the case of R1 and R2 as the *interchange case*. I have introduced this case, and argued for a certain conclusion as to how it should be physically interpreted, as a stage in the development of my argument against physical realism. And its relevance to the issue of realism is already, in general terms, apparent. Realism accepts the existence of a relevant external reality and equates it with the core physical reality—with that portion of the total physical reality covered by the combination of what is physically fundamental and what is wholly constituted by what is physically fundamental. So, in the hypothetical case envisaged, a realist construal of the situation would have to equate all the ontological and factual ingredients of E with their counterparts in the core physical reality, and this would mean, in particular, identifying the external space S with physical space. But on my interpretation of the situation these two spaces cannot be numerically the same since, relative to a correlation between corresponding points and regions, they have different topologies. Specifically, the region, R1, in S which corresponds to the Oxford region in physical space is not located in the region corresponding to the region in which Oxford is located, but rather in the region corresponding to the region in which Cambridge is located, and the region, R2, in S which corresponds to the Cambridge region in physical space is not located in the region corresponding to the region in which Cambridge is located, but rather in the region corresponding to the region in which Oxford is located. So the realist construal of the situation would fail.

Precisely what follows from this, and how serious it is for the realist, are things that I shall consider presently. But, first, I want to set the point I have been developing in a broader context.

The interchange case is just one example of a whole class of envisageable cases in which the structure and organization of the relevant external reality differ, in a certain respect, from what all our empirical evidence suggests and what, in response to this evidence, we ascribe to the physical world in our ordinary and scientific thinking. But within this class it is also an example of a more specific range of cases that meet three further conditions, and which, as a result of meeting them, are of particular relevance to the anti-realist argument I want to develop.

The first condition is that the relevant difference between the structure and organization of the external reality and those of the putative physical world which the empirical evidence suggests has nothing to do with the fact that, at the level of ordinary observation, the empirical evidence suggests that the world has, in certain respects, a distinctively sensible character. As we have seen, the ultimate status of such evidence depends on whether it is evaluated from a realist or an idealist standpoint. From a realist standpoint, the evidence gets ultimately overruled by the fact that a scientific investigation of the situation reveals that the ascription of distinctively sensible features to items in the physical world has no explanatory value (either with respect to how things behave within the world or with respect to how things sensibly appear to us). From the standpoint of canonical idealism, where the character of the world is fixed by the way things are disposed to empirically appear at the human viewpoint, the fact that things are systematically disposed to sensibly appear to us in certain ways can be allowed to make its own distinctive contribution to this character. But what, in the present context, matters is that any issue over the evidential status of sensible appearance is not relevant to the kinds of case that I am trying to identify. In these cases the empirical evidence that is relevantly out of line with the external reality relates to structural and organizational aspects of the empirically suggested world that have nothing to do with the presence of qualitative items of a distinctively sensible sort.

The second condition is that the difference between the structure and organization of the external reality and that of the putative physical world that is empirically suggested is not so great, overall, as to lend any plausibility to the conclusion that there is, under what is envisaged, no physical world at all. It is not difficult to think of differences whose extent might tempt us to that conclusion. For example, if what we envisaged as the external reality contained only a two-dimensional space, then, even if it were organized in a three-dimensional way, we might pause to wonder whether we had

strayed too far from the structure of a physical world to allow us to retain our ordinary belief in the existence of such a world. Certainly, this is likely to be the response of someone who approached the issue from an initial position of realism. Even more obviously, if what we envisaged as the relevant external reality was not world-like at all, but consisted of some supernatural personal being, or group of such beings, with direct volitional control over the character of our sensory experiences, we would be likely to need some persuading to think that this was reconcilable with the existence of the physical world. I am not saying that these responses to the envisaged cases would be justified: this is not a matter that I can appropriately address at this stage. But for the purposes of my present argument I need cases that are clear-cut—cases that the realist cannot just set aside by claiming that they do not call for physical interpretation at all.

The third condition, and the one most crucial to the whole point of the interchange case, and to the way in which I am claiming we should interpret it, is that the reason why there is the difference between how things stand in the external reality and how things empirically appear is that the external reality is nomologically organized as if, by the standards of what would be required for organizational uniformity, it had a different structure from the one that it actually has, and it is this organizationally simulated structure which gets projected on to the human empirical viewpoint. This means that the difference is not just accidental—contingent on the particular external conditions obtaining—but nomologically guaranteed, so that there are no possible circumstances in which, without an alteration to the laws, the true character of the external situation would become empirically manifest.

Where the structural and organizational difference between the envisaged external situation and how things empirically appear satisfies these three conditions, I shall speak of it as a case of *relevant deviance*.

We have so far only identified one example of a case of relevant deviance, but it is not difficult to think of further examples—indeed, the potential list is infinite. Some of these examples would again be concerned with the geometry of space. For instance, and very similar in general character to the interchange case, we could suppose the external reality to be organized, in respect of both its internal behaviour and its causal relations with the human mind, exactly as if a certain spherical region of space were rotated, relative to its surroundings, by a certain amount in a certain direction. Or, changing the nature of the external reality more radically, we could

suppose that the external item corresponding to the notional space of our ordinary beliefs consists of two separate three-dimensional spaces organized as if they were joined, each space corresponding to one of two portions into which the notional space can be exhaustively divided. Other examples might concern different forms of structure, such as relations of qualitative sameness and difference. For instance, we could suppose that in the external reality there are two types of particle which are nomologically required to exchange their current spatial distributions and dispositional properties every hour, so that everything in the reality is organized, internally and in its causal relations with human mentality, exactly as if each of the types in hours 1, 3, 5, 7, ... is the same as the other in hours 2, 4, 6, 8, ... In each of these cases, the envisaged organizational anomaly ensures that the structure and organization of the external reality differs in a certain respect from what the empirical evidence suggests. And, assuming that this is the only relevant difference, we can then take all three cases to satisfy the three specified conditions for relevant deviance: they are explicitly devised so as to satisfy the first and third, and, in each case, it is clear that the extent of the difference between how things are envisaged to be externally and how they empirically appear is not such as to threaten belief in the existence of a physical world.

These cases are like the interchange case, in being examples of relevant deviance, and if I was right in my physical interpretation of the interchange case, it is clear that we need to adopt a similar mode of interpretation in these further cases too. In other words, in each case we need to take the physical structure to coincide with the organizationally simulated structure of the external reality, rather than with its actual structure. So, if we suppose the external reality to comprise two spaces organized as if they were joined, we need to conclude that there is only one physical space, reflecting the distinctive (single-space simulating) character of the organization. And if we suppose that two types of external particle are nomologically required to exchange their current distributions and dispositional properties every hour, we need to trace the spatio-temporal paths of the relevant physical particles in a way that achieves organizational uniformity, rather than in a way that matches the external relations of qualitative sameness and difference. In fact, it is clear that this is how we need to decide the issue of physical structure in all cases of relevant deviance. The underlying point is simply a generalized version of what has already featured in the earlier discussion. The physical structure must conform to the organizationally simulated structure of the external

reality because it is this structure, not its actual structure, that is projected on to the human empirical viewpoint, and because our ordinary understanding of what is involved in each form of physical structure is conditioned by the requirement that the world which that structure characterizes be our world empirically.

V

Let us now return to the question of how all this bears on the issue of physical realism. If my argument has been sound, we know that in any case of relevant deviance there is some respect in which the structure of the physical world differs from the structure of the relevant external reality. This already seems to pose a problem for physical realism. Realism is committed to accepting the existence of a relevant external reality and identifying it with the physical world in its primary core; and this involves taking the structural aspects of this reality to form the corresponding aspects of physical structure. So, for cases of relevant deviance, the correct account of how things stand physically and the realist account are in direct conflict. We have already noted this conflict in the example of the interchange case: a realist construal of the situation identifies physical space with the external space *S*, while, on what I am taking to be the correct account, these spaces are not identical, since the locations of the *S* regions *R*₁ and *R*₂ are different from the locations of the corresponding physical regions of Oxford and Cambridge.

This result looks awkward for the realist view, but perhaps it does not, as it stands, amount to a refutation. Strictly speaking, the realist thesis, as we have formulated it, is only concerned with the *actual situation*: it takes the physical world to be something whose existence *is* logically independent of facts about human mentality and *is* philosophically fundamental. But the cases of relevant deviance are, in the context of our discussion, only *hypothetical*: they represent ways in which we can coherently envisage how the situation *might be*. So, the realist could still insist that such cases do not *in fact* obtain and that his account of the physical world is, *as things are*, correct. Nor would such a response be merely perverse. For, as we noted in the interchange case, there are good grounds for supposing that the kind of deviance envisaged does not obtain: as well as running counter to what all the empirical evidence suggests, there would be something inherently puzzling about a situation in which the external organization was anomalous in this kind of way. So, perhaps the

realist can afford to concede that cases of relevant deviance are not amenable to a realist construal, and just insist that a realist account is correct for how things actually are, or at least for how we are entitled to take them to be.

Whether the realist can really afford to limit the relevance of his claims in this way is, I think, very doubtful. There may be good reasons for concluding that the external reality is highly unlikely to be subject to the kinds of relevant deviance envisaged. But we certainly cannot exclude the possibility that it is, and it would surely be awkward for the realist if his view of the world had to be, to some extent, speculative in this way—if its correctness had to depend on the outcome of some issue of fact that we cannot definitely settle. In any case, it would be odd if realism were the right philosophical account of the world as it actually is, or as we can reasonably suppose it to be, but did not generalize to how things would stand if one of the less probable outcomes obtained. We surely expect realism, as a *philosophical* thesis, to be a response to our fundamental conception of the physical world, rather than to what we take to be the most plausible account of the specific circumstances of its existence.

I am inclined to think, then, that the conclusions we have already reached about the hypothetical cases of relevant deviance are enough to undermine, or at least to seriously damage, the realist view. But, for the sake of argument, let us suppose that I am wrong about this, and that the realist can afford to respond to the situation in the way suggested—dismissing the cases of deviance as ones that we need not seriously entertain. Even so, I think we can show that the underlying difficulty for his position remains. For, when we look into the issue more closely, I think we can see that the factors which make cases of relevant deviance resistant to a realist construal carry over to cases to which the realist's approach seems better suited. In particular, they carry over to the case which represents the realist's best scenario. This is the situation where—ignoring any respects in which evidence at an observational level suggests the presence of a world with distinctively sensible features—the structure and organization suggested by the actual and potential empirical evidence exactly match the structure and organization of the external reality. In other words, it is the situation where the external reality systematically, and in every detail, projects its own structure and organization on to the human empirical viewpoint. I shall speak of this as the case of *zero deviance*. What makes this case so propitious from the realist's standpoint is that, with respect to all aspects of structure and organization that feature in his understanding of the world, it contains nothing that could even be thought to indicate that

there is a qualitative difference between the external and physical realities, and so nothing that presents an immediate obstacle to identifying the two. Indeed, the considerations that obliged us in the cases of deviance to conclude that the two realities differed in structure and organization would here oblige us to conclude that, in all relevant respects, they coincide. So, if we can suppose the actual situation to be one of zero deviance and still show that the realist cannot provide an acceptable account of it, that will suffice to refute realism altogether.

In order to be able to evaluate how realism fares with the case of zero deviance, we need to begin by looking a little more closely at what is going on in the cases of relevant deviance. So far, in discussing these cases, we have focused almost exclusively on the issue of how we should decide what structural properties to ascribe to the physical world, given a supposition of the deviance in question. And we have reached the conclusion that we should take the physical structure to conform to the organizationally simulated structure of the external reality—the form of structure that this reality projects on to the human viewpoint—rather than to its actual (intrinsic) structure. What we have not yet considered, in any detail, is what consequences this conclusion has for our understanding of the *metaphysics* of the situation. We have noted, of course, that the qualitative difference between the relevant aspects of the physical and external structures excludes a realist construal of the situation: it excludes an identification of the physical and external realities that the differing structures characterize. But we have not yet offered any positive account of how these distinct realities are metaphysically related; and, although I do not want, at this stage, to pursue this issue in depth, there are certain straightforward points that I want to highlight, because of the light they shed on the crucial case of zero deviance.

Let us start by focusing on the familiar example of the interchange case. The question we had to settle, in this case, was whether, when we suppose the external reality to be as the example envisages, we should take the arrangement of regions in physical space to accord with how things are systematically disposed to appear at the human empirical viewpoint, so that we assign the Oxford region to a location in Oxfordshire and the Cambridge region to a location in Cambridgeshire, or should, rather, take the physical arrangement to mirror the arrangement in the external space, where the region corresponding to Oxford is enclosed by a region corresponding to a portion of Cambridgeshire, and the region corresponding to Cambridge is enclosed by a region corresponding to a portion of Oxfordshire. The conclusion we

reached, for reasons elaborated, was that we need to model the physical arrangement on how things are disposed to empirically appear, thus leaving a sharp discrepancy between the positioning of the Oxford and Cambridge regions in physical space—the positioning of the two cities in the geographical layout of England—and the positioning of the corresponding regions in the external space. The question I now want to ask is: Given this discrepancy, how are we to account for the obtaining of the physical arrangement? What gives the regions of physical space the arrangement they have? Of course, the points and regions of physical space, as of any genuine space, would not be the points and regions they are if their locations were different: in that sense, the character of the physical arrangement is logically guaranteed by the identities of the regions involved. But what is at issue is how we are to account for the fact that the arrangement of physical regions has the character it does in relation to the arrangement of the corresponding external regions. So, in the specific area of interest, what brings it about that the physical region corresponding to R1 is located in surroundings that correspond to the surroundings of R2, and that the physical region corresponding to R2 is located in surroundings that correspond to the surroundings of R1?

It is clear that the obtaining of the physical arrangement does not logically stem from the external arrangement on its own. If it did, there would be no accounting for how the two arrangements come to differ. The external arrangement is logically relevant to the obtaining of the physical arrangement: it logically contributes to it. But another factor that also logically contributes to the obtaining of this arrangement (to the fact that the physical regions are arranged in the way they are in relation to the arrangement of the corresponding external regions) is the way in which the external reality is functionally organized by the relevant laws—laws that cover both its internal behaviour and its causal links with human mentality. The nature of the contribution which this organization makes to the physical outcome is easy to discern. It is this organization which, as it were, operates on the external arrangement and converts it into the revised form of arrangement that gets projected on to the human empirical viewpoint; and so it is this organization which, in partnership with the external arrangement, determines what form of arrangement (in relation to the corresponding external arrangement) has the right empirical credentials—the right kind of empirical immanence—to qualify as the actual (intrinsic) arrangement in our world. We can also see that, between them, the external arrangement and the deviance-creating organization logically *suffice* for the obtaining of the physical arrangement.

Or, at least, that they do so in the framework of those endowments of the human mind that make us empirically receptive to the orderly character of our sensory experiences (for without that receptiveness there would be no empirical viewpoint to form the target of the relevant projection), and with the backing of any other factors that may be logically needed for the existence of a physical world. The reason why I add this second qualification is not something which I want to pursue at this stage; nor do I need to for the purposes of the present discussion. But I shall eventually argue that the mere presence of a world-like external reality with an appropriate internal character, and with appropriate modes of projection on to the human empirical viewpoint, is not enough to secure the real existence of a physical world.

The same situation holds for all cases of relevant deviance. So, if we envisage an external reality which contains two three-dimensional spaces organized as if they were joined, we must see the unispacial structure of the physical world as logically stemming—in the framework of certain further factors—from the combination of the presence of the two external spaces and their unispacial organization. Likewise, if we envisage an external reality in which two types of particles are nomologically required to exchange their current spatial distributions and dispositional properties every hour, we must see the spatio-temporally continuous histories of the corresponding physical particles as logically stemming—in the framework of certain further factors—from the combination of the discontinuous histories of the external particles and the compensating aspects of the organization (aspects which, as it were, make the discontinuities, overall, functionally inert). And this is how things stand in cases of relevant deviance quite generally. To set out the situation precisely, let us speak of the *exclusive* external organization when we want to refer to how the external reality is organized internally (in respect of what takes place within it), and let us speak of the *inclusive* external organization when we want to refer to how the external reality is organized both internally and in its causal relations with human mentality. (So, in the interchange case, the exclusive organization is what is created by the E-laws, and the inclusive organization is what is created by the E-relevant laws.) Then, for any case of relevant deviance, there is a certain respect in which the structure of the physical world differs from the structure of the external reality. And, with respect to that difference, the situation is such that (i) as well as logically depending on the external structure, the obtaining of the physical structure logically depends on relevant aspects of the inclusive external organization; (ii) the way

in which the inclusive organization logically contributes to the obtaining of the physical structure is by combining with the external structure to project a revised form of structure on to the human empirical viewpoint; and (iii) taken in the framework of the relevant endowments of the human mind, and with the backing of any other factors that may be logically needed for the existence of a physical world, the combination of the external structure and the inclusive organization logically suffices for the obtaining of the physical structure.

With this account in place, let us now return to the crucial case of zero deviance, where the form of structure that is projected on to the human viewpoint coincides, in all relevant respects, with the structure of the external reality. This case, as I have said, represents the realist's best scenario. For we are now obliged to take the relevant physical and external structures to be qualitatively the same, and this seems to free the way for a realist construal of the situation, in which the physical and the external realities are identified. But the further points that have now come to light about the cases of relevant deviance put a different complexion on things. For once we have accepted that, in the cases of deviance, the character of the physical outcome logically depends not just on the character of the external structure, but also on the nature of the inclusive organization, we surely have to accept that there is an analogous dependence on this organization in the case of zero deviance. Thus, if in the cases of deviance the organization logically contributes to the physical outcome by combining with the external structure to ensure the projection of a certain revised form of structure on to the human viewpoint, then surely in the special case of zero deviance we must see the organization as analogously logically contributing to the physical outcome by combining with the external structure to ensure the projection of the same (unrevised) form of structure on to the human viewpoint. It is surely clear that the two types of case are exactly parallel, and that the contributory role of the organization in the context of zero deviance, where, as it were, it endorses the external structure, is just a limiting case of its role in the context of the relevant kind of deviance.

In fact, the contributory role of the organization in the case of zero deviance becomes quite explicit when we take account of two further points. The first is that, even when we assume the actual situation to be one of zero deviance, we can envisage various ways in which the organization *could have been* such as to create a case of relevant deviance; and, crucially, we can envisage ways in which such deviance could have occurred without change to the relevant

aspects of the external structure. For example, even when we assume the functional topology of the external space to be wholly in line with its intrinsic topology, and, in consequence, assume this intrinsic topology to be wholly in line with what gets projected on to the human empirical viewpoint, we can still envisage how, without change to this intrinsic topology, or indeed to any aspect of the space's intrinsic geometry, things could have been organized, both in respect of what takes place within the space and in respect of causal relations with human mentality, exactly as if two congruent regions were interchanged. And, obviously, there is a whole range of further examples of a similar kind. The second point is that, in cases of relevant deviance, as we noted, the inclusive external organization not only logically *contributes* to the relevant physical outcome, but also, in partnership with the external structure, and in the framework of the relevant endowments of the human mind, logically *ensures* it, or, at least, does so with the backing of any other factors that may be logically needed for the existence of a physical world. Thus, in the interchange case, with the relevant endowments and relevant backing in place, the combination of the external arrangement of regions and the distinctive, deviance-creating organization logically determines the character of the arrangement of physical regions, in relation to the arrangement of the corresponding external regions, and, in particular, ensures the reversed locations of the Oxford and Cambridge regions in relation to their external correlates. Putting these two points together, we can see that, even in the case of zero deviance, the character of the inclusive organization must logically contribute to the physical outcome. For if the organization had, instead, been such as to create some form of relevant deviance, though without affecting the relevant aspect of the external structure, and without disturbing either the relevant endowments of the human mind or the presence of any other factors that are logically needed for the existence of a physical world, then the structure of the physical world would have been correspondingly different. And this means that the character of the actual physical structure must be logically due, in part, to the fact that the actual organization is not of this deviance-creating sort.

It is clear, then, that, even when we take the actual situation to be one of zero deviance, which provides the realist with his best scenario, a realist construal of the situation fails. The realist wants to say that the external reality forms the physical world in its primary core. But the external reality will only be identifiable with the core physical reality if its structure can be equated with the structure of this reality, and the phenomenon of organization

dependence that has now come to light means that such an equation is excluded. Thus, whether there is deviance or not, what gives the physical reality its various forms of structure are not the corresponding forms of external structure on their own, but the combination of those external forms and the inclusive external organization. For it is this organization which, by operating on the external forms of structure, in the framework of the endowments of the mind, determines what gets projected on to the human empirical viewpoint, and it is the nature of this projection that fixes the character of the forms of physical structure. The only thing that is distinctive about the case of zero deviance is that the forms of structure organizationally projected involve no revision of the forms that are externally realized, so that the physical forms of structure are made to match their external correlates. So, even in the case of zero deviance, the structural aspects of the physical reality cannot be equated with the corresponding aspects of the external reality, since they depend, for their obtaining, on something further. They are put in place by a combination of these external aspects and the inclusive organization.

VI

Physical realism has been shown to be untenable. The argument I have employed to show this has been complex, and, before going any further, I think it would be helpful if I provide a summary of the steps that have brought us to this point.

1. I defined the *core physical reality* as that portion of the total physical reality which is covered by facts that are either physically fundamental or wholly constituted by facts that are physically fundamental. I also stipulated that something qualifies as a *relevant external reality* if and only if (i) its character is, through and through, logically independent of facts about human mentality, (ii) the facts that feature in it are in no instance constitutively reducible to facts that lie outside it, and (iii) it is responsible for the systematic control of human sensory experience. By taking the physical world to be ontologically independent of the human mind and something whose existence is philosophically fundamental, the realist is committed to taking the core physical reality to be a relevant external reality.

2. We can envisage hypothetical situations in which the structure and organization of the relevant external reality differ from those of the putative

world which is empirically suggested at the human viewpoint. And, within the total range of such hypothetical situations, we have identified a particular group whose members satisfy three further conditions, namely: first, the difference in question has nothing to do with respects in which the empirical evidence may suggest the presence of distinctively sensible features in the world; second, the extent of the difference is not such as to lend plausibility to the conclusion that, if what is envisaged obtains, there is no physical world; and third, the reason for the difference is that the external reality is nomologically organized as if, by the standards of what would be required for organizational uniformity, it had a different structure from the one it actually has, and it is this organizationally simulated structure that gets projected on to the human empirical viewpoint. Situations which satisfy these three conditions are what I have termed cases of *relevant deviance*.

3. In cases of relevant deviance the core physical reality and the relevant external reality cannot be identified, since the structure of the physical reality conforms to what is empirically projected, rather than to the structure that is externally realized. So, in the interchange case, in which everything in the external reality is organized, both internally and in relation to human mentality, exactly as if the regions R1 and R2 were interchanged, the locations of the Oxford and Cambridge regions are respectively in Oxfordshire and Cambridgeshire, in line with how things appear empirically, rather than in the surroundings that correspond to the surroundings of their external counterparts. This means that realism cannot accommodate such cases.

4. This result creates problems for the realist, but may not undermine his position entirely: he could still insist that realism provides the correct account of the actual situation, or of what we are entitled to take to be such. But, by looking into the metaphysics of how things stand in cases of relevant deviance, we have seen that, even in the case which represents the realist's best scenario, where there is no relevant disagreement between the character of the external reality and how things are disposed to appear at the human empirical viewpoint, and where, in consequence, we have to take the physical and the external realities to relevantly coincide in structure and organization, the two realities still cannot be identified. For, even in this case—the case of *zero deviance*—the physical reality logically derives its structure not from the structure of the external reality on its own, but from the combination of the external structure and the inclusive external organization, or, more precisely, from this combination, together with the support of the relevant

endowments of the human mind, and any other factors that may be logically needed for the existence of a physical world. In other words, the two realities cannot be identified, because the structure of the physical reality logically depends, for its obtaining, on more than just the structure of the external reality. But if realism cannot handle the case that represents its best scenario, it fails altogether.

Throughout this argument I have assumed that the realist accepts a certain account of the nature of the core physical reality—an account which represents it as composed of certain kinds of ontological and factual ingredient at its fundamental level. But, as I stressed at the outset, this assumption was simply to provide a concrete point of focus, and not because the success of the argument depended on it. The argument can be adapted to any account of the physical reality that represents it as having a spatial structure and an internal organization. As I see it, this means that it can be adapted to any account that the realist could offer that has any claim to plausibility.

This refutation of physical realism is not quite the end of the present story. For, as I indicated, I intend my argument against realism to serve as an argument for the idealist alternative. But before I turn to that, and indeed partly by way of preparation, I want to bring out something further about the failure of the realist position. Specifically, I want to bring out the full extent of this failure.

VII

The realist is committed to taking the core physical reality to form a relevant external reality, and my argument against him has turned on the claim that, even if there is such an external reality, it cannot be identified with the physical reality. But the physical and external realities are complex: the fact that they cannot be identified does not entail that they have no components in common. In that sense, the realist could still, in theory, hope to salvage something from his position by continuing to insist that *certain* physical entities are external entities or *certain* physical facts are external facts. He could hope to be able to say that, while physical realism is false, it is not, as it were, false through and through, and that a realist account of physical phenomena—ontological or factual—is defensible in places. However, it only takes a little reflection to see that, once physical realism itself has been

discredited, even this semi-realist position is untenable. To give the position its best chance of succeeding I shall, in examining it, continue to assume that there is a relevant external reality, and one that is world-like in its general character, so that there are items on the external side that we can at least think of as candidates for identification with physical items.

Being world-like in its general character, the external reality, like the physical reality, contains a three-dimensional space, and I want to begin by considering the relationship between this external space and physical space. We already know that these spaces are numerically distinct, since, even in a case of zero deviance, the arrangement of regions and points in physical space logically depends, for its obtaining, on more than just the arrangement of the corresponding external regions and points. But given that these spaces are distinct, it is also clear that they are distinct through and through: no region or point in one is numerically the same as any region or point in the other. For the identities of the regions and points of a space cannot be separated from the network of spatial relations in which they stand to one another: to envisage a region or point as occurring in two distinct spaces would be to envisage it—absurdly—as having two identities. In any case, the considerations that show that the external and physical spaces are numerically distinct would also show, quite directly, that this distinctness applies to all their elements. For the organization dependence that we found to characterize the arrangement of regions and points in physical space will also characterize the arrangement of regions and points in any region of physical space—thereby distinguishing that region from any external correlate.

Next, we must notice that, for an occupant of a concrete space, its occupancy of that space is the very form of its existence—its mode of being, as we may put it—and this means not only that it cannot exist without location in that space, but also that it is incapable of simultaneously occurring in a different concrete space, in which it would possess an additional form of existence; for existence in a different space would simply create a different individual. So, given that the physical and external spaces are wholly distinct, the occupants of each must be entirely distinct as well. The same point will hold for any other type of concrete space-located item. Thus, given the complete distinctness of the physical and external spaces, it is clear that no physical spatial event is identifiable with an external spatial event, and no physical instance of a region-characterizing property is identifiable with an external instance. This allows us to conclude that, apart from the dimension of time, the physical and external realities are, ontologically, entirely distinct: no

concrete physical entity—whether an element of space or something located in space—is numerically the same as any external entity.

Finally, given that the physical and external spaces are entirely distinct, and that the entities located in these different spaces are, in all instances, distinct as well, we know that the two realities are also entirely distinct in their factual ingredients. For any physical fact will either be a fact about physical space or a fact about something existing or occurring within it. So, we know that the physical and the external realities are distinct through and through: they contain no common components, ontological or factual, other than ones that are covered by the mere dimension of time. It follows that the failure of realism is total. A realist account fails not only with respect to the physical world as a whole, but also with respect to every ingredient of the world that is of a distinctively physical kind. The physical reality resists a realist construal at every point.

VIII

Given that the core physical reality is entirely distinct from any relevant external reality, how are we to make metaphysical sense of its obtaining at all? Addressing this question will bring the argument to its ultimate conclusion, in which the considerations that have obliged us to reject a realist view of the world, and to reject it through and through, will oblige us to accept the idealist alternative. The question itself, of course, presupposes that there is a physical reality whose obtaining is metaphysically at issue, and, as we shall see, this presupposition could be challenged. But for the time being I shall put any doubts over the existence of this reality on one side.

It is clear that, given its total resistance to a realist construal, we shall now have to think of the core physical reality, and indeed the physical reality as a whole, as something which is, through and through, constitutively sustained. We shall have to accept that all physical facts are ultimately constituted by non-physical facts—facts that are devoid of physical entities and physical properties—and that, for all physical entities, their existence is ultimately constituted by non-physical facts. It is not that we are forced to this conclusion by logic alone. But given that the physical reality is not an external reality, and is indeed separate from any such reality at every point, the only alternative to taking it to be something constitutively sustained would be to identify it with something within the realm of human mentality—something which

draws its ontological ingredients from human-mental entities, and draws its factual ingredients from human-mental states of affairs. And there is no such identification that is remotely plausible. Indeed, without a gross distortion of what we know about the realm of human mentality, I do not think we can even find something in that realm which comes near to exemplifying the structure and organization of the physical reality.⁵

We know that the core physical reality is, through and through, constitutively sustained by how things stand in some non-physical reality. And in the kinds of case on which we have been mainly focusing, where we assume the presence of a world-like relevant external reality, and suppose the situation in relation to that reality to be one of zero or relevant deviance, the general form of the sustainment involved is easy to discern. We have already seen how, in such cases, the structure of the physical reality is logically determined by the combination of the structure and inclusive organization of the external reality, or, more precisely, by this combination, together with the support of the relevant endowments of the human mind, and any other factors that may be logically needed for the existence of a physical world. So, given the presence of these endowments, and, if they are needed, these further factors, the arrangement of regions in physical space is logically determined by the combination of the arrangement of the corresponding regions in the external space and the way in which the external reality is organized, internally and in relation to human mentality, in respect of this arrangement. Let us, for convenience, refer to the relevant endowments of the mind, and any other factors that may be logically needed for the existence of a physical world, as the *supporting factors*. Now logical determination need not be constitutive, since the obtaining of what is logically determined may not satisfy the requirement of involving nothing over and above the obtaining of the determining factors. To take a case I mentioned when I first set out the conditions for constitution, if an essentially omnipotent God prescribes that a certain type of event will happen at a certain time, this prescribing will logically determine the subsequent occurrence of the relevant

⁵ Some of the things that Berkeley says in the course of expounding his doctrine of *esse est percipi* could be thought to suggest a position which equates the physical world with something in the realm of human mentality. Specifically, they could be thought to suggest a position which equates the world with an organized collection of our sensory ideas (of the internal objects of our sensory perceptions). For an identification of the obvious inadequacies of such a position, and for a discussion of the overall thrust of Berkeley's claims and arguments about the physical world, see my 'Berkeley on the physical world', in J. Foster and H. Robinson (eds.), *Essays on Berkeley* (Oxford: Oxford University Press, 1985), 83–108.

type of event, but the fact of its occurrence will not be constituted by the fact of the prescribing, because its obtaining will involve something genuinely additional to the fact of the prescribing.⁶ But in the present instance it is obvious that the logical determination in question is constitutive. The facts of physical structure not only logically stem from, but are constituted by, the structural and organizational facts about the external reality, together with the relevant supporting factors. And once we have accepted that this is the form of the constitutive sustainment for the structural aspects of the physical reality, there is no way of avoiding the conclusion that this is the form of the sustainment for all aspects. There is no way of avoiding the conclusion that, in the kinds of case under consideration, the structure and inclusive organization of the external reality, with the backing of the relevant supporting factors, are responsible for the constitutive sustainment of the entire physical reality in its primary core.

This raises the further question of how, in these cases of zero or relevant deviance, the external reality and its inclusive organization, together with the supporting factors, are equipped to constitutively sustain the physical reality. What, as it were, forms the metaphysical mechanism of this sustainment? How are we to see the sustainment as working? Once again, it is the determination of physical structure that provides the model. Thus, we saw that, in the cases of relevant deviance, the reason why the physical structure turns out to be, in the relevant respect, qualitatively different from the external structure is that the external structure and the inclusive organization, taken in the framework of the relevant endowments of the human mind, combine to ensure the projection of a different form of structure—the organizationally simulated structure—on to the human empirical viewpoint, and this projected form, in the context of its projective provenance, is then what fixes the relevant aspects of physical structure. And exactly the same story holds in the case of zero deviance, except that the form of structure projected by the external structure and inclusive organization is that of the external structure itself. It is clear that it is in similarly projective terms that we have to understand the constitutive sustainment of the entire core physical reality. The way in which, with the assistance of the supporting factors, the external reality and the inclusive organization suffice for the constitutive sustainment of the physical reality is by disposing things to appear systematically worldwise, and

⁶ See Ch. 1, Sect. II. An analogous example could be constructed by envisaging a situation in which the occurrence of a later event is logically determined by the combination of earlier conditions and laws of nature.

worldwise in a specific complex of ways, at the human empirical viewpoint, thereby projecting a certain form of world on to that viewpoint.

With the backing of the supporting factors, the external reality and its inclusive organization constitutively sustain the core physical reality by projecting a certain form of world on to the human empirical viewpoint, and they project this form of world by disposing things to appear systematically worldwise, in the relevantly specific ways, at this viewpoint. But to think of these factors as disposing things to appear thus at this viewpoint is the same as thinking of them as forming the concrete embodiment of the sensory organization. For the sensory organization is that system of control over the course of sensory experience that disposes it to conform to its world-suggestive pattern, and it is precisely by systematically controlling sensory experience in ways that dispose it to conform to this pattern that the external reality and its inclusive organization dispose things to present themselves, at the human viewpoint, in ways that empirically suggest the presence of a certain kind of world. The significance of this, in the context of our current discussion, is clear. It means that, in the kinds of case we are presently envisaging, where there is zero or relevant deviance, the metaphysical situation meets the requirements of phenomenalist idealism, which takes the physical world to be something whose existence is constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. And it means that it meets the requirements of such idealism in its canonical version, where the experiential facts that play this central constitutive role are ones that concern the content of the sensory organization.

To complete the argument, we now only need to take one further step. The case of zero deviance, as I have stressed, is the case that provides the realist with his best scenario—the case that represents the relationship between the external reality and how things are disposed to empirically appear at the human viewpoint in the way that is best suited to a realist view of the world. If even in this case it turns out that the requirements of canonical idealism are met, it is inconceivable that they should not be met in all cases. It is inconceivable that the sensory organization should ever fail to play the central role that the idealist assigns to it in the constitutive creation of the physical world, whatever the nature of the external factors that lie behind it, and, indeed, irrespective of other circumstances altogether. This is not to deny that the nature of what underlies the sensory organization may also be constitutively relevant to the existence of the physical world. We might

want to say that, unless this organization has a certain kind of provenance, or concrete realization, the overall situation will not suffice for the creation of a physical world. For example, we might want to say that, in order to secure the creation of a physical world, the sensory organization has to stem from the presence of an external reality of a world-like kind. I shall consider these sorts of issue later. All I am insisting at present is that, once we have acknowledged that the sensory organization centrally contributes to the constitutive creation of the physical world in the hypothetical circumstance of zero deviance, we cannot avoid concluding that it makes this same creative contribution whatever the circumstances, so long as there exists a world whose constitutive creation can be at issue.

If this is right, the only way in which we could now refuse to accept the canonical idealistic view of the physical world would be by refusing to accept the existence of such a world at all. And it is this that brings us to the next phase of our investigation. The existence of a physical world is something I have been taking for granted: I have confined my attention to the issue of how its existence is to be ultimately understood. But we shall shortly see that the very conclusion we have reached, in the framework of the assumption of the world's existence, could be thought to call that assumption into question. It could be thought that, if we are obliged to replace our ordinary realist understanding of the world by an idealist understanding, that, in itself, makes it impossible to suppose that there is a genuine world to have an understanding about. It makes it impossible to suppose that there is something that qualifies as a real physical world in any recognizable sense.

5

The Challenge of Nihilism

I

Physical realism has been discredited. It does not leave the physical world with the right kind of empirical immanence to count as our world in the requisite sense—a sense which allows it to be the world which our ordinary physical beliefs are about. This failure of realism does not just reside in its incapacity to provide a satisfactory account of physical-item perception—an incapacity which leaves the realist's world beyond the reach of our perceptual awareness. It also stems from the fact that realism does not accord the way things are disposed to empirically appear at the human viewpoint the requisite authority with respect to the character of the physical world. Thus, in the case on which we mainly focused, it allows for the possibility that, although there is a physical world, and one with which we have the normal means of sensory contact, the arrangement of regions within it is, in some respect, at variance with what all the actual and potential empirical evidence suggests. And allowing this possibility is in conflict with our empirical understanding of the nature of physical spatial arrangement.

The only way of giving the world the requisite empirical immanence is by embracing phenomenalistic idealism, which takes the physical world to be something whose existence is ultimately constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. And the only satisfactory way of embracing this type of idealism is by adopting the specific form of it that I have labelled *canonical* idealism, where the central world-creative role is assigned to the sensory organization. This organization is the system of control over the course of human sensory experience that disposes it to systematically conform to its world-suggestive pattern, and, in the context of certain endowments of the human mind—endowments that ensure the mind's empirical receptivity to the orderly character of its sensory experiences—this

conformity creates the systematic appearance of a certain kind of world at the human empirical viewpoint. The claim of canonical idealism is then that, either on its own or as part of a larger complex, the sensory organization, in combination with the relevant endowments of the mind, constitutively creates the physical world, and determines the character of its primary core, precisely by disposing things to empirically appear systematically worldwide in this way. This primary core of the world—the *core physical reality*—is that portion of the total physical reality which is covered by physical facts that are either physically fundamental (not wholly or partly constituted by other physical facts) or wholly constituted by facts that are physically fundamental.

It is clear how canonical idealism ensures the requisite empirical immanence of the physical world. Sensory experience is able to put us into awareness contact with physical items because, in the context of whatever other factors may be constitutively relevant, it is the organization of such experience that logically creates them. And, *ex hypothesi*, the way things are disposed to present themselves empirically at the human viewpoint is accorded the requisite authority with respect to the issue of how things physically are.

If the only way of giving the physical world the requisite empirical immanence is by embracing phenomenalist idealism, and if the only ultimately satisfactory way of developing this position is by adopting canonical idealism, it might seem that the issue of the physical world is now settled. But the situation is not so simple. What creates the complication is that, while canonical idealism is needed to secure the empirical immanence of the world, and thereby avoid the defects of realism, the idealist's position itself is open to certain seemingly powerful objections. The most obvious objection, and the one on which I shall mainly focus, stems from the fact that the idealist account seems to be in clear conflict with our basic understanding of the ontological status of the world in relation to human mentality—an understanding that we cannot give up without undermining the belief that the item whose status is at issue qualifies as a physical world at all. The conflict seems so blatant that the whole idea of canonical idealism, and phenomenalist idealism in general, is likely to strike us as absurd. How could we begin to think of something as a genuine physical world—a world of real three-dimensional space and material objects—if it logically depends for its very existence on facts about human sensory experience? I have tried to show, by philosophical argument, that the realist view of the world is untenable. But, on the face of it, the idealist alternative is just a non-starter.

There are two factors that individually give rise to this impression of a conflict. The first is that, when we consider the nature of the physical world in its own terms, and the circumscribed place that human subjects seem to occupy within it, it is hard to see how we can make sense of the suggestion that facts about human sensory experience centrally contribute to its existence, or that its existence logically depends, in any way at all, on facts about human mentality. The second factor is that it is hard to think of something that is idealistically created in the way envisaged as having the right kind of objectivity, in relation to the human mind, to count as a real world. For given the central role in the idealistic creation that is assigned to the sensory organization, and given that the organization plays its creative role by disposing things to appear systematically worldwise at the human empirical viewpoint, it is hard to see how anything could be created other than some kind of virtual reality. On the surface these two factors seem closely related: both are concerned with the ontological relationship between the physical world and human mentality; both turn on the *prima facie* impossibility of reconciling our basic understanding of this relationship with the kind of ontological dependence on human mentality that the world is given under canonical idealism. But, as we shall see, the issues they raise are, in fact, quite different, and if the idealist is to have any prospect of defending his position, he will need to approach them in quite different ways.

What resources are available for defending canonical idealism on either count will occupy our attention for much of the remaining discussion. But I want to begin by taking a closer look at the situation that would arise if the difficulties for the idealist position turned out to be insuperable and the position had to be rejected.

II

As I have formulated them, both physical realism and canonical idealism are positions advanced in the framework of the assumption that there is a physical world; where they differ is in their accounts of the ontological status of this world, in itself and in its relation to human mentality. But the assumption of the existence of a physical world is not philosophically uncontroversial. The most familiar way in which it has been challenged is by the arguments of the radical sceptic, who, in the traditional debate, points out that we have no way of standing outside the perspective of our empirical viewpoint to check

on the extent to which, if at all, our sensory experiences accurately represent the external situation. This sceptical challenge does not positively deny the existence of a physical world, but simply disputes our entitlement to believe in its existence. Nor does it downgrade the ontological status of the world along the lines of the idealist. For it accepts that, if a physical world exists, it is to be conceived of in a realist fashion, as having an existence which is both logically independent of human mentality and philosophically fundamental. Indeed, this realist conception of the world, or at least the realist's claim of independence, is an explicit part of the sceptic's traditional case, since it is because he assumes that a physical world would have to be something external to the human mind that he takes our belief in it to be problematic. Correspondingly, positions along the lines of phenomenalistic idealism have often been seen as offering a line of defence against the sceptic, though, as we shall see, the idealist's resources in this respect are usually misunderstood.

It is not, however, the possibility of scepticism, but the threat of something much stronger, that gives the issue of the existence of the physical world its particular relevance to our present discussion. We have seen that physical realism has to be abandoned and that the only alternative account of the world that remains a serious option is that of canonical idealism. But suppose it now turns out that canonical idealism itself is vulnerable to objection. Suppose, to take the case on which we are focusing, we are forced to conclude that the idealist account must be rejected because of its conflict with our basic understanding of the ontological status of the world in relation to human mentality. This would mean that we could no longer find any acceptable account of the physical world at all, since any attempt to accommodate all that our basic understanding of the world requires would involve us in a contradiction. Thus, responding to the need for the world to be empirically immanent, we would be forced to say that, if there is a physical world, it has to be conceived of in an idealist way, while, responding to the need for the world to have the requisite mind independence and objectivity, we would be forced to say that, if there is a physical world, it has to be conceived of in a non-idealist way. Since the correct account of the world cannot be both idealist and non-idealist, the conclusion which we would be forced to embrace, taking everything into account, would be that of total physical nihilism, which rejects the existence, and even the possibility, of a physical world altogether. We would still be free to envisage a relevant external reality with a world-like structure and organization, and such a reality, of course, would satisfy all that is required of a physical world in respect of independence

and objectivity. And we could also, perhaps, make a place in our ontological scheme of things for an item along the lines of the idealistically created world, which would embody all that is required of a physical world in respect of empirical immanence. But we would be prevented from recognizing the existence of a physical world itself, since we would be prevented from thinking that there could be anything that satisfies the requirements of such a world on both fronts.

From a common-sense standpoint this would obviously be a counter-intuitive and unpalatable outcome: it is doubtful whether, in our ordinary thinking, we are even psychologically capable of relinquishing our belief in a physical world. But, given that physical realism has been refuted, and that the only remaining alternative to physical nihilism is canonical idealism, it might be thought that, from a philosophical standpoint, being forced to a nihilist conclusion would not be as unsettling as we might initially expect—that it would not affect, as greatly as it ostensibly appears to, our overall view of how things stand. There are two things that could help to create this impression.

In the first place, simply by its rejection of the realist's fundamentalist claim, there is already an element of nihilism in canonical idealism. For, by taking the world to be something whose existence is ultimately constituted by facts of a non-physical kind (facts that are devoid of physical entities and physical properties), the idealist is committed to saying that, when we get down to the level of what is metaphysically fundamental, the physical reality, with all its factual and ontological ingredients, disappears. And since it is part of the definition of the relation of constitution that where a fact *F* is constituted by a fact *F'*, or set of facts *S*, the obtaining of *F* involves nothing over and above the obtaining of *F'* (the members of *S*), we can see why this loss of the physical at the level of what is metaphysically fundamental is, in a certain respect, nihilistic. It is, of course, only *in a respect* nihilistic: there remains a theoretically important distinction between eliminating something at the level of what is metaphysically fundamental and eliminating it altogether.

Second, even when we acknowledge the importance of this distinction, there are a number of areas where it seems that nothing in our view of the situation would turn on whether we ended up as idealists or as nihilists. In both cases it seems we would find ourselves accepting the same account of how things stand in the realm of human sensory experience, and, in particular, of how our experiences are orderly in a way that (in the context of the relevant endowments of the human mind) makes it empirically appear to us that we are mobile percipients in a three-dimensional world. In both cases

we would presumably want to accord to this world-suggestive orderliness a major role in the formation of our ordinary and scientific physical beliefs, and, even if we ended up, as philosophers, denying the existence of a physical world, we would presumably be prepared to acknowledge the utility of these beliefs in capturing the empirical significance of the orderliness and equipping us to lead interesting and constructive lives. In both cases, too, we would presumably accept that the orderliness stems from, and reflects, some unitary system of control—the sensory organization—that disposes experience to be orderly in the relevant way, and in both cases we would be likely to suppose that the source of this control—the organizing agency—was something external to, and ontologically independent of, the human mind. Moreover, in both cases there would be the option, already mentioned, of taking the external source of control to be a form of reality with the same structure and organization as the realist's physical world—a reality that constrains the course of experience to conform to a world-suggestive pattern by, as it were, imposing its own character on it. Set within the context of so much common ground, and potential for common ground, the area of disagreement between the idealist and nihilist positions could come to seem quite narrow.

It might be thought, then, that, while idealism and nihilism are importantly different, the issue between them is not as crucial as it ostensibly appears to be, and that there is not a huge amount that we would stand to lose if we found ourselves unable to defend canonical idealism against objections. But this would be to overlook an important point. It is true that, once the nihilist has accepted the putative facts about human sensory experience that centrally feature in the idealist's account, his overall philosophical position need not greatly differ from that of the idealist. The only unavoidable difference would be that, while the idealist sees the organization of sensory experience, in the context of the other relevant factors, as sufficing to make it true that there is a physical world, the nihilist would see it as merely sufficing to make belief in a physical world a useful fiction. But where, in addition, the situations of the idealist and the nihilist crucially differ is over their entitlement to claim that the relevant experiential facts obtain. For, by denying the existence of the physical world, the nihilist deprives himself of any decent grounds for supposing that our sensory experiences are collectively orderly in the relevant world-suggestive way, or that they are controlled by the relevant form of order-imposing organization.

The basic point is that we are only entitled to think of ourselves as in a position to gain significant information about the course of human sensory

experience, and so equipped to discern its world-suggestive orderliness, if we are already entitled to accept the existence of the physical world and to rely on our ordinary beliefs about it. I am not now thinking of the point I made in Chapter 3, that the relevant orderliness of sensory experience consists in its systematic amenability to physical interpretation, so that it is only by reference to this interpretation that we can even specify what it is about our experiences that makes them orderly. This point, in itself, would not mean that we could not deny the existence of a physical world, but still be entitled to think of ourselves as equipped to discern the experiential facts that allow a physical interpretation to be empirically successful. The point I am now making is that, if we are to be justified in thinking that we have significant access to these experiential facts, and so can discern the orderliness that invites physical interpretation, we have to be already and independently justified in taking our access to the physical world for granted; and, of course, a nihilist denies that there is a physical world for us to have access to at all.

The point has two aspects. The first and most obvious aspect concerns our knowledge, or putative knowledge, of the sensory experiences of others. Assuming we possess it, such knowledge, and knowledge of the mental states of others quite generally, is not conveyed to us by direct mind-reading, but by our appropriately interpreting the behaviour, utterances, and physical circumstances of the subjects involved. It is not that, in normal cases, we find ourselves consciously inferring the other-minds conclusions from premisses about the relevant physical factors. Typically, we find ourselves directly acquiring the relevant other-minds beliefs when the appropriate physical factors obtain and are perceptually manifest. This is clearly so, for example, in cases where we are in conversation with someone else, who verbally communicates to us his thoughts and experiences. We do not infer what the other person is thinking or experiencing by focusing on the physical character of his utterances, working out what they mean, and assuming that he is trying to give us correct information. We just understand the utterances straight off, and are thereby led to acquire the appropriate beliefs about his mental condition. But even when the relevant information is thus directly received, it is only available for such reception because of the presence of the physical world and the perceptible physical factors to which we cognitively respond. If, at the level of philosophical reflection, we were to drop our belief in the physical world, we could no longer think of ourselves as in a position to gain other-minds information in this kind of way, since we could no longer think of ourselves as having this kind of communicative contact with other subjects.

The second aspect of the point is less obvious. Leaving aside other-minds knowledge, the only kind of information that someone might be equipped to acquire about the course of human sensory experience would be exclusively concerned with himself. And in this area we might at first suppose that the nihilist will fare better. For in the case of our own minds we have direct introspective access to our current mental condition—an access that does not depend on our responding to physical evidence or physical cues—and so the nihilist might hope that, by introspectively acquiring and retaining knowledge of our sensory experiences, each of us will at least be in a position to discern the relevant form of world-suggestive orderliness, and infer the presence of an appropriate sensory organization, with respect to his own case. But, in fact, even this modest hope would be in vain. For, whatever may be possible in principle, in practice the bulk of what a human subject knows, or can discover, about his own experiential life is inextricably tied to the framework of his physical beliefs. Thus, to take my own case, which I assume to be typical, I certainly believe, and would ordinarily take myself to know, that my past experiences have, by and large, exhibited an orderliness suggestive of the presence of a three-dimensional world in which I live and move as a perceptive inhabitant. And I believe that it is only because my experiences have exhibited this orderliness, and because it has shaped the formation of my beliefs, that I have come to accept the existence of such a world, and to think of its character in the terms that I do. But the reason why I now believe these things is not that I possess, or possess anything remotely approaching, a comprehensive recollective record of my past experiential biography—a record that enables me to discern the orderliness by directly surveying the composition of this biography and establishing its amenability to physical interpretation. It is rather that, not having much in the way of a recollective record, I nonetheless find myself with a rich stock of putative information about the physical world and my place as a predominantly veridical percipient within it, and, on this basis, I assume that my experiences, being predominantly veridically perceptive, have exhibited the world-suggestive orderliness that this would require. The knowledge of my experiential past that I now possess independently of what I take to be my present physical knowledge is relatively meagre and, divorced from the framework of that knowledge, evidentially insignificant. Even where I think that I can directly recall some past experience—for example, directly recall how things visually or tactually appeared to me on a particular occasion—I rely on my knowledge, or putative knowledge, of the physical circumstances

of its occurrence in order to gauge its temporal relations to other episodes in my experiential past, and its temporal distance from the present.

It turns out, then, that, while the nihilist may initially assume that he can take over the idealist's account of how things stand in the realm of human sensory experience—that he can legitimately recognize the same world-suggestive orderliness and infer the same form of sensory organization—this assumption is misconceived. For almost all the knowledge that we think we have about the experiential realm either directly rests on knowledge we think we have about the physical world, or depends for its epistemic credentials on the assumption of the existence of a physical world and our epistemic access to it. The upshot is that, by denying the existence of a physical world, the nihilist leaves us, at the level of philosophical reflection, in something close to an epistemic void. He leaves us without grounds for supposing that we can discover anything about the mental lives of other human subjects, or, indeed, for supposing that such subjects exist at all. He leaves us without grounds for thinking that we can find out more than an insignificant fraction about our own past experiential lives. And, as well as requiring us to reject all our ordinary physical beliefs, he leaves us without grounds for supposing that, by reflecting the world-suggestive ways in which our experiences are ordered and organized, they even serve a useful purpose. As for what might exist beyond the realm of human mentality, by so extensively undermining the credentials of what we ordinarily take ourselves to know about the course of sensory experience, and the other aspects of human mentality, the nihilist leaves us without any basis for even rational conjecture. All in all, if we were to embrace physical nihilism, our epistemological resources would be hardly any better than those of a solipsist of the present moment.

III

There is no denying that the epistemological outlook under physical nihilism is bleak. It might seem strange, however, that I am stressing this point in the context of a comparison between nihilism and *idealism*. If the comparison were between nihilism and *realism*, then the point would be well taken. But surely, it will be said, the idealist, too, will find himself in exactly the same epistemological difficulties. For since he is excluding the physical world from the metaphysically fundamental reality, surely he, too, like the nihilist, will have to start from epistemic scratch, without the benefit of physical

knowledge behind him. Surely he, too, will have to start from information about human sensory experience and work from there to try to establish the presence of the factors that he sees as sufficing for the idealistic creation of a physical world. And if this is his situation, then he will encounter exactly the same problem—the problem that, without being entitled to rely on his ordinary physical beliefs, he will have no grounds for supposing that sensory experience exemplifies the relevant form of orderliness and is controlled by the relevant form of organization, and so no grounds for supposing that the factors needed for the idealistic creation obtain.

Certainly, in traditional discussion phenomenalist idealism and other positions along similar reductive lines are often represented as, in part, epistemological theories, which are offered in response to radical scepticism about the physical world, and attempt to put our belief in a world on a firmer foundation. On the one side, the radical sceptic is represented as insisting that a belief in a physical world—a world realistically conceived—is mere speculation, since we cannot get behind our sensory data to check on their causal origins and representational accuracy. On the other side, and in the context of the sceptical challenge, phenomenalist idealism, at least when developed in its purest form, is viewed as a method of epistemological defence. For if the idealist takes the physical world to be wholly created by facts about human sensory experience—and this would be the purest form of his position—he seems to bring physical facts within the scope of what our experiences are equipped to reveal. The world becomes, as it were, relocated on *this side* of the veil of perception, where we seem to be ideally placed to discern its character. We now know that, when it is cast in this anti-sceptical role, phenomenalist idealism, even in the pure form just envisaged, would be a failure: it would fail for the same reasons as created the epistemological problem for the nihilist. For if it were forced to start by suspending our ordinary physical beliefs, and had to rely, for any physical conclusions it reached, on what we can independently know about sensory experience, there would be hardly anything to serve as its epistemological premisses, and its experientially reductive account of the physical would give it no advantages over realism. There is no getting round the fact that the bulk of what we know, or think we know, about the realm of human sensory experience depends on our taking our knowledge of the physical world to be already in place.

In fact, however, the idealist does not need to approach the epistemological issue of the physical world in this sort of way. In excluding the physical world

from the domain of what is metaphysically fundamental, and in taking it to be ultimately constituted in the way envisaged, he is not making a claim about the foundations of knowledge. He is saying that physical facts are ultimately sustained by, and their obtaining involves nothing over and above, facts of a different kind, and he is taking these sustaining facts to be wholly or partly to do with how things stand in the realm of human sensory experience. But he is not committed to saying that our epistemic access to physical facts must be mediated by our access to the sustaining facts. Nor is he committed to saying that it is only after we have independently acquired knowledge of the sustaining facts that we are entitled to credit ourselves with access to the physical. As far as his idealism goes, he is free to insist that our access to the physical facts is epistemologically basic, and that we can legitimately take it to be so without having first to forgo any reliance on the knowledge it provides. And this would allow him to say that the physical information that is thus independently available can be employed to help in the establishing of the experiential facts that feature in his idealist account. In effect, he is free to combine his metaphysically reductive construal of the physical world with the epistemological aspects of our common-sense outlook, in which, as well as conceiving of the world in a realist fashion, we assume that physical states of affairs are unproblematically open to empirical scrutiny through ordinary observational means.

Indeed, the idealist is in a better position than the realist himself to take his stand on this common-sense epistemological view. For realism, as we have seen, is not able to preserve our common-sense assumption that we have perceptual access to the physical world, and without such access it is not possible for physical states of affairs to be open to observational scrutiny in the ordinary sense. The idealist's great advantage, epistemologically, is not that he is better equipped to take up the sceptical challenge and show how, through our independent access to the non-physical facts that constitutively underlie the world, our physical beliefs can be put on a firmer foundation than the one on which we normally rely. It is, rather, that, by avoiding the realist's problem of perception, he has the opportunity of dismissing the demand for a firmer foundation as misconceived. He has the opportunity of taking his philosophical stand on the epistemological position that we accept in our ordinary thinking—of insisting that we are equipped to gain direct knowledge of our physical environments through perception, and that this is something which, despite the sceptic's challenge, we are entitled to take for granted. This, in my view, is the right epistemological stance for the idealist to

adopt, just as it would be the right stance for the realist, too, if only he could make sense of physical-item perception. (Of course, even the idealist will not be in a position to adopt it if he is unable to defeat the challenge of the nihilist.)

This raises the question of how, in the idealist system, we can think of environmental states of affairs as able to reveal themselves to a subject perceptually, given that they are not part of what is metaphysically fundamental, and so do not have any fundamental influence on the human mind. In detail, the answer to this would be complex, and, at certain points, there is more than one approach that the idealist might take, but, in basic outline, the situation is, I think, reasonably clear.

We need to begin by reminding ourselves of a point that I stressed in an earlier context.¹ This is that sensory experiences of a perceptual kind—and it is only such experiences that are qualitatively equipped to be physically perceptive—have a conceptual content that makes them phenomenologically seem to their subjects to be presentational perceptions, in perspective, of items in a three-dimensional spatial environment. Thus, a visual perceptual experience, by its conceptual content, phenomenologically seems to its subject to be the presentational awareness of an environmental arrangement of colours, viewed in the perspective of his own spatial position, and a tactual experience, by its conceptual content, phenomenologically seems to its subject to be the presentational awareness of a certain form of spatial contact between a part of his body and another environmental object. As for the nature of the relevant conceptual content, I have already explained that my own view of the nature of perceptual experience is along the lines of the traditional sense-datum theory, which takes the core of any such experience to consist in the presentation of a sensible item (a sense datum or sense quale) that is internal to the realm of sensory awareness; and, in the context of that basic approach, I take the relevant conceptual content to be an element of interpretation that is directed on to the presented item and represents it as something environmental. But this specific account of the content is not something on which I need to insist in the present context.

Because perceptual experiences have this conceptual content (however understood), and thereby have this phenomenological feel, they purport to inform the subject—invite him to believe—that his current environment is sensibly characterized in the relevant way. So, if his visual experience phenomenologically seems to the subject to be a presentational awareness of a

¹ In Ch. 3, Sect. VI.

certain array of colours at certain distances and in certain directions from him, then it will thereby purport to inform him—invite him to believe—that this is how, in relation to him, things in his current environment are. The fact that an experience issues a doxastic invitation to the subject does not, of course, ensure that the invitation will be accepted and the relevant environmental belief formed. It does not even ensure that the subject will feel any inclination to accept the invitation and form the relevant belief. Even in the context of everyday life, where we tend to accept sensible appearance at face value, there are various factors that can persuade a subject that his experience is partly or wholly non-veridical, and that the putative information it carries needs to be accepted with some qualification or rejected entirely.

For a subject to gain epistemic access to an environmental state of affairs through perception, his perceptual experience must be linked with the obtaining of that state of affairs in a way that allows it to qualify as a perceptual registering of it. In the framework of realism, we have no way of understanding how such perceptual registering would be possible, since we have no way of understanding how perceptual awareness can reach to the world realistically construed. In the idealist's system we can take such registering to be secured by the combination of three factors. The first factor is that the experience purports to inform the subject of the current obtaining, in perspectival relation to him, of a state of affairs of the relevant type. So, if the relevant state of affairs is that of there being a certain arrangement of colours at various distances and in various directions from him, the experience is, by its content, such as to purport to inform the subject that his current environment is thus arrayed. The second factor is that the experience occurs as a result of those aspects of the sensory organization, or of the concrete factors underlying them, that are distinctively involved in the constitutive sustainment of the relevant state of affairs. These aspects of the sensory organization are the ones that are distinctively responsible, in the context of the relevant endowments of the human mind, for disposing things to give the systematic appearance, at the human empirical viewpoint, of the obtaining of such a state of affairs at the relevant time. This factor creates an ontological link between the experience and the relevant state of affairs, and it is this link that allows the state of affairs to come within the scope of the subject's awareness. The third factor is that the way in which the occurrence of the experience results from the relevant aspects of the sensory organization (or from the concrete factors underlying them) exemplifies a type of experience-generative process by which, quite generally, perceptual experiences tend to occur as a result of

aspects of the sensory organization (or the factors underlying them) that are distinctively involved in the constitutive sustainment of environmental states of affairs of types with which the experiences are content-linked—those types of whose obtaining the experiences purport to inform the subject. This factor ensures that, although the state of affairs is not ultimately involved in the production of the experience, the informational content of the experience can be thought of not merely as qualitatively fitting, but as concretely reflecting the character of the state of affairs. Combined with the second factor, it is this that allows the experience to count as a perceptual registering of the state of affairs. The precise way in which the second and third factors work out will depend, of course, on the concrete form of the sensory organization—on what it is that the idealist takes to be ultimately responsible for controlling the course of human sensory experience and disposing it to conform to its world-suggestive pattern. We have already touched on this topic, but it is one that we shall need to explore in more detail later.

There is still the question of what entitles the idealist to take his stand on the epistemological position of common sense and dismiss the sceptical challenge in the way I am advocating. We can, as we have just seen, provide an account of how, in the idealist's system, environmental states of affairs can become open to epistemic scrutiny through perception. But if the idealist cannot establish, or find grounds for believing, that human sensory experience exemplifies a world-suggestive orderliness except by relying on a prior knowledge of the existence and character of the physical world, what gives him the right to be confident that our perceptual experiences do standardly provide genuine environmental information or that there is a physical world at all? Why is it unreasonable for the sceptic to insist that, if we are to be justified in retaining our ordinary belief in the existence and epistemic accessibility of the world, this belief must be put on a firmer foundation? For the time being I shall put this further issue on one side, and simply assume that the common-sense position is the correct one to adopt. I shall return to the issue at the end of the next chapter, after the idealist position has been fully developed and fully vindicated in other respects.

IV

There is much at stake, then, epistemologically, over whether, having already rejected realism, we end up endorsing canonical idealism or nihilism with

respect to the physical world. The idealist, as well as preserving our knowledge of the physical world, can use that knowledge as a basis for reaching well-grounded conclusions about how things stand in the realm of human sensory experience. So, although he takes certain kinds of organizational fact about sensory experience to be the central component in what constitutively creates the physical world and endows it with its core character, he does not have to find some independent way of establishing the obtaining of those facts—prior to recognizing the existence of the world—but can, rather, appeal to what is independently known about the physical world as a basis for concluding that the experiential realm is organized in the constitutively appropriate way. In contrast, by rejecting the existence of a physical world, the nihilist would condemn us to something approaching an epistemic void. For without the basis of what we ordinarily take ourselves to know about the character of the world, there is very little that we would be entitled to conclude about the character of the sense-experiential realm, or, indeed, the realm of human mentality quite generally. And without much in the way of well-grounded beliefs about the realm of human mentality, there would be no basis for even rational speculation about what might lie beyond.

If there is much at stake over whether we accept canonical idealism or physical nihilism, there is much at stake over whether we can find a satisfactory way of defending the idealist position against the seemingly decisive objection already identified, that the idealist account is excluded by our basic understanding of the ontological status of the physical world in relation to human mentality. As we have seen, this objection divides, in effect, into two objections, which focus on different ways in which the idealist account and our basic understanding seem to come into conflict. Thus, on the one hand, there is the objection that, when we consider the nature of the world in its own terms, and the circumscribed place that human subjects seem to occupy within it, it is hard to see how we can make any sense of the suggestion that facts about human sensory experience centrally contribute to its existence, or that its existence logically depends on facts about human mentality in any way at all. On the other hand, there is the objection that something that was idealistically created in the way envisaged—with the central creative role assigned to the sensory organization—would not have the requisite objectivity, in relation to the human mind, to count as a real world. These two objections will need to be dealt with separately. I shall consider the first one now, and leave consideration of the second until the next chapter.

Whatever the correct philosophical account of it, it is surely undeniable that the physical world, if it exists, is a world of three-dimensional space and mobile material occupants. Certainly, this is something that we have been taking for granted in all the discussion so far, and it is surely clear that something that was not of this spatio-material type would simply not qualify as the physical world in any recognizable sense.² Now within the spatio-material world there is a certain group of complex material objects with which human subjects are intimately associated, and, indeed, it is by virtue of this association that human subjects qualify for the title 'human', as that title is commonly understood. The objects in question are biological organisms belonging to the animal species *homo sapiens*; they are what I shall speak of as *human organisms*, and they incorporate all and only the material aspects of what we mean by *human beings*. In normal circumstances—and, for simplicity, I shall ignore any other—the association of human subjects with human organisms is one-to-one: each subject (subject of mentality) is relevantly associated with just one human organism, and that organism is relevantly associated with only that subject. The nature of the association is philosophically controversial. Many philosophers take the association to be one of identity: human subjects simply are human organisms and are material through and through. Others hold that, while human subjects are not purely material, they are partly so, and have the properties of a human organism as the material part of their nature. Others, again, think that human subjects, or at least those subjects to which human mentality fundamentally belongs, are wholly non-physical entities, without material composition or location in physical space. On this Cartesian view the association of subject with organism is normally assumed to be, at least primarily, a functional one—a matter of there being some psychophysical arrangement which equips the organism to causally affect the mental states and activities of the subject and equips the subject to causally affect the physical states and activities of the organism. Certainly, this is how Descartes himself saw things.³ Of course, even on the physicalistic and semi-physicalistic views of the mental subject,

² The insistence that the physical world contains material objects does not exclude supposing, along the lines envisaged earlier (Ch. 3, Sect. II), that such objects are ontologically derivative entities, whose existence is ultimately constituted by the spatio-temporal distribution and functional organization of certain region-characterizing properties, though it would only be in the framework of a realist view of the world that such a supposition could be entertained.

³ See especially his Sixth Meditation and *The Passions of the Soul*, in *The Philosophical Writings of Descartes*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch, i and ii (Cambridge: Cambridge University Press, 1984).

the philosophers concerned will recognize causal interaction between mind and organism. But the psychophysical provision for modes of interaction will not play such a prominent role in defining the nature of the subject–organism association, and in the full-blooded physicalistic case the interaction will be entirely within the physical domain.

The problem for the idealist is that there seems to be no plausible way of understanding the association between human subjects and human organisms that is compatible with his own philosophical view of the world. Clearly, in order to be able to assign the role he does to the sensory organization in the constitutive creation of the world, he has to conceive of the relevant subjects of sensory experience in a Cartesian fashion, as wholly non-physical and ontologically independent of all things physical. So, he cannot construe subjects as organisms or as entities which have the properties of an organism as part of their nature. But even from the standpoint of the Cartesian view (and I shall be looking more closely into the idealist's commitment to this view in the next chapter) there does not seem to be a way for him to make sense of the psychophysical situation as it would be standardly interpreted. The association between the non-physical subject and the physical organism is, as I have said, normally taken to be primarily a functional one, and, ignoring any issue about the nature and status of the physical world, there is no denying that this presents itself, initially, as the most plausible view. Certainly, there is a range of familiar ways in which a functional association seems to manifest itself through detectable forms of psychophysical causation, such as the way in which input to the organism's sensory receptors tends to induce sensory experiences in the subject, and the way in which attempts at physical action by the subject tend to elicit motor responses in the organism. And the neurophysiological evidence seems to point to the presence of a complex interplay between mind and brain in all areas of mentality. But once it is accepted that human subjects have this functional relationship with human organisms, it seems impossible to think of the physical world as ontologically dependent on the human mind, as the idealist claims. It seems impossible to avoid thinking of the realm of human mentality as just a circumscribed portion of a larger psychophysical reality, whose physical portion has the logical capacity to exist without it, or at least (to echo a point explained earlier⁴) has this capacity apart from any respects in which causation from the mental realm to the physical may logically contribute to

⁴ In Ch. 3, Sect. V.

the identities of certain *physical particulars*. The special weight that the idealist puts on the role of the sensory organization in the constitutive creation of the world makes the problem particularly acute. Within the context of the envisaged functional relationship, it is the physical world itself that is responsible for this organization, since the character of the subject's sensory experiences is controlled by the organism's sensory input, and the character of this input reflects the character of the organism's environment. And it is surely a manifest incoherence to suppose that the thing that is responsible for organizing our sensory experiences derives its very existence, in part, from the obtaining of that organization.

Given the apparently clear-cut conflict between the idealist's account of the physical world and the supposition of a functional relationship between human subjects and human organisms, it seems that the only option left for the idealist would be to deny that such a relationship obtains. And such a denial, though counter-intuitive, would not be utterly absurd. What creates the *prima facie* impression of a functional relationship, and what leads us to suppose that there is one when we consider things from a philosophically neutral standpoint, is that there is a certain one-to-one pairing of subjects with organisms such that, within the domain of normally endowed subjects and organisms, there are certain conspicuous forms of reasonably reliable correlation between subject events and organism events within the same pair. So, for example, given a normally endowed subject–organism pair, when the organism ocularly encounters a certain kind of environmental scene, the relevant subject can normally be relied on to have, a moment later, a representationally appropriate form of visual experience, and when the subject mentally makes to perform a certain kind of bodily action, the organism can normally be relied on to exhibit the appropriate form of motor response. It is hardly surprising that, with no philosophical axe to grind, we come to interpret these correlations as regularities of causal interaction—as a manifestation of a stable psychophysical arrangement that empowers the organism to causally affect the subject and the subject to causally affect the organism in the relevant ways. And it is then only natural that we should extend this functional understanding of the relationship between the subject and the organism to cover the subject's mentality on a broader front. But the fact of these correlations between what takes place in the subject and what takes place in the organism does not *prove* that there is causal traffic between them. Even within the framework of realism, some of those who have accepted a Cartesian view of the mind, but have thought that there are difficulties

over understanding how the physical organism and the non-physical subject can come into causal contact, have been prepared to deny that such contact exists. In my view, the supposed difficulties for Cartesian interaction are illusory, as I have tried to show elsewhere.⁵ At the same time, I think that the argument I have developed here does succeed in establishing that, if there is a physical world, its existence needs to be understood in the envisaged idealist way. So, if there is a problem for the idealist in making sense of a functional relationship between subjects and organisms—a relationship which makes provision for forms of causal interaction—perhaps he can afford to appeal to that argument and deny that such a relationship holds. Perhaps he can afford to dig his heels in on his idealist account of the world and draw the conclusion to which that account apparently commits him.

If there were no better line of defence available to the idealist, I suppose I would have to accept this proposal, but I find it unattractive. We are assuming that, whatever the relationship between human subjects and human organisms, there is a rich provision of causal interaction within the physical world itself. And, indeed, I think that to deny the presence of causal processes within the world would be tantamount to denying that there was a physical world at all in any recognizable sense. On the idealist view these forms of physical causation, like everything else that physically obtains, are constitutively sustained by aspects of the sensory organization, in the context of certain other factors, and the constitutive role of these aspects is to dispose sensory experience to provide the systematic appearance of such causation at the human empirical viewpoint. But, for the most part, this systematic empirical appearance of causation within the physical realm takes the form of the systematic appearance of certain forms of correlation within that realm. For example, the empirical appearance of a world whose material occupants interact gravitationally (responding to a causal force of mutual attraction) consists in the empirical appearance of a world in which certain aspects of the behaviour of these occupants, in relation to one another, correlate in a certain way with the quantity of their matter and the distances between them. This commits the idealist to saying that, in the context of whatever other factors are relevant, the sensory organization constitutively sustains the relevant forms of physical causation by disposing sensory experience to provide the systematic empirical appearance of the relevant forms of physical correlation. But if the sensory organization plays this constitutive role in the case of

⁵ John Foster, *The Immaterial Self* (London: Routledge, 1991), ch. 6.

physical—intra-world—causation, there is an awkwardness in denying it an analogous role in the psychophysical case, where it also disposes things to present the systematic empirical appearance of a world whose character, in combination with the distribution of types of mentality over subjects and times, secures causation-suggesting forms of correlation. It is true that the idealist has a reason for refusing to accept the analogy: the recognition of the world as having an internal functional organization does not conflict with the idealist's understanding of how the existence and character of the world are ultimately constituted, while, on the face of it, the recognition of a functional relationship between subjects and organisms does. But there is still something awkward about having to appeal to a mode of idealistic sustainment that works in the one case and not in the other.

Fortunately, I think there is a much better way in which the idealist can deal with the problem that has arisen. This problem stems from an apparent conflict between the way in which his idealism represents the relationship between the physical world and human mentality and the way in which it seems that we have to understand this relationship when we think about the world in its own spatio-material terms, and focus on the special association between human subjects and human organisms. Thus, on the idealist account, the physical world is ontologically dependent on the human mind, since organizational facts about human sensory experience centrally feature in the factors that constitutively create it. But when we focus on the apparently functional relationship between the human subject and the human organism, it seems we have to think of the world, or at least the core physical reality, as ontologically independent of the human mind (or as independent apart from any respects in which causation from mind to world affects the identities of physical particulars). It seems to me, however, that, advanced on this basis, the claim of ontological independence is not something that the idealist needs to dispute, and that the apparent conflict referred to is illusory. Certainly, the idealist is committed to saying that the physical world is something whose existence is ultimately constituted by facts of a more fundamental kind, and that these constitutive facts centrally include ones about human sensory experience. This is the very core of his idealist position, which sets it in opposition to the two claims of the realist. But this, it seems to me, does not prevent him from acknowledging that, as part of the psychophysical reality whose physical portion is thus ultimately constituted, human subjects feature as entities that are functionally attached to human organisms, and that, *in that context*, they, and the facts about their mentality,

logically contribute nothing to the existence and character of the physical world beyond what even a realist himself would accept. The crucial point is that human subjects and their mental lives feature in two quite different places in the idealist's account, and the way in which the idealist needs to conceive of their relationship to the physical world is quite different in the two cases. In the context of the idealist's account of the ultimate metaphysical status of the physical world, the organization of human sensory experience plays a central role in the constitutive creation of the world, so that, in that context, the realist's thesis of ontological independence must be denied. But, with respect to what gets idealistically created, human subjects and the physical world form interactive partners in an integrated psychophysical whole, and, if we want to do justice to the nature of this partnership, it would quite wrong to characterize the physical component as logically depending for its existence on the psychological. In other words, what we have to recognize is that, from the standpoint of canonical idealism, the issue of the ontological relationship between the physical world and human mentality arises in two quite different forms, according to whether we consider how things stand with respect to the nature of the psychophysical reality whose obtaining is idealistically sustained, or consider how things stand with respect to the nature of this sustainment, and the correct way of characterizing the relationship depends on which form of the issue applies.

Confirmation of this approach comes by considering another example where the idealist appears, at first sight, to be in difficulties over the relationship between human mentality and the physical world. Let us continue to assume, what we have been implicitly taking for granted throughout the previous discussion, that the time dimension which features in the physical world is also part of the metaphysically fundamental reality—the reality which contains human mental subjects and the other entities and facts that are constitutively involved in the idealistic creation. Now suppose that, within this reality, the first human subjects come into existence at a certain time t , along with that system of control over the course of human sensory experience (the course of experience in both those first human subjects and their successors in the human community) which forms the sensory organization. By disposing things to appear systematically worldwise at the human empirical viewpoint, this organization, in combination with whatever other factors are thought of as relevant, will then be what, on the idealist account, serves to logically create the physical world, and the precise content of this disposing will be what determines the character of the world in its

primary core. Let us further suppose that, when scientifically evaluated in the appropriate way, the way things are disposed to empirically appear at the human viewpoint unequivocally supports the hypothesis that the physical universe came into existence at a time t' , millions of years before t . The question is: how, given these suppositions, should the idealist date the start of the universe? It is clear that he is committed to saying different, and ostensibly conflicting, things, according to how he approaches the issue. Looking at the situation of what is idealistically created, he is committed to saying that the physical universe comes into existence at t' —the time that is fixed by how things are disposed to empirically appear at the human empirical viewpoint—even though t' is much earlier than the time, t , when human subjects arrive on the scene. In other words, he is committed to saying that what is created as a world for human subjects—created in part by the way their sensory experience is organized—is created as something whose existence precedes them. On the other hand, looking at the situation in terms of the idealistic creation, he is committed to saying that the universe (complete with its internal history) only starts to exist at t , the time when human subjects and the sensory organization come into existence, since it is this organization, in combination with the other constitutively relevant factors, that logically creates it. So, approaching the issue of the temporal origin of the universe from different standpoints, the idealist finds himself committed to quite different answers. But it is surely clear that this does not expose some real inconsistency in his position: it does not reveal that he is committed to a contradiction. Rather, it shows that the issue itself can be interpreted in two quite different ways, according to whether one is considering how things stand with respect to the created physical reality or how things stand with respect to the mode of its creation, and that what it is correct for the idealist to say varies with the interpretation. With respect to what is created, it is correct for him to say that the universe comes into existence at t' , and with respect to its mode of creation, it is correct for him to say that it comes into existence at t .

This is exactly analogous to the idealist's situation in the case of the issue of ontological independence. With respect to the nature of the reality that is idealistically sustained, the idealist should be happy to accept that the physical world and human mentality are interactive partners in a functional whole, and that, in the context of that partnership, the world is ontologically independent of the mind. Indeed, he should be happy to accept that, in the context of that partnership, the world is what controls the course of human

sensory experience. This is entirely compatible with his insistence that, with respect to the idealistic sustainment, both the existence of the physical world and its functional partnership with human mentality are ultimately logically derived from more fundamental factors, in which the organization of human sensory experience centrally features.

What these cases, in effect, show is that, from the standpoint of the idealist position, in which the physical world is excluded from the domain of what is metaphysically fundamental, there are two quite different frameworks of thought, or assertion, within which claims about the physical world and its relationship to human mentality can be made, and the same claim—same in terms of its conceptual content—can sometimes have different truth-values in the different frameworks. Thus, on the one hand, there is what we might call the *mundane* framework, in which the physical world, and the larger psychophysical reality of which it forms a part, are specified in their own terms, without reference to what ultimately underlies them. All that we ordinarily and scientifically want to say about the physical world and its functional links with human mentality belongs to this framework. On the other hand, there is what we might speak of as the *transcendental* framework, in which the situation of the physical world and the larger psychophysical reality is considered from the standpoint of their ultimate metaphysical status. It is in this framework that the idealist advances his distinctive thesis about the world, which represents it as constitutively created by facts of a more fundamental kind, in which facts about the sensory organization centrally feature. What is crucial is that, in trying to evaluate the idealist position, we should not conflate these frameworks and suppose that, in order for the position to be true, the way in which the idealist transcendently characterizes the situation has to match the way in which things stand mundanely. It was just such a conflation that created the false impression that the idealist account of the constitutive creation of the world was in conflict with what we have to accept when we focus on the nature of the world in its own spatio-material terms and try to do justice to the special association of human subjects with human organisms. This impression disappears as soon as we recognize that how things need to be specified when the world is considered in its own terms does not have to coincide with how things need to be represented in the final metaphysical story. The idealist can accept that, in the mundane framework, it is correct to characterize the world and the human mind as causally interacting, just as it is correct to recognize forms of causal interaction within the world itself, and he can acknowledge that, as

interactive partners, the world and the mind are ontologically independent. But he can continue to insist, as his philosophical account of the ultimate metaphysics of the situation, that facts about human mentality play a crucial role in the constitutive creation of the world.

When he considers the character of the psychophysical reality in its own terms, even the idealist must recognize that its physical component is ontologically independent of its psychological. And, trivially, there is nothing in how the world is to be described in its own terms that involves the reduction of something physical to something wholly non-physical. In this way, we might be tempted to say that even the idealist will have to embrace physical realism at the mundane level. But this would be a mistake. Realism, like idealism, is a philosophical theory about the ultimate ontological status of the physical world, both in itself and in relation to human mentality. Its claims do not pertain to how the world is to be described in its own terms, even though, in a sense, they reproduce the character of the mundane situation in the status they assign. What we *can* say is that, by acknowledging the mind independence of the world at the mundane level, the idealist is able to accommodate one of the factors that helps to give realism its initial appeal, and is thereby able to neutralize one of the factors that helps to make idealism itself seem initially unacceptable.

V

As we noted, there are two ways in which an objector to canonical idealism might take the position to be excluded by our basic understanding of the ontological status of the physical world in relation to human mentality. I have tried to defend the idealist with respect to one form of the objection—the form which appeals to the apparent difficulty of reconciling the idealist account of the constitutive creation of the world with what we need to accept when we focus on the world in its own spatio-material terms, and try to do justice to the special association of human subjects with human organisms. What remains to be considered is the other form of the objection, which claims that something constitutively created in the way the idealist envisages would not have the requisite objectivity to qualify as a real world. I see this objection as posing the biggest challenge to canonical idealism, and indeed to phenomenalist idealism in any form. But, as I have already indicated, I am going to leave discussion of this challenge until the next chapter.

For what remains of this chapter I want to turn to a quite different area where the idealist might be thought to be in difficulties. The issues here are of a more technical kind than those we have just been considering. They are also slightly peripheral to the main themes of the overall discussion. Nonetheless, they do bear on a central aspect of the idealist position, and if this position is to be vindicated, and nihilism avoided, the idealist needs to have an effective way of dealing with them.

The mundane framework is, as I have stipulated, that in which the physical world and the psychophysical reality of which it forms a part are specified in their own terms, without reference to what ultimately underlies them. But, importantly, this does not mean that accepting the idealist's account of how the existence and character of the world are ultimately constituted should make no difference to our understanding of the nature of the mundane situation—our understanding of the sorts of thing that physically and psychophysically obtain to be mundanely specified. And it is already clear, from our earlier discussion, that it should. Thus, under realism we cannot justify the ascription to physical items of distinctively sensible qualities—qualities such as sensible colour and the distinctively sensible forms of extension and shape—since such ascription would not help to explain the ways in which physical items are disposed to sensibly appear, or to explain anything else in the empirical domain. But under canonical idealism where the character of the world is logically drawn from the ways in which things are disposed to appear at the human empirical viewpoint, such qualities can be thought of as achieving physical realization through the dispositions to sensible appearance themselves. Again, under realism both physical space and its material occupants are credited with forms of intrinsic content that are empirically inscrutable—forms whose nature cannot be discovered through empirical investigation, and cannot even be transparently specified in physical terms. But by taking the character of the world to be exclusively drawn from how things are disposed to appear at the human empirical viewpoint, the idealist restricts what physically obtains to what is capable of revealing itself empirically, so that inscrutable forms of physical content are automatically excluded. There is also an important respect in which accepting an idealist account of the constitutive creation of the world should affect our view of how, in the mundane framework, things stand psychophysically. Under realism, as we have seen, we are forced to conclude that physical-item perception is impossible, since there is no way of understanding how our perceptual awareness can reach to objects that are

ontologically independent of the mind. But under canonical idealism the very way in which the world is constitutively created brings it within the scope of our perceptual awareness, and the perceptual accessibility of the world, along with concrete instances of physical-item perceiving, becomes part of what has to be recognized in the mundane framework. It has to be recognized in this framework, even though it is only in the transcendental framework that we can discern the factors that enable us to understand how it obtains.

There is no denying, then, that the idealist's account of what ultimately underlies the physical world has implications with respect to the mundane character of the world and of the larger psychophysical reality of which it forms a part. But there is one aspect of this which seems to create a logical problem—a problem that challenges the very coherence of the idealist position—and it is on this that I now want to focus. The apparent problem stems from the specific way in which physical facts are supposedly idealistically sustained.

Canonical idealism takes the physical world to be something whose existence and character are constitutively sustained by factors of a different kind, in which the sensory organization centrally features. The basic idea is that, in the context of whatever else may be constitutively relevant, the sensory organization constitutively creates the physical world by disposing things to appear systematically worldwise at the human empirical viewpoint, and it constitutively determines the character of the world, in its primary core, by the precise content of this disposing. But if this is what ultimately determines the character of the world, it seems that the character of the world could turn out to be, in certain respects, ultimately *underdetermined*. For it seems that the situation might turn out to be one in which how things are disposed to present themselves empirically would suffice to endow the world with a character, but would not suffice to endow it with a character that was fully determinate. But the suggestion that the world might have a character that is not fully determinate seems to make no sense: it seems to be in conflict with a basic principle of logic. Hence the apparent problem, and the resulting challenge to the coherence of the idealist position.

To illustrate, let us focus on a particular hypothetical case. It seems that, compatibly with what we know about the physical universe and about the content of the sensory organization, we can suppose that there is some physical region and time such that the sum total of the actual and potential empirical evidence available at the human viewpoint fails to provide a decisive indication of whether there is any carbon in that region at that time. This failure could

occur either because the evidence gives significant but inconclusive support to each of the alternatives or because it simply fails to bear on the question at all. Whichever of these factors may be involved, let us take R and T to be such a region and time. Now on the idealist account how things stand with respect to R at T, to the extent that it is ultimately determined at all, is logically determined (in the context of the other relevant factors) by how things are disposed to appear, in that respect, at the human empirical viewpoint. So, given that the actual and potential empirical evidence neither decisively supports the hypothesis that there is some carbon in R at T nor decisively supports the contradictory hypothesis that there is no carbon in R at T, the idealist is committed to saying that there is ultimately nothing that suffices to make either hypothesis true. In other words, he is committed to saying that, in respect of the presence or absence of carbon in R at T, the character of the world is ultimately indeterminate. But, as a matter of simple logic, it seems that this conclusion of indeterminacy cannot be correct. For, in the space of logical possibilities, there is no middle position between the situations of carbon presence and carbon absence. If one of these situations fails to hold, the other, as a matter of logic, must hold. So it seems that the idealist is committed to a logical incoherence.

One way in which the idealist might respond to this challenge would be to say that, contrary to what was suggested, the sort of case just envisaged could not occur. He would not, of course, be able to claim that there is some kind of logical restriction on the way in which our sensory experiences are organized, and on the resulting way in which things are disposed to present themselves at the human empirical viewpoint. He cannot rule out a priori the possibility of a situation in which the sum total of the actual and potential empirical evidence is unequivocal in seeming to indicate the existence of a world, but not unequivocal about every aspect of its character. But what he might still insist is that if we suppose the situation to be of that sort, in which the way things are disposed to present themselves empirically fails to assign a determinate character to the world whose presence it seems to indicate, the conclusion we need to draw is not that the character of the world is ultimately indeterminate, but that there is no real physical world at all. Or, putting the point the other way round, he might say that, in so far as we are happy to accept the existence of a physical world, we must also accept that the way things are disposed to present themselves empirically does succeed in assigning a determinate character to it, simply because canonical idealism is correct and there is no coherent alternative that such idealism allows.

I do not find this response plausible. The difficulty with it is that, if it were correct, then, at the level of what is metaphysically fundamental—the level of the factors that ultimately determine what physically obtains—the difference between the existence and non-existence of a physical world could turn out to be very slight. Thus, let us suppose, as the actual situation, that there is a physical world of the sort we ordinarily believe in, and that the sensory organization is such that, from the standpoint of canonical idealism, the way things are disposed to appear at the human empirical viewpoint fully covers the existence of this sort of world and assigns it a fully determinate character. Now consider a possible situation obtained from the actual situation by changing the content of the sensory organization in one minor respect and holding constant as much of the rest of the situation as that change allows. The change to the sensory organization is such that, if E is the total new way in which things are disposed to appear worldwide at the human empirical viewpoint, and if W is the notional world whose existence E unequivocally seems to indicate, there is one small region of W and one short period such that the actual and potential evidence provided by E does not bear on the question of whether there is any carbon within that region at any time during that period. Given that a physical world exists in the actual situation, it is very hard to think that this slight change to the way things are disposed to empirically appear would serve to eliminate it. Yet it is just such an elimination that an idealist who responds to the problem in the way envisaged would have to accept. He would have to accept that, because this change brings us to a situation in which the idealistically relevant factors no longer succeed in assigning a fully determinate character to a certain minute portion of the putative physical reality, they no longer manage to sustain a physical reality at all. This is not a position I would be happy to accept.

In my view, the correct solution to the problem comes not by denying that the sort of physical indeterminacy envisaged can occur, but by getting clearer about just what the indeterminacy involves. The situation, as we shall see, is complicated. But what we need to begin by noticing is that, even in a case where the relevant sort of indeterminacy occurs—where how things are disposed to appear at the human empirical viewpoint leaves some aspect of the character of the world open—there is still a sense in which canonical idealism represents the physical situation as fully determinate.

Let us continue to focus on the same example. Within the physical world whose existence and character are idealistically sustained there is, we are

supposing, a region R and a time T such that the sum total of actual and potential empirical evidence (supplied by how things are disposed to appear at the human empirical viewpoint) fails to reveal whether there is carbon in R at T. Given the nature of the idealist account, this has the consequence that the factors on which the character of the world idealistically depends neither suffice to make it true nor suffice to make it false that there is carbon in R at T, and this seems to entail that, in the relevant respect, these factors fail to give the world a determinate character. But here we must be careful. Certainly, in one obvious sense, the factors do fail to give the world a determinate character: of the two determinate alternative characters that are available in the relevant respect, they fail to endow the world with either. But there is also a sense in which these factors do, in the relevant respect, assign a determinate character to the world. For, inasmuch as they suffice to sustain the existence of a world which contains the region R at the time T, they automatically suffice to sustain the existence of a world in which either there is carbon in R at T or there is not. After all, even if the way things are disposed to empirically appear does not reveal precisely how things stand with respect to R at T, it also in no way suggests that how things stand defies the laws of logic, and if how things stand conforms to the laws of logic, then one of the physical alternatives—which are jointly logically exhaustive—must hold. (The situation here is not relevantly different from one in which, with respect to two physical alternatives that are not logically exhaustive, the way things are disposed to empirically appear decisively supports the hypothesis that one of them holds without indicating which.) So, in one sense, canonical idealism represents the relevant physical situation as ultimately indeterminate, in that the idealistically relevant factors neither suffice to sustain a fact of carbon presence nor suffice to sustain a fact of carbon absence, while, in another sense, it represents the physical situation as ultimately determinate, in that these factors suffice to sustain the fact that things are, determinately, either one way or the other. This puts a different complexion on the idealist's situation. For on neither count is there any obvious way in which the idealist's position comes into conflict with any principle of logic. There is no logical difficulty in discerning a middle position between there being an idealistic sustainment of a fact of carbon presence and there being an idealistic sustainment of a fact of carbon absence, and it is just that middle position that forms the case of physical indeterminacy. And, crucially, although there is no logical room for a middle position between carbon presence and carbon absence, the idealist

is not committed to saying that there is. Rather, because the idealistically relevant factors suffice to make it true that there either is or is not carbon in R at T, he is committed to saying that that is how things physically are. Quite generally, the idealist can allow for cases in which the idealistically relevant factors leave open certain aspects of the character of the world, and in that sense leave the character of the world as ultimately indeterminate, without having to suppose that, when such cases occur, we are forced to abandon the law of the excluded middle in the mundane specification of the physical situation.

It might seem that we already have, in this, a solution to the whole problem—a solution which shows how the idealist can allow for cases of indeterminacy without falling foul of the principles of logic. But the situation is not quite so simple. Certainly, the point just made solves one aspect of the idealist's problem: it shows how he can accommodate cases of indeterminacy without having to suppose that how things physically stand is in breach of the law of the excluded middle. But there is a further way in which his allowing for cases of indeterminacy could be thought to come into conflict with the principles of logic. For, irrespective of any issue over the law of the excluded middle, there is an argument which purports to show that, in the context of the idealist's account of the constitutive creation of the world, the very notion of the kind of indeterminacy in question is implicitly self-contradictory.

Let us introduce a sentential operator 'It is idealistically true that', or for short 'I-true that', to mean 'The idealistically relevant factors suffice to make it physically the case that'. And, staying with our example, let us also introduce a two-place predicate 'C', meaning '... contains carbon at time ...', so that we can then abbreviate the sentence 'There is carbon in R at T' to 'C(R,T)'. Using 'Not' as short for 'It is not the case that', we can then express the sense in which canonical idealism represents the physical situation as indeterminate by the combination of the claims:

- (1) Not I-true that C(R,T).
- (2) Not I-true that Not C(R,T).

and can likewise express the sense in which it represents the situation as determinate by the claim:

- (3) I-true that either C(R,T) or Not C(R,T).

It is with the implications of claims (1) and (2) that the argument is concerned. It is central to the idealist thesis that the physical world has no

character other than what the idealistically relevant factors give it: whatever is physically the case is so because the idealistically relevant factors make it the case. In particular, then, it seems that the idealist is obliged to say that if there is carbon in region R at time T, this is so because the idealistically relevant factors suffice to make it the case that there is carbon in R at T. In other words, it seems that he is obliged to assert:

(4) If $C(R,T)$, then I-true that $C(R,T)$.

And, similarly, it seems that he is obliged to say that if there is no carbon in R at T, this is so because the idealistically relevant factors suffice to make it the case that there is no carbon in R at T. In other words, it seems that he is obliged to assert:

(5) If Not $C(R,T)$, then I-true that Not $C(R,T)$.

But then, by steps of *modus tollendo tollens*, (4) combined with (1), and (5) combined with (2) yield respectively:

(6) Not $C(R,T)$.

(7) Not Not $C(R,T)$.

which are contradictories. So, according to the argument, the combination of claims (1) and (2), which expresses the supposed fact of physical indeterminacy, together with certain clear consequences of the idealist thesis, generates a contradiction.

At first sight this argument seems cogent, but on closer scrutiny I think we can see that it is fallacious. The fallacy resides in the supposition that the idealist is obliged to assert the conditionals (4) and (5). For if the idealist accepts that this is a case where the idealistically relevant factors leave open the carbon situation of R at T (in other words, accepts (1) and (2)), he should regard the truth-values of both (4) and (5) as indeterminate. This can be seen by reformulating the conditionals as the equivalent disjunctions:

(4*) Either Not $C(R,T)$ or I-true that $C(R,T)$.

(5*) Either $C(R,T)$ or I-true that Not $C(R,T)$.

On the assumption that (1) and (2) are true, the second disjunct of each of these disjunctions is false, and so the truth-values of the disjunctions will be the truth-values of their first disjuncts. In other words, the truth-value of (4*) will be the truth-value of 'Not $C(R,T)$ ', and the truth-value

of (5*) will be the truth-value of 'C(R,T)'. But, given his acceptance of (1) and (2), the truth-values of both of these will be, for the idealist, indeterminate. This shows that, given his acceptance of (1) and (2), the truth-values of both (4) and (5) will also be, for the idealist, indeterminate. So, an idealist who accepts the indeterminacy expressed by the combination of (1) and (2) is not obliged to assert (4) and (5), and so can reject the argument.

This still leaves the problem that, when we focus on the conditionals in their own terms—independently of a truth-functional evaluation—they seem to be making claims whose truth is indisputable from an idealist standpoint. After all, there is no getting round the fact that the idealist has to think of the world as deriving its whole character from the idealistically relevant factors, and this seems to commit him to saying that, however things are carbon-wise with respect to R at T, they are so because these factors suffice to make them so. So, how can he avoid agreeing to (4) and (5)? But there is an illusion here. When we focus on these conditionals in their own terms, we inevitably understand them as inviting us to envisage a situation in which things are as specified by the antecedent, and as then asserting the consequent in the framework of the supposition that that situation obtains. So, we understand (4) as inviting us to envisage the situation of there being carbon in R at T, and as then asserting, in the framework of the supposition that that is how things are, that the idealistically relevant factors suffice to make it the case that there is carbon in R at T. And we analogously understand (5) as inviting us to envisage the situation of there not being carbon in R at T, and then as asserting, in the framework of the supposition that that is how things are, that the idealistically relevant factors suffice to make it the case that there is no carbon in R at T. But the supposition that a certain type of physical situation obtains is the same as the supposition that it definitely and unequivocally obtains: it leaves no room for the qualification that the matter might be ultimately indeterminate. So, in the framework of the relevant suppositions, thus understood, the assertions of the relevant consequents are, from the idealist standpoint, clearly correct: it is clearly correct to assert, in the framework of the supposition that C(R,T), that it is idealistically true that C(R,T), and clearly correct to assert, in the framework of the supposition that Not C(R,T), that it is idealistically true that Not C(R,T). In this way, it comes to seem to us that the conditionals themselves are ones that the idealist is obliged to accept. And, of course, when the conditionals are understood in

the way envisaged, he *is* obliged to accept them. For, thus understood, the conditionals become equivalent to:

- (8) If definitely $C(R,T)$, then I-true that $C(R,T)$.
 (9) If definitely Not $C(R,T)$, then I-true that Not $C(R,T)$.

where the force of the ‘definitely’ operator is explicitly to exclude indeterminacy. But when the conditionals are reformulated as (8) and (9), the steps of *modus tollendo tollens* from (1) and (2) no longer yield a contradiction. The conclusions they yield—that Not definitely $C(R,T)$ and that Not definitely Not $C(R,T)$ —merely endorse the claim of indeterminacy expressed by (1) and (2).

This disposes of the argument. And, given the failure of the argument, it seems to me that the problems for the idealist over the issue of indeterminacy are at an end. The idealist can allow for cases of indeterminacy without having to abandon the law of the excluded middle in the mundane specification of the world. Thus, he can allow for a case in which the idealistically relevant factors do not suffice to determine whether there is carbon in a particular region at a particular time, but still insist that, even if such a case were to obtain, it would not affect the truth of the mundane claim that either there is carbon present or there is not. And, as far as I can see, he can allow for such cases without coming into conflict with the principles of logic at any point. It remains true that, considered in the abstract, there seems to be something strange about the suggestion that the character of the world may be, in certain respects, indeterminate, so that there is ultimately no answer to the question of whether things are this way or that. But that is simply because when we consider the suggestion in the abstract, we locate it in the framework of our ordinary realist view of the world, which requires us to be able to make sense of the world in its own physical terms. And there is no way of making sense of physical indeterminacy in physical terms.⁶ Any strangeness disappears once it is clearly understood that the indeterminacy that concerns us is tied to the context of an idealist account of the world—an account that represents the world as deriving its existence and character from more fundamental

⁶ Even in the framework of physical realism, of course, there can be indeterminacy in how our physical concepts apply to the world. For example, given the vagueness of the relevant concepts, there may be ultimately no answer to the question of whether a particular person is sufficiently lacking in hair to count as *bald*, or to the question whether a developing oak is sufficiently large to count as a *tree*. But this does not involve any indeterminacy in the character of the world itself.

factors—and that we do not have to be able to make sense of it except in the terms of that account.

One final point. The fact that the relevant form of physical indeterminacy does not affect the law of the excluded middle in the context of the mundane specification of the world still leaves the question of whether, in this same context, it affects the principle of bivalence, which asserts that any proposition must be either true or false. Thus, in the particular case on which we have been focusing, where the idealist accepts claims (1), (2), and (3), should he say that, because of the indeterminacy ascribed to the world by the conjunction of (1) and (2), the proposition that $C(R,T)$ is neither true nor false, or should he, rather, say that, because (3) leaves no room for a physical state of affairs between carbon presence and carbon absence, the proposition that $C(R,T)$ is either true or false? The answer, I think, is that, given the nature of his position, the idealist needs to recognize two pairs of concepts of physical truth and falsity, and the status of the principle of bivalence depends on which pair of concepts is at issue. Thus, on the one hand, there are what we might speak of as the mundane concepts of physical truth and physical falsity. These are our ordinary concepts of physical truth and falsity, which we employ when we think of the physical world in its own terms, without reference to what metaphysically underlies it. They require the assertability of all instances of the schemata ‘The proposition that p is true if and only if p ’ and ‘The proposition that p is false if and only if Not p ’, where ‘ p ’ holds place for a proposition-expressing sentence of the physical language. With these concepts in place, the idealist is committed to accepting the principle of bivalence for physical propositions simply through accepting the law of the excluded middle. So, his acceptance of (3) becomes a commitment to accepting that the proposition that $C(R,T)$ is either true or false. On the other hand, there are what we might speak of as the transcendental concepts of physical truth and physical falsity. These are the concepts of physical truth and falsity that stem from the idealist account, so that a physical proposition that p qualifies as true just in case it is idealistically true (the idealistically relevant factors suffice to make it the case that p), and qualifies as false just in case its negation is idealistically true (the idealistically relevant factors suffice to make it the case that Not p). With these concepts in place, the idealist’s recognition of the possibility of physical indeterminacy obliges him to abandon the principle of bivalence. So, in the particular case in question, his acceptance of both (1) and (2) commits him to concluding that the proposition that $C(R,T)$ is neither true nor false. It is crucial that the idealist should not conflate these

alternative pairs of concepts of physical truth and falsity. Otherwise, he will be led back into contradiction. He will find himself forced to move from a premiss of indeterminacy (such as expressed by the conjunction of (1) and (2)) to the conclusion that things are physically neither one way nor the other (as expressed by the conjunction of (6) and (7)).

6

The Issue of Objectivity

I

If there is to be a physical world at all, it has to conform to the thesis of phenomenalistic idealism: it has to be something whose existence is ultimately constituted by facts about human sensory experience, or by some richer complex of non-physical facts in which such experiential facts centrally feature. For in no other way can the physical world have the empirical immanence it needs if it is to qualify as *our* world—as a world *for us*. We have also seen that, within the range of options allowed by phenomenalistic idealism, the only position with any prospect of acceptability is what I have labelled *canonical* idealism, which assigns the central world-creative role to the sensory organization—to that system of control over the course of human sensory experience that disposes it to conform to its world-suggestive pattern. More precisely, the thesis of canonical idealism is that, either on its own or as part of a larger complex, the sensory organization, in combination with certain endowments of the human mind, constitutively creates the physical world by disposing things to appear systematically worldwide at the human empirical viewpoint, and determines all the details of its specific character, or at least the specific character of its primary core, by the specific content of that disposing.

We saw in the last chapter, however, that canonical idealism itself, as indeed any form of phenomenalistic idealism, seems, at first sight, to be open to a decisive objection. For, on the face of it, the idealist account of the world is in blatant conflict with our basic understanding of the ontological status of the world in relation to human mentality. One aspect of this—the aspect on which we particularly focused—is that when we consider the nature of the physical world in its own terms, and the circumscribed place that human subjects seem to occupy within it, it is hard to see how we can make sense of the suggestion that facts about human sensory experience centrally contribute

to its existence, or that its existence logically depends on facts about human mentality in any way. Thus, within the physical world there is a certain group of complex objects, namely human biological organisms, with which human subjects are in some way intimately associated, each subject being associated with a particular organism, and each organism with a particular subject. Whatever the precise nature of this association, it is hard to deny that, in each instance, it involves, or makes provision for, a complex functional relationship between the subject and the associated organism, whereby, in a variety of ways, the mentality of the subject is equipped to affect the condition of the organism, and the condition of the organism is equipped to affect the mentality of the subject. But once it is accepted that human subjects have this functional relationship with human organisms, it seems impossible to think of the world as ontologically dependent on the human mind, as the idealist claims. Even if we think of human subjects as ontologically independent of the physical world, as the idealist position requires, it seems impossible to avoid accepting that human mentality is just a circumscribed portion of a larger psychophysical whole, whose physical portion exists independently of it. The role that the idealist assigns to the sensory organization makes the apparent problem particularly acute. Within the context of the putative functional relationship, it is the physical world itself that is responsible for the sensory organization, since the subject's sensory experiences are controlled by the organism's sensory input, and the character of this input reflects the character of the organism's environment. But the idealist takes the sensory organization to play the central role in the constitutive creation of the world, and it is surely incoherent to suppose that what is responsible for organizing our experiences derives its very existence, in part, from the obtaining of that organization.

Having identified this apparent problem for the idealist, I tried to show that, on closer scrutiny, it could be seen to be illusory. The crucial point was that human subjects and their mental lives feature in two quite different places in the overall idealist account, and the idealist needs to conceive of their relationship to the physical world in correspondingly different ways in the two cases. On the one hand, in the context of the idealist's distinctive philosophical theory, the organization of human sensory experience plays a central role in the constitutive creation of the world, and so, in that context, the ontological dependence of the physical world on the human mind is an essential ingredient of the idealist's position. On the other hand, with respect to what gets idealistically created, human subjects and the physical world

form interactive partners in an integrated psychophysical whole, and if the idealist is to do justice to the nature of this partnership, he cannot think of it as one in which the physical partner logically depends for its existence on the psychological. In other words, for the idealist, the issue of the ontological relationship between the physical world and human mentality arises in two quite different forms, according to whether he considers how things stand with respect to the nature of the psychophysical reality that is idealistically sustained, or considers how things stand with respect to the nature of this sustainment, and the correct way for him to characterize the relationship depends on which form of the issue applies. The first form of the issue arises in what I termed the *mundane* framework, where the concern is with the character of the physical world and the larger psychophysical reality of which it forms a part considered in their own terms—without reference to what ultimately underlies them—and in this framework it is correct for the idealist, as for the realist, to acknowledge the ontological independence of the world. The second form of the issue arises in what I termed the *transcendental* framework, where the concern is with the nature of the ultimate metaphysical situation, and in that framework the idealist, in contrast with the realist, takes facts about human mentality to centrally contribute to the world's existence.

This disposes of the particular problem—the particular way, on which we focused, in which the idealist's account of the world seems initially to be in conflict with our basic understanding of the ontological status of the world in relation to human mentality. But, as I also noted, there is a further, and quite different, way in which this conflict seems to arise—a way that has nothing to do with the special relationship between human subjects and human organisms. For, on the face of it, something that was idealistically created in the way envisaged would not have the right kind of objectivity, in relation to the human mind, to count as a real world at all. It is true that, by assigning the central role in the creation of the world to the sensory organization, rather than to the world-suggestive orderliness it ensures, the canonical idealist moves things, with respect to objectivity, in the right direction. For at least the sensory organization is something that transcends the actual course of human sensory experience. But, given the way in which the organization plays its putative world-creative role, this transcendence does not remove the basic problem. If, as canonical idealism claims, the sensory organization makes its constitutive contribution to the existence of the world by disposing things to appear systematically worldwise at the human empirical viewpoint, it is hard to see how what (if anything) gets created could count as anything

more than a virtual reality. It is hard to see how we could think of it as forming a genuine world.

Whether the idealist has the resources to deal with this further problem is not something I have yet considered, and it is on this issue that I want to focus in this final chapter. I have already remarked that I see the problem over objectivity as posing the biggest challenge to the idealist position.

There is a great deal at stake here. If it turns out that the challenge cannot be met, we shall be forced to conclude that there is no physical world at all. We shall be forced to say that there can be nothing that has both the empirical immanence needed to qualify as *our* world and the objectivity needed to qualify as a *real* world, and so there can be nothing that satisfies all that the existence of a physical world requires. This, in itself, would be a very unpleasant outcome. It would also undermine our epistemological situation in ways we might not have anticipated. We might have supposed that, even if we had to give up our belief in a physical world, our knowledge that there was an experiential simulation of a world—that our sensory experiences were organized so as to provide the systematic empirical impression of a world—would serve our theoretical and practical purposes almost as well. We might have supposed that we could continue to recognize the theoretical and practical appropriateness of our ordinary belief in a physical world, and, by and large, of our specific beliefs about it, but just bear in mind that what secures this appropriateness is not that the beliefs are actually true, but that they encapsulate the distinctively world-suggestive ways in which our experiences are organized. But, as we saw, once we deny the existence of a physical world, we undermine our entitlement to suppose that this experiential organization obtains. For almost all the knowledge we could think of ourselves as having about the experiential realm either directly rests on knowledge we take ourselves to have about the physical world, or depends for its epistemic credentials on the assumption of the existence of a physical world and our epistemic access to it. If we had to reject the existence of the physical world, we would find ourselves, at the level of philosophical reflection, not only bereft of our physical beliefs, but in an epistemic wilderness.

A great deal, then, hangs on whether the idealist can find a satisfactory way of dealing with the problem of objectivity. One thing that he cannot do is simply to appeal, once again, to the distinction between the mundane and the transcendental frameworks. He cannot just point to the fact that, in the mundane framework, where the world is described in its own terms, and

where it has to be represented as an interactive partner of human subjects within a larger psychophysical reality, the world's objectivity in relation to human mentality is secure. For the issue of objectivity that now concerns us is not to do with how things stand mundanely, but with how they stand transcendently. It is not the issue of whether, within the psychophysical reality that is supposedly idealistically sustained, the physical portion has the required objectivity in relation to the psychological; as we have seen, there is no denying that, in terms of the make-up of that putative reality, the physical portion is ontologically independent of the human mind and so enjoys the requisite objectivity. Rather, the relevant issue is whether, when we stand back from the perspective of the mundane framework, we can see the factors that the idealist takes to be metaphysically fundamental as sufficing to create a world that is real. After all, even in the case of something that is indisputably no more than a systematic illusion of a world, like what is depicted in the film *The Matrix*, the illusory world will be objective in its own terms. The illusion will represent the situation of a world whose existence is logically independent of the mentality and existence of the subjects involved.

As I have already indicated, the problem of objectivity and the issue of whether the idealist has the resources for dealing with it are what I shall be particularly focusing on in this chapter. But, before I go any further with this, I want, briefly, to turn away from the topic of objectivity to say something about the idealist's view of the human mind. Although this topic is not directly to do with the problem of objectivity, and is one that I would need to cover even if that problem had not arisen, the idealist's view of the mind will, as it happens, turn out to be crucial to the solution to the problem that I shall eventually propose.

II

In taking the physical world to be constitutively created by facts of a different kind, the idealist is committed to thinking of these latter facts as ones whose obtaining is logically independent of the existence of the physical world. He cannot think of any of these constitutively relevant facts as ones to whose obtaining physical facts in any way logically contribute; for if he did, the lines of logical dependence would be rendered viciously circular. This means, in particular, that the idealist has to ascribe world independence to the realm of human sensory experience, whose organization plays the central role in the

idealistic creation. He has to think of sensory experiences and their subjects as wholly non-physical in nature, and he has to think of their existence, and the obtaining of all the constitutively relevant facts about them, as owing nothing to how things stand in the physical reality to whose sustainment they centrally contribute.

In specifying the idealist's commitment in these terms, I am assuming, what I have been tacitly assuming throughout, that the idealist accepts the common-sense view of the ontological structure of the mind, which takes experiences, thoughts, beliefs, emotions, and other concrete items of mentality to be, and to be irreducibly, the token states, acts, and activities of persisting mental subjects. Strictly speaking, this is not the only option available. Some philosophers, following David Hume, reject any fundamental ontology of mental subjects and construe individual minds as nothing more than organized collections of appropriately related mental items.¹ The idealist's reductive account of the physical world does not itself prevent him from adopting this view: he could take the sensory organization to play the central role in the constitutive creation of the world, but think of the minds on which this organization is imposed as composed of nothing but sensory experiences and other mental items. Even so, the view seems to me to be clearly mistaken. My main objection to it is not the one that is most commonly made, that the Humean has no satisfactory way of defining the unity of the mind—of specifying what it is for different mental items to belong to the same mind. In fact, I think that, by exploiting the fact that total experiences are extended in time, and the fact that successive items in a stream of experience mereologically overlap, it is possible to construct an adequate definition of the unity of the mind in purely Humean terms.² My main objection to the Humean position is the more basic one, that the very notion of subjectless mentality is unintelligible. Thus, I can no more understand how there could be a thought without a thinker, a belief without a believer, or an experience without an experiencer, than I can understand how there could be speech without a speaker, or motion without something that moves. So, even though the Humean option is available to the idealist, it is one that I shall continue to ignore.³

¹ Hume sets out his view in his *A Treatise of Human Nature*, ed. L. Selby-Bigge, 2nd edn., rev. P. Nidditch (Oxford: Oxford University Press, 1978), Bk. 1, Pt. 4, sect. 6.

² I am thinking of my definition of the relation I term *co-personality* in *The Immaterial Self* (London: Routledge, 1991), ch. 8.

³ For a fuller discussion and criticism of the Humean view see my *The Immaterial Self*, 212–19.

The idealist is obliged to think of the realm of human sensory experience as logically independent of the physical world, but he is not obliged to ascribe such world independence to the realm of human mentality as a whole. For where forms of mentality are not involved in the idealistic creation, there is no incoherence in supposing that physical factors logically contribute to their realization. And, indeed, whatever views we have of the metaphysical status of the world and of the nature of the human mind, it is uncontroversial that some forms of mentality logically depend for their realization on physical factors, simply because they inherently involve a psychological relation between the subject and something physical. For example, it is logically impossible for someone to be thinking about Venice, or to be hearing the chimes of Big Ben, or to be fond of gold, without the existence (though not necessarily the *current* existence) of the physical item which forms the object of his mental state or activity. Likewise, it is logically impossible for someone to know that it is currently raining, or that there was a frost last night, without the obtaining of the physical state of affairs that is known. What is controversial, of course, is whether these physically relational forms of mentality are, or are ever, psychologically fundamental. Some philosophers, the so-called *internalists*, think that whenever a subject stands in a psychological relation to a physical item, his fundamental psychological state is one which is not inherently physically relational, and the relational fact is one that is constituted by the subject's being in this state, together with certain further facts. Other philosophers, the so-called *externalists*, think that, at least in certain cases, the psychological relationship between a subject and a physical item is a fundamental aspect of the situation and not something which decomposes, at the psychological level, into other factors. We have already focused on a specific form of the externalist–internalist issue in our initial discussion of the nature of perception, where the decompositional view represented the approach of the internalist, and the fundamentalist view represented the approach of the externalist. (The labels 'internalist' and 'externalist' were, of course, introduced by philosophers who accepted a realist view of the physical world.)

In taking the subjects of mentality to be wholly non-physical, the idealist is committed to thinking of all mental states, acts, and activities as wholly non-physical on their mental side—as non-physical, as it were, in their mental substance. This means, I think, that the only forms of mentality that he could take to be world-dependent would be, as in the examples above, ones that were inherently physically relational. It might be thought that, in the case of these physically relational forms, the idealist's reductive account of

the physical world would prevent him from adopting an externalist position. For how could the physical world supply items to feature in psychologically fundamental relationships if its existence is ultimately constituted in the way that the idealist envisages? But this is an area where, as before, we must be careful to bear in mind the distinction between the mundane and the transcendental frameworks. It is true that, in the transcendental framework, where he offers his account of the ultimate metaphysical status of the world, the idealist will not be able to recognize psychologically fundamental relationships between subjects and physical items, since this account represents the world, and all its ontological ingredients, as constitutively created by facts of a different kind, in which facts about human sensory experience centrally feature. And since facts about sensory experience are psychological facts, this obliges him to accept that any psychological relationship between a subject and a physical item ultimately decomposes into more fundamental factors at the psychological level, even if the mode of decomposition is not of the sort standardly envisaged by the internalist. But this still leaves him free to suppose that some of these psychological relationships are fundamental relative to the psychophysical reality of which the created world forms a part—that, within that reality, they do not break down into further factors. In other words, it leaves him free to accept externalist accounts of them at the mundane level, where we are concerned with the character of the world and the larger psychophysical reality in their own terms.

To what extent the idealist *should* accept such accounts, at the mundane level, is a further question, though not one that I shall investigate here: for the most part, it would not have any relevance to the topic of idealism itself. Such thought as I have given the issue inclines me to the view that there are a number of areas where a decompositional approach would not do justice to the nature of the relational mentality involved, and where mundane externalism should be embraced. Certainly, this is so in the case of physical-item perception, where, as we already know, a decompositional account, at the mundane level, would not allow us to think of the relevant physical items as objects of perceptual awareness at all. And in this area, as we have seen, idealism at the transcendental level is able to support the cause of mundane externalism. For what idealism is able to show, by the form of its reductive account of the physical world, is how a genuine, and mundanely irreducible, perceptual awareness of physical items is possible—something that cannot be shown when we confine our attention to the character of the psychophysical reality in its own terms.

However he sees things as turning out in the psychophysical reality that is idealistically sustained, the idealist is committed to thinking that at the level of what is metaphysically fundamental human mentality is entirely logically independent of the physical world; and, of course, he is committed to thinking of this world-independent mentality as including the sense-experiential mentality that features so centrally in his theory. By according this world independence to human mentality in its metaphysically fundamental form, the idealist's view of the human mind coincides with that of the Cartesian dualist—though, of course, the dualist holds this view within the framework of a realist understanding of the world. I say that it coincides with the view of the *Cartesian* dualist, because, as I have already made clear, I am assuming that the idealist endorses the common-sense view of the ontological structure of the mind and takes mental items to be, and to be irreducibly, the token states, acts, and activities of persisting mental subjects. It was this acceptance of a fundamental ontology of mental subjects that distinguished Descartes's dualist conception of the mind from that of Hume.

In the current philosophical scene the Cartesian view of the human mind, as indeed any radically dualist view, is very unfashionable. Almost all current philosophers take it for granted that the fundamental subjects of human mentality are entities whose natures are, at least in part, physical, and the overwhelming majority take these subjects to be biological organisms, whose natures are purely physical. Again, a majority of current philosophers not only take the fundamental subjects of human mentality to be purely physical, but additionally accept some form of physicalistic account of the mentality itself, either by directly construing it as physical in nature (so that mental states, acts, and activities become purely physical states, acts, and activities), or by taking its occurrence to be in some way reducible to physical facts. The suggestion that, at the metaphysically fundamental level, human subjects may be wholly non-physical and their mentality logically independent of the physical world is, by and large, no longer considered a serious option. Assuming that current philosophers have some rational basis for their dismissal of the Cartesian view, there seems to be, here, the prospect of a further problem for the idealist. Not only does his idealist theory seem, on the face of it, inadequate to secure the requisite objectivity of the world, but it also employs a conception of the mind that is now standardly regarded as wholly discredited.

There are two reasons why I think that the idealist need not be worried by this new challenge to his position.

In the first place, although most of the philosophy of mind over the last fifty years has been in radical opposition to the Cartesian view, and to any radically immaterialist account of the mind, this opposition has been advanced within the framework of a realist view of the physical world, and cannot be detached from that framework. Thus, when these opposing philosophers insist that the fundamental subjects of mentality are entities with wholly or partly physical natures, and when some of them additionally insist that the mentality itself is to be ultimately understood in physical terms, they are taking it for granted that the physical world is something whose existence is both logically independent of the human mind and philosophically fundamental, and their claims are only intelligible within that framework. But, whatever the final verdict on canonical idealism, we have already established, and established without making prior assumptions about the ultimate nature of the mind, that physical realism is untenable, and so we know that the framework for these physicalistic and semi-physicalistic accounts of the mind is no longer available. It does not follow from this, of course, that the philosophers who have offered these accounts do not have well-founded objections to the non-physicalistic alternatives. But, in fact, like the accounts themselves, these objections too, at least when they have any initial plausibility, implicitly depend on a realist conception of the world: they turn on the supposed problems of maintaining a radically non-physicalistic understanding of human subjects and their mentality in the context of an understanding of the world as something mind-independent and fundamental.

Second, I have tried to show elsewhere that, even from the standpoint of physical realism, the Cartesian view of the mind can be seen to be correct.⁴ Thus, adopting a realist conception of the world, I have tried to show that there is no form of physicalistic or reductive account of human mentality that does justice to the nature of such mentality—in particular, no such account that does justice to the character of conscious experience as introspection reveals it. And I have tried to show that there is no way of understanding how mentality can be non-physical and irreducible without the subjects to which it fundamentally belongs being non-physical too. In addition to offering this positive case for the Cartesian view, I have tried to answer what I take to be the most important of the objections brought against it. In particular, I have

⁴ This is the central aim of my *Immaterial Self*. See also my essay 'A brief defence of the Cartesian view', in K. Corcoran (ed.), *Soul, Body, and Survival* (New York: Cornell University Press, 2001), 15–29, where I cover some of the main points more succinctly.

tried to show, in some detail, that the dualist has satisfactory ways of dealing with the various difficulties that are thought to beset the dualist's conception of causal interaction between mind and body—both the supposed difficulties in making sense of dualistic interaction at all and the supposed difficulties over reconciling it with the theories of physical science. This is not the context in which to rehearse my arguments on these issues, whether the arguments which seek to establish the truth of the Cartesian view or the arguments to defend it against its critics, but I still consider them cogent.

I shall assume, then, that, whether or not he can successfully deal with the problem of objectivity, the idealist is at least secure in his Cartesian view of the mind—the view required by the role he assigns to human mentality in the constitutive creation of the physical world.

III

Canonical idealism, as I have defined it, assigns the central role in the constitutive creation of the world to the sensory organization, but it does not claim that this organization is creatively sufficient on its own. Part of the reason why I have left room for other constitutively relevant factors is already clear from our previous discussion. The sensory organization is that system of control over the course of human sensory experience that disposes it to conform to its world-suggestive pattern, and the way in which this contributes to the constitutive creation of the world is by disposing things to appear systematically worldwise at the human empirical viewpoint. But, as I have all along recognized, the disposing of sensory experience to conform to the relevant pattern only disposes things to appear systematically worldwise at this viewpoint because it occurs within the context of certain endowments of the human mind—endowments that make us empirically sensitive to the orderly character of our sensory experiences and equip us to register its world-suggestiveness. I have not tried to settle precisely what the relevant endowments are. All I have insisted, I hope uncontroversially, is that they include a natural doxastic bias in favour of regularity—a natural propensity of the mind to try to represent things as working in a uniform way. So, I have said that, although the orderliness of our sensory experiences does not, in itself, involve much in the way of regularity, part of the reason why it empirically invites us to think of ourselves as perceptive inhabitants of a unitary three-dimensional world (why it makes things appear systematically

worldwise at the human empirical viewpoint) is that this way of representing our situation is what would be needed to achieve the regularity that our minds doxastically seek. But whatever endowments we take to be relevant, the idealist has to include their presence, along with the sensory organization, in the factors that play a role in the constitutive creation of the world.

But there is also a quite different reason why I have defined canonical idealism in a way that leaves room for other constitutively relevant factors, and this reason is directly concerned with the issue of objectivity: it is directly to do with a certain way in which the idealist might hope to be able to give the idealistically created world the requisite objectivity in relation to the human mind.

Under physical realism we would take the sensory organization to result from the way in which the physical world itself causally controls the course of human sensory experience. The idealist can accept that this is how things are to be described *mundanely*, where we are concerned with the character of the idealistically sustained psychophysical reality in its own terms. But he cannot think of the physical world as what is *ultimately* responsible for the sensory organization, since he excludes the world from the domain of what is metaphysically fundamental; indeed, he takes the sensory organization to centrally feature in the factors from which the physical world derives its existence. Even so, the idealist is still at liberty, in his account of what is metaphysically fundamental, to suppose that there is some form of external reality that causally controls the course of human sensory experience, and is thereby responsible for the sensory organization; and since it is hard to avoid thinking that there is something that is responsible for this organization, and not easy to think of anything in the realm of human mentality to which such responsibility could be ascribed, the idealist is very likely to favour this view. One specific version of the view, as we have seen, would be to postulate an external reality that mirrors the structure and functional organization of the realist's physical world, and which is functionally linked with the realm of human experience by the analogue of psychophysical laws.

Now, for the sake of argument, let us suppose that the idealist does recognize the presence of some form of external controlling reality, of whatever nature. He then has two options as to how he could take this reality to feature in his overall account. On the one hand, he might say that the only relevance this reality has to his idealist position is in forming the source of the sensory organization, and that it makes no constitutive contribution to the existence of the physical world other than what is mediated by the

contribution of the organization itself. This would mean that, in giving his account of the constitutive creation of the world, any reference to this reality would become redundant once the constitutive role of the organization had been specified. On the other hand, he might say that, in addition to what is mediated by the sensory organization, the existence of the external reality and the control it exercises over our sensory experiences make an independent constitutive contribution to the existence of the world—a contribution that combines with the contribution of the organization and the relevant endowments of the human mind. This would mean that, to understand the nature of the idealistic creation, we would have to take account not only of the way in which—in the framework of the relevant endowments—the sensory organization disposes things to appear systematically worldwide at the human empirical viewpoint, but also of the external factors from which the disposing organization stems. It is the availability of this second option that is relevant to the issue of objectivity. The *prima facie* problem for the idealist, as we have seen, is that something that is constitutively created in the way he envisages does not seem to have the kind of objectivity it needs, in relation to the human mind, to qualify as a real world. It seems that, with so much weight resting on how things are disposed to appear to us through the medium of our sensory experiences, the most that could idealistically result would be a virtual reality—the experiential simulation of a world. If the idealist takes the sensory organization and the endowments of the mind to be the only factors involved in the constitutive creation of the world, it is hard to see how he could begin to deal with this problem. But the availability of the second option opens up the prospect of a possible solution. For, by taking the sensory organization to have an external source of a suitable kind, it may be possible for him to provide a sort of objective underpinning to the world-suggestive way in which things are disposed to empirically appear—an underpinning that renders this system of appearance in some sense objectively appropriate. And once he has secured such an underpinning, there is the prospect that, by taking it to play an essential role in the idealistic creation (so that it is only as a result of this underpinning that the sensory organization is thought of as idealistically effective), he can secure the requisite objectivity in the item created.

In other words, the suggestion is that the idealist may be able to achieve a satisfactory account of the situation by dividing the work of the idealistic creation into two components, with complementary roles. One component, with which we are already familiar, would cover the contribution of the

sensory organization and the relevant endowments of the human mind, and would turn on the way in which these factors dispose things to appear systematically worldwise at the human empirical viewpoint. Its role would be to ensure that what gets idealistically created has the right kind of empirical immanence to form a world for us. The other component, which is what is new, would cover the contribution of the external source of the sensory organization, and would turn on the fact that this source provides an objective underpinning to the way in which things are disposed to empirically appear. Its role would be to ensure that what gets idealistically created has the right kind of objectivity, in relation to the human mind, to qualify as a real world.

I shall speak of this way of trying to deal with the problem of objectivity as the *externalist strategy*. Whether this strategy can be successfully implemented remains to be seen. What seems clear to me is that there is no other strategy worth considering. There is, it seems to me, no prospect at all of being able to regard the sensory organization and the endowments of the mind as sufficient, on their own, for the constitutive creation of a genuine world. And, once it is agreed that there has to be some additional constitutive factor, I cannot see what else the idealist could hope to turn to other than the presence and controlling role of an external reality.

Before we look into the prospects of the externalist strategy in more detail, there is an important procedural point that needs to be recognized. To successfully implement the strategy, the idealist needs to be able to envisage a form of external controlling reality which, when its presence and mode of control are added to the other factors that he takes to be constitutively involved, will ensure the constitutive creation of a real world. But in judging whether a given form of reality will suffice for this purpose, we must be careful to approach the issue in the right way. In the abstract, it would be easy to think that, whenever a fact F is constituted by a fact F' , or set of facts S , a knowledge of the obtaining of F' (the members of S), combined with a proper understanding of what the obtaining of F would involve in its own terms, would, on its own, oblige us to recognize the consequential obtaining of F . Applied to the particular case that concerns us, this would entail that, for a given account of the constitutive creation of the world to be acceptable, it has to be the case that a knowledge of the obtaining of the factors that it takes to be constitutively relevant, combined with a proper understanding of what the existence of a physical world would involve in its own terms, would, on its own, oblige us to recognize the consequential existence of such

a world. But this view of what the acceptability of an idealist account requires is mistaken, and it is important that we should see that this is so.

I can best bring out the point by focusing on an analogous case that featured in our earlier discussion. The case I have in mind is one in which a realist about the physical world entertains the hypothesis that the fundamental physical reality is ontologically composed merely of space and certain region-characterizing properties, and that the persisting space occupants of our ordinary conceptual scheme are ontologically derivative entities, whose existence is constituted by the spatio-temporal distribution and functional organization of the properties.⁵ The precise way in which this hypothesis would need to be developed depends on what kinds of persisting occupant are taken to be, within the domain of all such occupants, relatively fundamental: in the earlier discussion I focused on a case in which these relatively fundamental occupants were taken to be spherical particles. But the details of the hypothesis need not here concern us. What matters is the combination of two things. The first is that, within the framework of physical realism, the hypothesis is surely coherent. There is surely nothing in our concept of a persisting space occupant that excludes the possibility that the occupants of our ordinary conceptual scheme are ontologically derivative in that way. Our ordinary belief in the existence of such occupants is surely tolerant of the hypothesis that what ultimately underlies their existence are distributional and organizational facts of the relevant kind. The second thing is that, if the realist were to set aside his prior knowledge of the existence of persisting space occupants, and just suppose himself to know that properties of the relevant kind are distributed and organized in the relevant way, then, even with a proper grasp of the notion of a persisting occupant, there would be nothing to oblige him to recognize a derivative ontology of such occupants. There would be nothing to force him to go beyond the conclusion that the properties are distributed and organized in an occupant-simulating way. But if it is coherent for the realist to hypothesize that the occupants of our ordinary conceptual scheme are ontologically derivative in the relevant way, but a knowledge of the relevant kinds of distributional and organizational fact would not, in itself, oblige a recognition of the existence of such occupants, then (irrespective of whether there is anything in its favour) the hypothesis cannot be rejected on the grounds that this obligation would not obtain.

⁵ See Ch. 3, Sect. II.

The same point applies in the case of canonical idealism and its capacity to deal with the problem of objectivity. If we want to decide whether a given idealist theory adequately deals with the problem—whether the factors it takes to be involved in the constitutive creation of the world make adequate provision for the objectivity of what is created—the procedure should not be to consider whether, without appeal to our prior knowledge of the physical world, the supposition that these factors obtain would oblige us to recognize the real existence of a world. That would be to impose too severe a test on the adequacy of the theory. Rather, the procedure should be to consider whether our ordinary belief in the physical world, which we hold quite independently of any beliefs about the metaphysically fundamental reality, is tolerant of the theory's account of how the existence of the world is ultimately sustained, or whether embracing the theory's account of how things stand fundamentally would force us to give up that ordinary belief. Tolerance, of course, would not ensure that the theory was correct. But it would ensure that the factors which the theory takes to be involved in the creation of the world are equipped to perform that creative role. And it would mean, in particular, that these factors make adequate provision for the objectivity of the world.

What we now have to investigate is whether the externalist strategy is able to deliver a theory of this kind.

IV

To try to implement the externalist strategy, the idealist has to begin by identifying a suitable type of external reality—a reality that he can plausibly take to be responsible for controlling the course of human sensory experience, and whose character and mode of control he can see as providing the right kind of objective underpinning to the way in which things are disposed to appear at the human empirical viewpoint. It is only after he has managed to secure this element of objectivity in the realm of empirical appearance that he can hope, by assigning it an essential role in the idealistic creation, to secure the requisite objectivity in the item created. One obvious possibility, here, would be to suppose that what controls our sensory experiences, and thereby determines the actual and potential empirical evidence available at the human viewpoint, is an external reality with a structure and organization that exactly match those of the notional world which this evidence suggests.

More precisely, it would be to suppose that the structure and organization of the external controlling reality exactly match what the evidence suggests at a suitably basic level of empirical enquiry. I add this refinement because at the level of ordinary observation the empirical evidence represents the world as having qualitative elements of a distinctively sensible kind, and neither these elements themselves, nor the aspects of structure and organization that concern them, could be replicated in the external reality. So, if, at the relevantly basic level, the empirical evidence, actual and potential, suggests the presence of a three-dimensional world, with certain kinds of space occupant, and governed by certain laws (both laws that concern its internal workings and laws that concern its causal links with human mentality), the idealist would, on this proposal, postulate an external reality of exactly that sort, and would take the world-suggestive character of the empirical evidence to be something that results from, and directly reflects, the character of this reality. In effect, the proposal coincides with something that we have already entertained on a number of occasions, when we have envisaged the possibility of the idealist postulating an external reality with the structure and organization of the realist's physical world.

It is clear how the presence and causal role of this kind of external reality could be thought to provide a suitable objective underpinning to the way things are disposed to empirically appear. For it would mean that how things are disposed to empirically appear, or at least appear in the relevant basic respects, would be a projection on to the human viewpoint of how they externally are; it would mean that, at least in those respects, the world-suggestive character of the system of appearance would embody a correct representation of the reality outside it. All, it might be thought, that the idealist would then need to do, in order to accommodate the objectivity of the physical world, would be to take the presence of this reality, and its role in controlling our sensory experiences, to be contributory factors in the idealistic creation. What was thereby created would be an empirically immanent world, since it would draw its character from how things are disposed to appear at the human viewpoint. But this created world would also, it might be thought, have a certain objectivity in relation to the human mind, since how things stand at the human viewpoint would be an accurate reflection of how they stand externally, and the role of the external reality in ensuring that reflection would be involved in the creation. This objectivity might well be thought adequate for the idealist's purposes, entitling him to claim that what the idealistically relevant factors

succeed in creating is a real world. Certainly, I used to think it adequate myself.⁶

If the idealist can succeed in endowing the idealistically created world with the required objectivity in this way, it is plausible to think that he could make the approach more flexible. For if he can give the system of appearance the right kind of underpinning by postulating an external controlling reality whose structure and organization exactly match those of the notional world that is empirically suggested, he could presumably also provide this underpinning by envisaging an external reality whose structure and organization merely functionally simulate those of this world. For example, if the empirical evidence (actual and potential) unequivocally suggests that, in the world we inhabit, there are three types of fundamental particle, the idealist could presumably afford to entertain the hypothesis that the external reality that causally lies behind that evidence is one that is exactly like that notional world, except that, corresponding to one of the three empirically suggested types of particle, there is a pair of types of external particle that are functionally indistinguishable. Or, to return to a familiar, and more radical, case, although the empirical evidence unequivocally suggests that the world we inhabit contains a single three-dimensional space, the idealist could presumably afford to entertain the hypothesis that the external reality contains two three-dimensional spaces, which are organized as if they were joined. In both these cases—and, of course, there are a host of others we could envisage of a similar kind—it would no longer be true that how things are disposed to empirically appear exactly coincides, in structure and organization, with the way they externally are. But how things are disposed to appear would still, in its way, exactly reflect the overall external situation, since the external reality would be organized exactly as if it had the structure and organization that were causally projected on to the human viewpoint. It is hard to think that the role of the external reality would serve the idealist's purposes in the original case, where there was an exact match, but fail to serve them in these new cases too. However, for the time being I shall set cases of functional simulation aside and focus exclusively on the case where the postulated reality exactly matches what the empirical evidence, at the relevantly basic level, suggests. It certainly seems that the idealist could afford to allow for the possibility of functional simulation, but there is no point in

⁶ See John Foster, *The Nature of Perception* (Oxford: Oxford University Press, 2000), 281–2. As will emerge, this is an issue on which I have radically changed my views.

trying to reach a definite verdict on this until we have established whether a matching reality would give the idealist what he needs.

The suggestion before us, then, is that the idealist can solve the problem of objectivity by (i) supposing that what ultimately causally controls the course of our sensory experience is an external reality whose structure and organization match the structure and organization of the notional world that the actual and potential empirical evidence, at a suitably basic level, suggests, and (ii) taking the presence of this reality and the mode of its control over sensory experience to play a supporting role in the constitutive creation of the world. The idealist will continue to insist, as the core of his position, that the central role in the creation of the world is played by the sensory organization, in the context of the relevant endowments of the human mind, and that it plays this role by disposing things to appear systematically worldwise at the human empirical viewpoint. But the new idea is that, if the sensory organization stems from a controlling reality of the kind envisaged, the world-suggestive way in which it disposes things to appear at our viewpoint will gain a kind of objective underpinning, since the content of the appearance, in the relevantly basic respects, will be an accurate reflection of the external situation. And the suggestion is then that, by postulating this kind of underpinning, and taking it to be an essential factor in the idealistic creation, the idealist can secure the requisite objectivity in the item created. Admittedly, the idealist would not be able to think of this underpinning as covering every aspect of the system of appearance, since, as I have made clear, the envisaged external reality only conforms to what the empirical evidence suggests at a relevantly basic level of empirical enquiry. Specifically, the distinctively sensible aspects of how things are disposed to empirically appear would not be underpinned by the external situation in the envisaged way, and in the idealist's system these aspects are allowed to contribute to the character of the created world. But the claim would presumably be that, if the system of appearance acquires a suitable objective underpinning in the respects in which it accurately reflects the external situation, the inclusion of this underpinning in the factors that constitutively create the world would suffice to confer the required objectivity on the world as a whole. Whether such a claim is, in the end, defensible is a further question, but, for reasons that will shortly emerge, not one that we need to pursue.

Leaving the last issue aside, does this approach give the idealist what he needs? Well, we might begin by noting that there is some plausibility in the idealist's supposing that sensory experience is controlled by an external reality

of the relevant kind. As I said earlier, it is hard to avoid thinking that there is something that is ultimately responsible for the sensory organization, and not easy to think of anything in the realm of human mentality to which such responsibility could be assigned. And once we are in the business of trying to account for this organization by appeal to something external to the mind, the postulation of a controlling reality that matches what is empirically suggested—that imprints its own character on the empirical evidence—offers a particularly straightforward option. But the crucial question is whether the postulation of this kind of reality can help the idealist with the problem of objectivity. Canonical idealism assigns the central role in the creation of the world to the sensory organization, operating in the context of certain endowments of the human mind, and it takes the organization to play this creative role by the way in which it disposes sensory experience to provide the systematic appearance of a world at the human empirical viewpoint. If the postulation is to help the idealist with the problem, the existence of the relevant kind of reality and its control of the course of human sensory experience must be such that, when incorporated into the factors that feature in the idealistic creation, they enable what is created to qualify as an objectively real world; and they can only do this if they provide the system of empirical appearance with the right kind of objective underpinning. It is in these terms that the suggested approach needs to be evaluated.

There is no denying that the approach has a certain initial plausibility. Intuitively, it is much easier to accept that the idealistically relevant factors suffice for the existence of a real world when they include the existence of an external reality of the relevant kind. But at least part of the reason for this is that, intuitively, our judgements still tend to be influenced by the realist outlook that characterizes our ordinary thinking—an outlook that takes the physical world itself to be an external reality. The result is that the inclusion of the relevant kind of reality among the idealistically relevant factors gives the intuitive impression of improving the prospects of the idealist's position, simply because it brings it closer to what is envisaged under realism. The idealist account comes to look more acceptable, because the kind of external reality envisaged, as part of what supposedly constitutively underlies the existence of the world, is something which, from the standpoint of our ordinary realist outlook, meets the actual requirements for such a world. Obviously, if we are to make a proper judgement about the situation, we need properly to distance ourselves from this ordinary outlook and address the issue in appropriately idealist terms.

When we do, it soon becomes clear that, with respect to the problem of objectivity, postulating an external reality of the envisaged kind does nothing to help the idealist at all, since it does nothing to give the system of empirical appearance the objective underpinning it needs. It is true that, if the structure and organization of the external reality match those of the notional world that is empirically suggested, the system comes to acquire a *kind* of objective underpinning. It acquires that underpinning simply by being an accurate reflection of the external situation—by the fact that what the system suggests, and invites us to believe, is, in respect of how things externally are, true. But veridicality in relation to the external situation is an objective underpinning of a *realist* kind. It is what would be needed if the role of the system was that of providing us with information about a reality that was ontologically independent of the human mind. It does not help to underpin the system of appearance in a way that supports its role in the context of the idealist's theory. In the context of that theory, the role of the system is, in effect, to be world-creative. For the system is put in place by the sensory organization, operating within the framework of the relevant endowments of the human mind, and it is by putting the system in place—by disposing things to appear systematically worldwide at the human empirical viewpoint—that this organization and these endowments, with the backing of any other factors that are needed, are taken to constitutively create the physical world. The kind of underpinning, then, that the system of appearance needs, in the context of the idealist's theory, is one that will enable it to play its role as the central element in what idealistically creates a world (or as that through which the sensory organization and the endowments of the mind are centrally involved in such a creation). And the particular respect in which it has to equip the system to play this role is that, when included as an idealistically relevant factor, it has to ensure that what is idealistically created has the requisite objectivity to qualify as a real world. It is impossible to see conformity to the external controlling reality as contributing anything to the capacity of the system to play that role. There is absolutely no basis for supposing that, when it is included as an idealistically relevant factor, such conformity would help to equip the system to give rise to a new reality that conforms to it—that it would do anything to enable the ways things are disposed to appear at the human empirical viewpoint to make it the case that that is how they are.

The suggested approach, then, fails, and fails at its most fundamental point. Postulating a controlling reality of the relevant kind is a straightforward way of

accounting for the sensory organization, and for the world-suggestive way in which it disposes things to empirically appear. But it does not help to give the system of appearance the kind of objective underpinning it needs from the standpoint of the idealist's theory, and so does not help the idealist with the problem of objectivity. Moreover, if the approach fails when developed in the straightforward way we have been assuming, where the external reality is required to exactly match what the empirical evidence, at a suitably basic level, suggests, it is obvious that it would also fail, for the same reason, if it were developed in its more flexible form, where functional simulation is allowed as an alternative. No kind of preservation of the character of the external situation in how things are disposed to empirically appear at the human viewpoint could assist the system of appearance in its role as world-creative.

We are not entitled to conclude from this that the externalist strategy itself is misconceived. It could still be that the right way for the idealist to deal with the problem of objectivity is by postulating an external controlling reality that provides an objective underpinning for the system of empirical appearance, and by then adding the controlling presence of this reality to the factors that feature in the idealistic creation; and, as I have indicated, I do not think that there is any alternative course worth considering. But if there is to be an effective way of pursuing this strategy, it will have to be by envisaging a form of underpinning that properly reflects the significance of the system of appearance in the idealist's theory. It will mean finding a form that is appropriate to the fact that the idealistic role of the system is not to provide some kind of representation of the external situation, but to centrally feature in the creation of a world.

V

There is, I think, a way of pursuing the externalist strategy that gives the idealist what he needs. It involves ascribing the ultimate responsibility for the sensory organization to God, and taking the fact of God's responsibility, in the context of the purposes that underlie it, to be what provides the right kind of objective underpinning for how things are disposed to empirically appear. I shall speak of this as the *theistic approach*, and it is this theistic way of trying to implement the externalist strategy that I now want to elaborate and defend.

We are looking for a form of underpinning for the system of empirical appearance that properly reflects the significance of that system in the idealist's theory—a form that is appropriate to the role of the system in the context of the idealistic creation—and so the central challenge for someone wanting to adopt the theistic approach is to show that God's responsibility for the sensory organization would provide an underpinning of that sort. I shall turn to this task presently. But we need to begin by getting clearer about what is involved in ascribing responsibility for the organization to God. And the first thing we need to note is that there are two very different forms that we could think of this responsibility as taking, leading to two distinct versions of the theistic approach. On both versions we are to think of God as having the ultimate causal control over the course of human sensory experience, and of deliberately exercising that control in a way that ensures the conformity of our experiences to the relevant world-suggestive pattern. And on both versions we are to think that God's purpose, in ensuring this conformity, is to ensure the systematic empirical appearance of a world at the human viewpoint. But the two versions offer very different accounts of the method by which God exercises his control and achieves his objectives.

The externalist approach that we earlier envisaged involved assigning control of human sensory experience to an external reality that functioned, in both its internal workings and its effects on human mentality, in a purely mechanistic fashion, under the governance of certain natural laws. It involved thinking of the external reality as having a controlling role of the same sort as the realist assigns to the physical world; in its simplest version, it involved taking this reality to be, in structure and organization, exactly like the realist's world. One way of pursuing the theistic approach would be to preserve the postulation of an external reality of this mechanistic sort, but add a recognition of God as the causal source of that reality. In other words, we would think of God not as controlling human sensory experience directly, but as causally creating an external mechanistic reality which, through its internal structure and laws, and through the laws that link it with human mentality, would do the controlling for him. And we would then think of this mechanistic reality, and its mode of control, as the immediate source of the sensory organization.

It is important that we should not think of this version of the theistic approach as merely a theistically developed version of the earlier approach. In the earlier approach the rationale for claiming that the controlling role of the external reality provided an objective underpinning to the way things

are disposed to empirically appear was that, under that control, the structure and organization of the type of world that the empirical evidence suggests (suggests, that is, at the relevantly basic level of empirical enquiry) would, point by point, preserve the structure and organization of the external reality itself, or perhaps the structure and organization that the reality functionally simulates. This, as we saw, would fail to provide the right kind of underpinning from the idealist standpoint, since conformity to the external situation would do nothing to enhance the status of the system of appearance with respect to the idealistic creation. An idealist who pursues the theistic approach in the way we are envisaging will not think of the mechanistic reality as providing, or as helping to provide, the required underpinning in that way. He will see the role of this reality as simply one of securing the appropriate sensory organization, thereby fixing the way things are disposed to empirically appear at the human viewpoint, and he will look exclusively to the role of God, as the causal creator of that reality and ultimate source of that organization, to provide the system of appearance with the underpinning it needs. In other words, he will insist that what objectively underpins the way things are disposed to appear is simply that God puts in place a sensory organization that is designed to dispose them in that way, and he will take the mechanistic reality to be relevant only as the concrete means by which God puts that organization in place.

The other way of pursuing the theistic approach represents God as controlling the course of human sensory experience directly. There is no additional external reality that God puts in place to causally mediate his experiential plans for human subjects. Rather, all our sensory experiences are directly caused by God himself, by immediately efficacious acts of will, and the reason why they conform to a certain world-suggestive pattern, and thereby make it empirically seem to us that we are inhabitants of a certain kind of three-dimensional world, is that God deliberately gives us experiences that are designed to achieve that result. This direct theistic control of our experiences does not, of course, involve the presence of a sensory organization in the terms in which we have been understanding this. There is not, in the ordinary sense, a system of control that governs the course of experience and disposes it to conform to its world-suggestive pattern: the only controlling agency is God, and the pattern is only achieved because of the systematic way in which God selects the types of experience to give us. So it might seem that this account of God's controlling role is not an available option under canonical idealism. But we can plausibly envisage that the method by which

God chooses the types of experience to give us is such as to provide something that we can count as a form of sensory organization for our purposes. God's objective, in choosing the types he does, is to ensure that our experiences are orderly in the relevant world-suggestive way and so succeed in creating the systematic appearance of a world at the human empirical viewpoint. Crucially, we can plausibly suppose that this objective involves a detailed evolving plan of the type of world to which he wants the empirical evidence to conform—a plan which, at any moment, fully covers the core character of the world up to that time, and is only open with respect to the future in so far as he wants to leave himself the freedom to adjust its details in response to our own behavioural decisions. Once we suppose that God's control of our sensory experiences is guided by this plan, we have to accept that, as well as rationally underlying his actual acts of control, the plan will embody ways in which God is rationally committed with respect to his control of experience in hypothetical and counterfactual cases too. For example, if the type of world that features in the plan includes the existence of a pen in my desk drawer over the past hour, then God's commitment to the experiential implementation of the plan will ensure, amongst other things, that were I now to decide to open the drawer and look inside, I would receive the appropriate types of drawer-opening and pen-seeing experiences, and that had I decided to open it ten minutes ago, I would have received those types of experience then. In this way, God's plan will have a rational role, in guiding his volitional acts and in forming his volitional policies, which, in its ultimate impact on the realm of sensory experience, is exactly like that of a sensory organization, as we have been conceiving of it. For by embodying God's commitment to ensuring conformity to the relevant world-suggestive pattern, its impact on this realm will be the same as if it were actually disposing it to conform to that pattern. In view of this, we can appropriately think of what is here envisaged as providing something that will serve as a sensory organization for our purposes, deeming God's planned control of experience to be a special form of organizing control.

An idealist who adopts this second version of the theistic approach would arrive at the position which, as I understand him, Bishop Berkeley was implicitly endorsing, or close to endorsing, in his *Principles of Human Knowledge*. Certainly, Berkeley held that our sensory experiences are directly brought about by divine volition, and, indeed, since he thought that we can make no sense of any kind of causal agency other than the volitional agency of a mental subject, he would have viewed the suggestion that there might be

a controlling reality of a mechanistic kind as unintelligible.⁷ As for his view of the physical world, it is true that some of his remarks suggest a rather crude, and clearly unacceptable, form of idealism, which equates the existence of a physical item with its being perceived. But there are also passages where he seems explicitly to allow that a physical item's existence may sometimes only involve its having a potential to be perceived (by someone in the right place at the right time).⁸ And, as I see it, the overall thrust of his argument points to the view that what secures the existence of the physical world is that God controls our sensory experiences in the relevantly orderly and purposive way—ensuring their conformity to the relevant world-suggestive pattern, and with the objective of securing the systematic appearance of a world at the human empirical viewpoint.⁹ In his *Three Dialogues*, published a few years after the *Principles*, Berkeley seems to offer a quite different account of the nature of God's involvement in the existence of the world.¹⁰ For he seems to be saying that the world exists not as something created by God's organization of human sensory experience, but as an internal object of God's understanding. This position is of no relevance to our own concerns. It is not a form of canonical idealism; and if idealism has to be non-realist, it is not a form of idealism at all, since a world in God's mind would be something whose existence was logically independent of human mentality and philosophically fundamental. I would only add that, if Berkeley really was putting forward this view, his doing so was, from his own standpoint, an aberration. For the view is in conflict with a central component of his basic philosophical outlook, namely his insistence that the world is something that falls within the immediate reach of our perceptual awareness. The centrality of this aspect of his outlook is something that Berkeley underlines at the end of the *Dialogues* themselves.¹¹

There are two very different ways, then, in which a canonical idealist could pursue the theistic approach. On the one hand, he could suppose

⁷ 'When we talk of unthinking agents, or of exciting ideas exclusive of volition, we only amuse ourselves with words' (*Principles of Human Knowledge*, in his *Philosophical Works*, ed. M. Ayers (London: Dent, 1975), sect. 28).

⁸ I particularly have in mind the passages in sects. 3 and 58.

⁹ The thrust in this direction is especially clear in those passages where Berkeley tries to explain why his philosophy does not detract from the reality of the world. See, in particular, sects. 29–36.

¹⁰ *Three Dialogues between Hylas and Philonous*, in his *Philosophical Works*, ed. Ayers. The apparent endorsement of this different account comes in the third of the dialogues.

¹¹ *Ibid.* 207. For a full discussion of Berkeley's various views about the physical world see my 'Berkeley on the physical world', in J. Foster and H. Robinson (eds.), *Essays on Berkeley* (Oxford: Oxford University Press, 1985), 83–108.

that God controls the course of sensory experience indirectly, by causally creating an external mechanistic reality to control our experiences directly and to form the concrete vehicle of the sensory organization. Let us speak of this as the *mechanistic* version of the approach. In its simplest form, this version would involve taking the relevant mechanistic reality to coincide, in structure and organization, with the type of world whose presence is suggested, at the relevantly basic level, by the actual and potential empirical evidence. In other words, it would involve taking the reality to have the structure and organization of the realist's physical world. On the other hand, the idealist could suppose that God controls the course of human sensory experience by direct volition, and in the framework of an evolving plan of the type of world to which he wants the empirical evidence to conform. Let us speak of this as the *Berkeleyian* version. This version, too, recognizes something that we can deem to be a sensory organization, and so think of as disposing sensory experience to conform to its world-suggestive pattern, but this organization will consist not in some mechanistic method of control that God has put in place, but simply in his commitment to giving us experiences that accord with the relevant plan. What is common to both versions—and what makes them ways of trying to implement the externalist strategy—is that they take the divine provenance of the sensory organization to be what gives the system of empirical appearance the right kind of objective underpinning for the purposes of the idealist theory. In other words, they take this provenance to provide a form of underpinning which, when included in the factors that feature in the idealistic creation, secures the requisite objectivity of the item created, so that we can legitimately think of that item as a real world.

This brings us to the crucial question. There are these two ways in which the idealist could take God to be the source of the sensory organization. But what we still do not know is what his rationale would be for claiming that, through the presence of this divine source, the system of appearance acquires the sort of objective underpinning that it idealistically needs. Why should the inclusion of this source in the factors that contribute to the creation of the physical world be thought to endow the world with the requisite objectivity? Whichever way we think of him as controlling human sensory experience, how can we come to see the idealist's God as anything more than the provider of a virtual reality? Why would he be more successful in creating a world than (at least as we would ordinarily interpret the situation) Descartes's malicious demon?

In trying to answer this question, I shall assume that the divine being who features in the theistic approach is one whose nature and status measure up to the full stature of the traditional Judaeo-Christian conception: that he is a personal being of limitless existence and power; perfect in his knowledge, rationality, and goodness; depending on nothing else for his existence or his attributes; and the creator or creative source of all else that concretely exists and obtains. Such a God has sovereignty over the entire realm of concrete reality external to himself, and this sovereignty involves not just the supremacy of his power to determine what is to exist or obtain, but also, given his status as creator and the perfection of his nature, the absolute moral authority to exercise that power, so that everything that falls under his control also falls within the sphere of his moral jurisdiction. This is the kind of being that Bishop Berkeley had in mind when he ascribed direct causal responsibility for our sensory experiences to God; and, for reasons that will now emerge, it is the kind of divine being that the theistic approach needs for a successful outcome, whether developed in the mechanistic or the Berkeleyian way.

The theistic approach represents God not just as deliberately ensuring that human sensory experience conforms to the relevant world-suggestive pattern, but as doing this with the express purpose of creating the systematic appearance of a certain type of world at the human empirical viewpoint, and this is so whether God's control of sensory experience is direct or employs a mechanistic intermediary. This means that, on the theistic approach, the world-suggestive way in which things are disposed to empirically appear is something that God ordains; and it means that, through that ordaining, it receives (as we might put it) God's authorization. Now the fact that something occurs with someone's authorization need not count for much: it need not provide any significant form of underpinning or objective support for what is authorized. But in the case we are envisaging, the one who authorizes is the Judaeo-Christian God, who is sovereign over the entire realm of concrete reality, and whose sovereignty, as I have stressed, gives him not just supremacy of power, but the absolute authority to exercise it. So, in particular, the content of how things are disposed to appear at the human empirical viewpoint comes within the sphere of his moral jurisdiction, as something that he has the indefeasible right to determine. For this reason, the fact that the system of appearance is authorized by God endows it with a normative status. How things stand empirically at the human viewpoint is not just how, as divinely controlled,

they are *made* to stand, but also how, as divinely authorized, they objectively *ought* to stand. Moreover, in authorizing the system of appearance, and securing its objectively normative status, God is implicitly authorizing it as a guide to belief, so that the world-characterizing beliefs that it invites us to acquire become ones that it is objectively appropriate to acquire. This appropriateness is not just appropriateness to the empirical evidence—that, of course, would obtain irrespective of God's involvement—but is an appropriateness that reflects the normative status of the system that embodies the evidence. It is objectively appropriate for us to form and hold the relevant beliefs because the empirical evidence that supports those beliefs is, under God's sovereign jurisdiction, evidence that we are objectively supposed to possess.

It is here, I think, that we can see how the theistic approach equips the idealist to deal with the problem of objectivity. What the idealist is looking for, in pursuing the externalist strategy, is some type of external source for the sensory organization that gives the way in which things are disposed to empirically appear the right kind of objective underpinning—an underpinning which, when included in the factors that feature in the idealistic creation, ensures that what is created has the requisite objectivity, in relation to the human mind, to qualify as a real world. We also know that, if an objective underpinning is to be of the right kind in that way, it must properly harmonize with the role that the system of appearance has in the idealist's theory: it must be an underpinning that is appropriate to the fact that the idealistic role of the system is not to provide information about the external situation, but (as put in place by the sensory organization and the relevant endowments of the mind) to centrally feature in the creation of the world. It was the failure to satisfy this condition that ruled out the earlier method of pursuing the externalist strategy. But, by giving the system of appearance its normative status, as how things objectively ought to be disposed to empirically appear to us, and by thereby securing the objective appropriateness of our believing in the type of world whose presence it empirically suggests, God's authorization of the system would succeed in giving it a kind of objective underpinning that measures up to what is idealistically needed. For this underpinning would not be a matter of securing the conformity of the system to some reality outside it, but would pertain to the status of the system in itself. It would, in this way, be just the kind of underpinning that was equipped to provide direct support for the system's role as world-creative. It would enable the idealist to insist that,

when God's authorization is added to the factors that feature in the idealistic creation, the objective status which that authorization confers on the system gives rise to an objective status in what is created, so that what is created counts as an idealistically real world, rather than as a mere virtual reality. It would be the kind of underpinning which, when taken to be an idealistically relevant factor, would give the idealist a proper basis for claiming that the world-suggestive way in which things are disposed to appear at the human empirical viewpoint is equipped to make it idealistically the case that that is how they are.

Of course, the suggestion that the idealist can solve the problem of objectivity in this way will not seem acceptable to someone whose philosophical intuitions continue to be controlled by a realist conception of the world, and who cannot detach the requirements for its objectivity from a need for mind independence. The kind of objectivity secured for the world by the theistic approach is an objectivity in explicitly idealist terms: it is an objectivity designed for a framework in which it is already assumed that, if there is a physical world, facts about human sensory experience centrally contribute to its existence. But, with realism already discredited, and with the positive case for an idealist understanding of the world already in place, the objectivity secured is, it seems to me, exactly what is dialectically required. It is exactly what the idealist needs in order to make good his position in the face of the apparent difficulty.

It is also important to bear in mind the procedural point I mentioned earlier. An objector to what I am claiming might say that, if we set aside our knowledge of the physical world and simply focus on the fundamental factors that the theistic approach takes to be idealistically relevant, there is nothing that obliges us to characterize these as genuinely world-creative. It is true that these factors suffice to give the system of empirical appearance a form of objective underpinning, and, unlike what was envisaged earlier, this form of underpinning properly harmonizes with the role of the system in the idealist's theory. But it could still be said that, if we confine our attention to these factors, there is nothing that forces us to see them as yielding a real world, rather than a virtual reality that is divinely ordained. Admittedly, granted that the Judaeo-Christian God is morally perfect, it could be replied that the possibility of there being a divinely ordained virtual reality can be excluded, since it would turn God into a deceiver: it would mean that God was systematically encouraging us to acquire beliefs that were false. (This point, of course, was crucial to Descartes's way of trying to defeat scepticism

about the physical world through a proof of the existence of God.¹²) But the objector would see this reply as begging the question. He would say that the idealist needs to show that authorization by the Judaeo-Christian God would enable the system of appearance to play the creative role that the idealist assigns to it before he is entitled to suppose that a morally perfect (and hence non-deceptive) being could be responsible for it.

This is where the procedural point becomes crucial. In order for the theistic approach to deal adequately with the problem of objectivity, it is not necessary that when we set aside our knowledge of the physical world and simply focus on the factors taken to be idealistically relevant, we should find ourselves compelled to recognize these factors as world-creative. All that is required is that our belief in a physical world be tolerant of the supposition that its existence is ultimately sustained in the envisaged way, with its character determined by how things are disposed to empirically appear at the human viewpoint, and with its objectivity stemming from the underpinning of this system of appearance provided by God's authorization. Given that we have already established that, if there is a physical world, its existence has to be sustained in an idealist way, and, given that the underpinning that the system of appearance derives from God is of the right sort to support the system's role as world-creative, it seems to me that this requirement is met. We can retain our ordinary belief in the physical world, characterized in its own spatio-material terms, while accepting that what ultimately underlies its existence are the sensory organization, the relevant endowments of the mind, and the special ordaining role assigned to God.

One final point. It is essential to the success of this solution to the problem of objectivity that it represents God not only as deliberately ensuring the conformity of our experiences to the world-suggestive pattern, but as doing this with the purpose of ensuring the systematic appearance of the relevant type of world at the human empirical viewpoint. It is this that allows us to think of the way things are disposed to empirically appear as something ordained by God, and as thereby receiving his authorization, and it is this authorization that provides the system of appearance with the appropriate form of objective underpinning. But once the idealist has decided to deal with the problem of objectivity in this way, he will inevitably see God's control of our experiences as having a further goal. A God of the Judaeo-Christian

¹² Descartes *Meditations on the First Philosophy*, in *The Philosophical Writings of Descartes*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch, ii (Cambridge: Cambridge University Press, 1984), Sixth Meditation.

type could not fail to know what is required for the idealistic creation of a world, and so could not fail to know that, if he organizes our experiences in the envisaged way and with the envisaged purpose, this will meet the relevant requirements. It is also unthinkable that if God does act in a way that meets these requirements, the meeting of them should turn out to be a consequence that he foresees, but does not intend. So, the idealist will inevitably think of God's ultimate goal in organizing our experiences in the way he does as that of providing us with an idealistically created world. He will think of God not just as deliberately arranging for things to appear systematically worldwise at our viewpoint, but as doing this as a way of ensuring that, at our viewpoint, a reality of the empirically suggested type obtains.

VI

By adopting the theistic approach, the idealist can solve the problem of objectivity. He can take God's authorization to provide a suitable objective underpinning for the system of empirical appearance, and, by then including this underpinning in the factors involved in the idealistic creation, he can secure the requisite objectivity in the item created. I shall also take it that this is the only way in which the idealist can deal adequately with the problem, since I cannot think of any other way in which he could arrange for the system of appearance to have an underpinning of an appropriate kind. I cannot think of anything other than the status conferred by God's authorization that could equip the system to play the world-creative role that the idealist assigns to it.

What we still have to consider are the relative merits of the two versions of the theistic approach, offering alternative accounts of the method by which God controls the course of human sensory experience. On one version—the mechanistic—God is represented as controlling our experiences indirectly, by putting in place an external mechanistic reality to control them directly. On the other version—the Berkeleian—God is represented as controlling our experiences by direct volition. We know that the idealist needs to adopt the theistic approach to deal with the problem of objectivity. But should he adopt it in its mechanistic or in its Berkeleian version? Or should he, perhaps, simply leave the matter open? Berkeley himself, as we noted, would have dismissed the mechanistic version as unintelligible. For he thought that we cannot make sense of anything having a capacity for causal agency other than a mental subject, and that we cannot make sense of any form of causal agency

by subjects other than that of active volition. But I certainly would not want to settle the issue on that basis. It does not seem to me that Berkeley offers any proper justification for his restrictive view of causation, and I do not find the view itself at all plausible. It seems to me that we have a perfectly good general concept of causation and that there is no philosophical difficulty in understanding how that concept can have application to events and processes in non-volitional contexts. In particular, whatever our ultimate verdict on the mechanistic version, I think we have a perfectly good understanding of how God could control the course of sensory experience by means of a mechanistic intermediary in the way envisaged.

In considering the issue of these rival versions of the theistic approach, there are two things that I need to stress in advance. The first thing, and it is already implicit in the previous discussion, is that the difference between them has no bearing on the effectiveness of the theistic approach in dealing with the problem of objectivity. The grounds for claiming that the approach is effective is that God's authorization of the system of appearance, with the objective normative status this confers on it, provides the system with the kind of underpinning it needs for its role in the idealistic creation, and this authorization, and the status it confers, are aspects of the approach on both versions. They are consequences of the fact that, however he arranges for it to obtain, the system of appearance—the way things are disposed to empirically appear at our viewpoint—is something which God ordains. It follows that, if there is any reason for preferring one version of the approach to the other, it must come from a different source. It also means that, whatever interest it may have in its own right, the issue is not, from the standpoint of the idealist project, of crucial importance.

The second thing I need to stress is that some of the reasons why a particular idealist might prefer one version to the other could be to do with his religious beliefs. For example, orthodox Christian belief places great emphasis on God's unconditional love for his human creatures and on the opportunities he provides for them to have a conscious and intimate encounter with that love in their daily lives. An idealist who holds that view of God may well find the Berkeleyan account of God's method of organizing human sensory experience more congenial than the mechanistic. It may strike him as harmonizing better with his view of the pastoral closeness of God to think of God as the direct provider of our sensory experiences, rather than as setting up a causal intermediary. It would enable him to suppose that, in every aspect of our daily experience of the world and our situation within it, we are in direct touch

with God's activity and purposes. Certainly Berkeley himself—in addition to his philosophical objections to non-volitional causation—found his account of God's causal immanence congenial for that reason.¹³ Now there is nothing improper in the idealist's allowing his religious beliefs to influence his decision between the two versions of the theistic approach. But is also clear that, in the context of this discussion, where we are addressing the issues of idealism from a purely philosophical standpoint, religious considerations should not be allowed to carry any weight. In this context, it is merely of passing interest if a particular account of God's causal role in relation to our sensory experiences happens to best suit some specific religious outlook.

Even when we restrict ourselves to strictly philosophical considerations, it might be thought that there is a straightforward reason for preferring the Berkeleyan account of the situation. For if God does not need to employ a causal intermediary in order to be able to control and organize our sensory experiences in the requisite way, it seems, on the face of it, odd if he should choose to do so, rather than take the ontologically simpler course of controlling and organizing them directly. It is not, of course, that there would be any extra effort involved in implementing the ontologically more complex procedure: causally creating and sustaining a mechanistic reality of a suitable kind would be as effortless to the Judaeo-Christian God as bringing about a single experience. But granted that God can achieve his purposes for human experience without a mechanistic intermediary, it would seem pointlessly cumbersome for him to introduce one; and we can surely be confident that a perfectly rational God would not gratuitously add to the realm of reality in that way.

I do not want to deny that the Judaeo-Christian God would not inflate the realm of reality gratuitously, but I do not think that, in the end, this point will help to resolve the issue. In part, this is simply because we do not have access to God's perspective on things, and so cannot exclude his having a reason for choosing the mechanistic method of control that lies beyond our comprehension. But, more importantly, it is also because, within the bounds of what we can comprehend, there is a further factor that could provide God with such a reason. For although the introduction of a mechanistic intermediary would make things ontologically more complex, it would also, developed in any plausible way, make them functionally tidier, since it would subsume the highly complex content of the sensory organization under the

¹³ See, for example, *Principles*, sects. 149–50.

much simpler organizational principles of the mechanistic reality and its links with the human mind. This gain in functional tidiness might, surely, be something that God would value. If he does, then God has both a reason to adopt the Berkeleian procedure and a reason to adopt the mechanistic. Nor, I think, can we make any progress, here, by considering whether God would set a *greater* value on ontological simplicity or functional neatness. I do not see how we could reach a verdict on this.

We cannot, it seems to me, decide between the two versions of the theistic approach by appealing to the ontological simplicity of what is envisaged by the one or to the functional neatness of what is envisaged by the other. As far as I can see, there are only two other considerations which could be thought relevant to settling the issue, and I shall look at them in turn.

In the first place, it could be thought that an idealist wanting to entertain the mechanistic version would be in difficulties over the question of what sort of thing, beyond its structure and organization, the mechanistic external reality might be. It might be claimed, very strongly, that, beyond its structural and organizational properties, the idealist could offer no suggestions at all about the nature of this reality. But this would need qualification. For the idealist has at least the option of supposing that the reality might be something mental. We have already, in the context of our earlier discussion of the inscrutability thesis, identified one way in which such a mentalistic account could be developed, namely by taking the reality to be a three-dimensional spatial one and taking the relevant space to be a concretely existing sense field. And this is not the only mentalistic possibility available.¹⁴ But even when the strong claim is dropped, it could still be insisted that the idealist cannot offer any *plausible* account of the sort of thing the relevant reality might be. For, although there are ways of representing this reality as mental, any mental reality would have to be ontologically composed of mental subjects and the things that take place within them, and it is difficult to think of the Judaeo-Christian God as employing mental subjects in this radically subservient role. It would seem morally inappropriate to the nature of a mental subject, as an idealist would conceive of it, if God were to mechanistically organize its mentality, or some relevant area of it, in a way that was purely designed to have the right kind of causal influence on the

¹⁴ The sense-field account was mentioned in Ch. 3, Sect. II. Another possibility, which develops an idea that featured in the following section, would be to think of the reality as composed of a group of mental subjects whose mentality was organized as if they were occupants of a three-dimensional space.

experiences of some other group of subjects. And if mentalistic accounts of the relevant reality are excluded, there are no others available.

But the question we need to ask is precisely how all this bears on the evaluation of the mechanistic version, and so on the issue between it and its Berkeleian rival. Let us assume, for the sake of argument, that a mentalistic account of the putative mechanistic reality is excluded, so that someone entertaining the mechanistic version would not be able to offer *any* positive account of the sort of thing, beyond its structure and organization, that the relevant reality might be. What problem, if any, does this create for the version? I suppose the worry would be that if we cannot form a conception of any type of non-mental reality that could play the role envisaged, we have no reason to suppose that such a reality is even objectively possible—that the abstract space of possibilities makes provision for a type of thing that is both non-mental and has the structure and organization required. But, while this may create some awkwardness for someone wanting to take his stand on the truth of the mechanistic version, it does not entitle us to conclude that the version is false or even unlikely to be true. For the fact that we have no positive reason to think that a non-mental reality of the relevant sort is objectively possible gives us no reason to think that it is not. It would not be surprising if there were a whole range of objective possibilities beyond the reach of human conception. Indeed, I think we would expect this to be the case. There is nothing here, then, that entitles us to reject the mechanistic version or even to conclude that the Berkeleian alternative is more plausible.

The second consideration that might be thought relevant to settling the issue picks up a point I mentioned when I first set out the basics of the idealist account.¹⁵ The realist takes the physical world to be what ultimately controls the course of human sensory experience. But he will also see it as involved in the functioning of the human mind over a much broader area. This involvement will turn on the special association of human subjects with human organisms, and the functional relationship between the brain of each such organism and the mind of the associated subject, and it will cover both ways in which the brain is involved in controlling what occurs in the subject's mind and ways in which, by the effects of his volitional acts on the brain, the subject is able to perform physical actions. An idealist can recognize this involvement at the mundane level, when he is characterizing how things stand within the psychophysical reality of which the idealistically

¹⁵ In Ch. 3, Sect. VI.

created world forms a part, but he cannot recognize it at the level of what is metaphysically fundamental, as characterized by his idealist theory. But even at this fundamental level, an idealist who postulated an external reality with the same structure and organization as the realist's physical world would equip himself with something which he could take to be functionally linked with human mentality in the same way as this world. And, as I pointed out in the earlier discussion, inasmuch as it would help to provide him with an ultimate explanation of the functioning of the mind, this might be seen as an advantageous outcome.

Once again, however, I do not think that the issue can be settled on this basis. To begin with, although in the framework of realism it is hard to deny that the brain has an extensive role in the ultimate control of human mentality, it is not implausible, from an idealist perspective, to suppose that the mind enjoys a much greater autonomy—that it is, to a much greater extent, equipped by its natural endowments to run its own affairs. Thus, whereas in the framework of realism it is reasonable to suppose that our psychological possession of beliefs is causally underpinned by neural structures in the brain—structures that in some way encode the content of those beliefs—an idealist could well find it natural to think of the human mind as having, at the metaphysically fundamental level, the resources to form and preserve beliefs by itself. And if he does take this view in the case of beliefs, he could equally find it natural to suppose that the mind has its own inherent dispositions to respond to its beliefs with appropriate desires and emotions. Of course, even the idealist cannot think of the mind, at the fundamental level, as entirely autonomous. He is already accepting that the realm of sensory experience is controlled by an external source, and there may well be other forms of mentality where the recognition of an external mode of control is needed. But if so, I see no reason why an idealist of a Berkeleian persuasion should not simply extend the direct controlling role he assigns to God in the case of sensory experience to cover these other areas. He will, in any case, presumably want to recognize God's direct involvement in sustaining our capacity for physical action. For it could only be through God's adjustments to his evolving world plan, in response to our volitional attempts to perform types of physical action, that these attempts could ever be successful. None of this, it seems to me, would create any problem for the Berkeleian.

The considerations that we have looked at do not provide clear-cut support for one of the versions of the theistic approach over the other. Since I cannot think of any other consideration which might enable us to settle the

issue—none, at least, that falls within the scope of our purely philosophical enquiry—I am brought to the conclusion that, on this question, agnosticism is the correct position to adopt. I am slightly disappointed to have to leave the situation open in this way. But I have already stressed that, since both versions are effective in dealing with the problem of objectivity, resolving the issue between them is not crucial for the purposes of the idealist project. (If it were, of course, the issue could be settled in those terms.)

Before leaving this topic completely, there is one final point that I want to make. On both the mechanistic and Berkeleian versions, I have envisaged the possibility of there being some form of externalistic control of forms of human mentality additional to sensory experience. But there is one category of mental events where I think that any externalistic control is logically excluded. The mental events in question are ones of the subject *doing* something, in contrast with events of something *happening to* him. I am thinking of such events as the taking of a decision, the deliberate framing of a mental image, and attempts to perform some kind of physical action. As I have explained in detail elsewhere, I take these events of mental agency to be (as I put it) intrinsically autonomous: their intrinsic natures logically preclude their being brought about by something else.¹⁶ If I am right, then no external control of human mentality can cover events of this sort. Not even God, with his sovereignty over the entire realm of concrete reality, has the power to make us do something in that sense, since it would not be a genuine doing if it were made to happen.

VII

The theistic approach solves the problem of objectivity. But it may now seem that, by employing it, the idealist lays himself open to a further objection. For the majority of current philosophers would insist that belief in the Judaeo-Christian God, or, indeed, in any form of supernatural being, is not rationally defensible. If they are right, then whatever help it would afford him in dealing with the issue of objectivity, the theistic approach is not something that the idealist is entitled to adopt. He is not entitled to adopt a certain solution to a problem if he is not entitled to recognize the entities that the solution postulates.

¹⁶ See my *The Immaterial Self*, 266–80.

The philosophical attack on Judaeo-Christian theism sometimes comes in a very strong form, where the attacker cites some factor which, he alleges, shows that such theism is positively untenable—for instance, by showing that it involves an incoherence, or by showing that it cannot be reconciled with things that we already know. For example, it is sometimes claimed that we cannot reconcile the perfect goodness and unlimited power of God with the amount of suffering in the world. Again, it is sometimes claimed that the notion of unlimited power itself is inherently paradoxical and resists coherent specification. Or again, it is sometimes claimed that we cannot think of God as omniscient without—unacceptably—having to give up our view of ourselves as free agents. Even the very idea of a supernatural personal being is likely to seem unintelligible to those philosophers who accept a physicalistic or reductive account of human mentality. With respect to this last case, I have already explained why I think that the idealist is entitled to his Cartesian view of the human mind, which takes the mental subjects involved to be wholly non-physical, and takes their mentality, at the level of what is metaphysically fundamental, to be logically independent of the physical world. And from the standpoint of this Cartesian view, the problems that a physicalistic or mentally reductive philosopher would have with the notion of a supernatural personal being would not arise.¹⁷ As for the more specific ways in which Judaeo-Christian theism may be thought to be untenable—ways that relate to alleged difficulties in supposing that there is a God of a distinctively Judaeo-Christian type—these are not issues that it would be appropriate to discuss here: dealing with them adequately would require a large-scale work in the philosophy of religion. All I can do is record my confidence that the putative difficulties can be overcome, and indicate that I am content, in this context, to proceed on that assumption.¹⁸

But even when we assume that there is nothing which positively excludes the theistic position, there is still the complaint that there are no good grounds for accepting it, and that without such grounds its acceptance would be irrational. There have been, of course, over the centuries, a number of attempts by theistic philosophers to demonstrate the credibility of their position. There have been attempts to show that we are rationally obliged

¹⁷ This is why I remarked earlier (p. 203) that the idealist's view of the mind would be crucial to the solution of the problem of objectivity that I would eventually adopt.

¹⁸ For a detailed discussion of the issues, and one that helps to eliminate many of the supposed problems for the theist, see Richard Swinburne, *The Coherence of Theism* (Oxford: Oxford University Press, 1977), *The Existence of God* (2nd edn., Oxford: Oxford University Press, 1991), and *The Christian God* (Oxford: Oxford University Press, 1994).

to accept the existence of God in order to have a satisfactory explanation of why there is a concrete reality at all. There have been attempts to show that belief in God is rational because it provides the most plausible explanation of certain specific aspects of the concrete reality that obtains, such as its orderly character or its inclusion of conscious rational beings like ourselves. There have been attempts to show that we need to accept the existence of a God who imposes requirements on human conduct in order to secure a metaphysically adequate source for moral obligation. There have even been attempts to show that the existence of God is logically guaranteed by our very conception of his nature. But none of these attempts has found much favour among modern philosophers; and in the eyes of many such philosophers the whole idea of trying to reach well-grounded conclusions about a supernatural realm—a realm that is, *ex hypothesi*, beyond the reach of empirical investigation in the ordinary sense—is misconceived. My own view is that, once we accept the intelligibility of theism, some of the arguments that have been developed in its support are reasonably strong ones (or can be made so with a degree of fine-tuning), and that their strength would be more generally acknowledged if philosophers were prepared to approach the issue in an open-minded way. But, again, this would form a major topic for separate discussion, and is not something that I can appropriately deal with in the present context.

The point I want to make, however, is that, although it would be reassuring for the idealist if he could gain support for his commitment to theism from independent pro-theist arguments, the availability of such support is not crucial to the success of his enterprise. He needs to accept the existence of the Judaeo-Christian God in order to make good his idealist account. But he does not need to offer a justification of theism beyond what it receives in the context of this account. To understand why, we need to begin by reminding ourselves of the epistemological framework in which the idealist is operating.

As we saw in Chapter 5, the idealist cannot hope to employ his idealist account of the physical world as a way of taking up and meeting the challenge of the radical sceptic.¹⁹ He cannot start, as the sceptic demands, by agreeing to call the existence of the physical world into question, and then, as a way of putting belief in its existence on a secure foundation, appeal to our knowledge of the facts that he takes to idealistically create it. For unless we can rely on our prior physical knowledge, we have no way of establishing,

¹⁹ See Ch 5, Sects. II–III.

or reaching well-founded conclusions about, the facts that are idealistically involved; we do not even have a way of gauging the relevant facts about human sensory experience. The only way in which the idealist (or anyone else) can hope to respond effectively to the sceptic is by dismissing his challenge as misconceived and taking his philosophical stand on the epistemological outlook of common sense. In other words, his only hope is to insist that we have direct epistemic access to the physical world through perception, and that we are entitled to take this access for granted, without having to start from a position where the existence of the world is in doubt. Since that earlier discussion I have been working on the assumption that the idealist does take his stand on the common-sense position in this way, and that he is right to do so. Of course, the position is only available to him if there is a way of understanding how epistemic access to the world through perception can be achieved in idealist terms. That is why I went on to give an account of the general form that such an idealistically realized access would take. In contrast, as we saw, the common-sense outlook is not available to the realist, since his position puts the world beyond the reach of perceptual awareness.

The framework in which the idealist is now working, then, and in my view is justified in working, is one in which he is taking for granted the existence of the physical world, and our access to it through perception, and is trying to show that his idealist account of it can be defended against objections. In particular, he is trying to show that, by adopting the theistic approach, he can ensure that what gets idealistically created has the requisite objectivity, in relation to the human mind, to qualify as a genuine world, rather than as just a virtual reality. Of course, if it turns out that he cannot solve the problem of objectivity in this way, or that his account is flawed in some other respect, then the idealist position will have to be rejected; and since realism has already been refuted, and an idealist account of the world is now the only one available, any rejection of idealism would necessarily lead to physical nihilism. But the point remains that the existence of the physical world is not something that the idealist needs to call in question in the process of trying to show that his position can be developed in a satisfactory way. The crucial consequence of this is that, so long as there is nothing that positively excludes Judaeo-Christian theism—and I am working on the assumption that there is not—and so long as canonical idealism is not flawed in some other way, the very fact that the idealist can only solve the problem of objectivity by adopting the theistic approach gives him all the justification that he needs in

his adoption of that approach. He does not have to appeal to an argument for the truth of theism before he can legitimately incorporate it into his idealist theory: he can simply appeal to the fact that, if there is a physical world, it has to be construed in an idealist way—a fact that has already been established by our earlier arguments—and then point out that there is no satisfactory construal which does not involve assigning a crucial role to the actions and purposes of God. In the framework of the assumption that there is a physical world, the recognition of these two facts is all that is required to establish the existence of God.

There is still one residual matter, here, that I need to deal with. I am insisting that the idealist is right to take his philosophical stand on the epistemological outlook of common sense, and so I am endorsing the claim that, even at the level of philosophical reflection, we are rationally entitled to take the existence and accessibility of the world for granted, without having to start from an initial position of doubt. What I have not yet explained—and this was an issue that I explicitly put on one side in the earlier discussion—is why I think we have this entitlement. Why do I think that we are justified, or at least justified as idealists, in dismissing the sceptical challenge in that way? It is obviously not enough, here, simply to point out that, unless we can rely on such an entitlement, we do not have any way of defeating the sceptic. The fact that scepticism cannot be rationally avoided without a reliance on the outlook of common sense gives us a reason for wanting that outlook to be correct, but does not prove that it is.

The reason why I think that we can properly take our philosophical stand on this outlook is, in essence, very simple. In our everyday lives it seems completely clear to us—and clear without our having to employ some justifying argument—that we are perceptive inhabitants of a physical world, and that we have direct epistemic access to that world through perception. Trivially, within its own context, this cognitive seeming is rationally authoritative and entitles us to accept the thing that seems to be the case. But what I also take to be true is that, where, in the context of our ordinary thinking, something cognitively seems completely clear, so that it is cognitively represented to us as a fact beyond dispute, then, even on philosophical reflection, this is something that we are rationally entitled to accept unless we can identify some positive reason why the veridicality of the seeming should be doubted. And in the case of how things ordinarily seem with respect to the existence and accessibility of the physical world, I think that no positive reason for doubt can be found. Of course, the

fact that something seems completely clear is not a logical guarantee that it obtains. Even from an idealist standpoint it is not difficult to think of possible scenarios in which things relevantly cognitively seem the way they do, but in which there is no physical world: just envisage a situation in which our sensory experiences conform to their world-suggestive pattern by mere chance, or in which there is no suitable objective underpinning for how things are disposed to empirically appear. But knowing that the veridicality of how things ordinarily cognitively seem is not logically guaranteed does not give us a positive reason for thinking it doubtful, and so it does not, as I see it, affect our rational entitlement to endorse the content of this seeming at the level of philosophical reflection. It does not even, I think, mean that this endorsement should fall short of full conviction: having the right to be sure of something does not entail being in a psychological state that logically guarantees the truth of the proposition in question.

In short, given that it ordinarily seems completely clear to us that we are perceptive inhabitants of a physical world, I take the onus to be on the sceptic to develop a positive case for distrusting the content of this seeming, and thereby to cast doubt on the existence of the world, and I cannot think of any way in which such a case could be developed. At least, I cannot think of any such a way within the framework of an idealist account of what the existence of a physical world would ultimately amount to. Under physical realism, of course, the outcome would be different: a realist view would prevent our endorsing the way things ordinarily seem, because it excludes the possibility of physical-item perception. I have not, admittedly, *proved* that the onus rests on the sceptic in this way; nor am I in a position to do so. All I have done, and all I am capable of doing on this front, is to identify the considerations which lead me to take the view of our epistemological situation that I do. If someone does not find these considerations compelling, there is nothing more I can say in my defence.

I conclude, then, that, with respect to the epistemological issue of the physical world, the idealist is justified in taking his philosophical stand on the outlook of common sense and dismissing the sceptical challenge as misconceived. And, with the assumption of the existence of the world in place, and the need for an idealist account of it already established, he is justified in developing this account in the one form that allows a satisfactory outcome—a form that secures the objectivity of the idealistically created world through the special role assigned to the Judaeo-Christian God.

VIII

Physical realism failed because it did not give the world the requisite empirical immanence. It left the world as something to which we have no perceptual access, and it did not accord the way things are disposed to empirically appear the right kind of authority with respect to how they physically are. Canonical idealism had no difficulty in securing the empirical immanence of the world, but faced a problem over its objectivity. Given that such idealism assigns the central role in the creation of the world to the sensory organization, and given that the organization plays this role by the way in which it disposes things to appear at the human empirical viewpoint, the problem was in seeing how what was idealistically created could count as anything more than a virtual reality—how it could have the right sort of objectivity in relation to the human mind to count as a real world.

I have argued that the canonical idealist can overcome this problem by assigning an appropriate role in the idealistic creation to the Judaeo-Christian God—a God who has sovereignty over the entire realm of concrete reality external to himself, and whose sovereignty involves not just the supremacy of his power, but also the absolute moral authority to exercise that power. The idealist must continue to assign the same central creative role to the sensory organization and the way in which, in the context of the relevant endowments of the human mind, it disposes things to appear systematically worldwide at the human empirical viewpoint. It is this that ensures that the idealistically created world is, in the required way, empirically immanent. But by taking the sensory organization to be put in place by God, and by taking the system of appearance it generates to be something which God expressly ordains, he can give this system a form of objective underpinning—the underpinning of divine authorization. And by then taking this underpinning to be an essential factor in the creation of the world, to be added to the sensory organization and the relevant endowments of the mind, he can secure the requisite objectivity in what is created: he can ensure that what gets created is not only ours empirically, but something that we can count as a genuine world. In this way the idealist can provide a theory of the world which meets what is required in all respects. As we saw, there are two very different ways in which the idealist could think of God as responsible for the sensory organization, one of which represents him as controlling the course of sensory experience directly, the other of which

represents him as exercising control through a mechanistic intermediary. Which method of divine control the idealist envisages makes no difference to the effectiveness of God's role in the context of the idealistic creation, and, for reasons which I explained, I have left the issue between them open.

My conclusion, then, is that the world is something whose existence is constitutively sustained by three factors: first, the sensory organization, which disposes sensory experience to conform to its world-suggestive pattern; second, the relevant endowments of the human mind, which make us empirically receptive to the orderly character of our sensory experiences, and thereby enable the sensory organization to dispose things to appear systematically worldwise at our empirical viewpoint; third, the ordaining role of God, which provides the right kind of objective underpinning for the way in which things are disposed to empirically appear. The first two factors, in combination, ensure the requisite empirical immanence of the world. The third secures its requisite objectivity. No other philosophical account of the physical world, it seems to me, can achieve this outcome. No other way of developing canonical idealism will yield an idealistically created world that we can legitimately think of as objectively real. And, if my argument against physical realism has been successful, no other account at all can make provision for a real world that is, in the required way, empirically ours—for a real world that forms an empirical world for us.

In describing the idealistically created world as a world *for us*, my main emphasis, throughout the discussion, has been on its empirical immanence. And this is natural, since the empirical immanence of the physical world is the factor which realism cannot accommodate, and which—to achieve it—obliges us to accept the idealist alternative. But the sense in which the created world is a world for us has two further aspects, and these aspects, too, deserve to be underlined.

The first aspect is one which I have already said something about in my initial exposition of the idealist position.²⁰ It is that what is created is a single world for all human subjects. An idealism that recognized different worlds for different subjects, or a single world that was necessarily restricted to a single subject (to oneself in the role of the proponent of the idealism), is something I would regard as tantamount to physical nihilism: I would not see it as preserving the existence of a physical world in any recognizable

²⁰ In Ch. 3, Sect. VI.

sense. As I pointed out in the earlier discussion, if the idealist is to ensure that what gets idealistically created forms a single world for all human subjects, he has to think of the sensory organization as an integrated system of control, which not only imposes suitably harmonious constraints on the sensory experiences of different subjects, but does this non-accidentally—as a result of something essential to the method of control it embodies. It is clear how a sensory organization with the divine provenance envisaged under the theistic approach would meet this condition, since the content of the organization would, at all points, reflect the guiding role of God's purposes, and these purposes would include the fixed objective of providing the relevant intersubjective harmony.

The second aspect is one I have not yet touched on, since it has not been directly relevant to the issues I have had to deal with. Identifying it now will enable me to complete the idealist picture.

In speaking of the created world as a world for *us*, I am signalling the multiplicity of the subjects involved, but I am also indicating that they form a community. It is not just that there is a single world for all human subjects taken individually, but that there is a common world which these subjects, as an integrated group, share. Now the recognition that human subjects form a community is something philosophically uncontroversial: it is not a point over which idealism and realism differ. But what is crucial is that we should understand how the communal nature of our situation fits into the idealist account. Specifically, what needs to be recognized is that we do not form a community, in any significant sense, at the level of what is metaphysically fundamental. It is true that we can envisage ways in which, at this fundamental level, there is provision for causal traffic between different subjects. Certainly, we can envisage such provision on the mechanistic version of the theistic approach, where we can suppose that subjects are functionally linked with the mechanistic external reality in the way in which, under realism, they are linked with the physical world, and so can suppose that they are equipped to affect one another via their effects on this reality. But none of this would, in itself, give the subjects involved anything that we can properly think of as a communal life. The communal life that we have, and are conscious of having, is only available to us in the context of the created world. All the multiplicity of ways in which we socially converse and interact depend on the presence of a physical world that we jointly inhabit, and on our recognition of the presence of this world and the presence of other subjects whom we can identify by reference to the human organisms with which they are associated.

In other words, our shared life as a community is something which, and which only, characterizes our situation mundanely, in terms of how we are placed within the psychophysical reality that is idealistically sustained. This does not, of course, detract from the reality of our communal life: it does not mean that our communal life is in some sense an illusion. At least, it does not do so if I have been successful in my defence of the reality of the idealist's world.

Given the depth of the personal relationships which our communal life enables us to form, we might well find ourselves with aspirations for a communal life which is not world-bound in this way and which holds at the level of what is metaphysically fundamental. And if, as many believe, there is a form of life beyond death, perhaps, in that new context, this aspiration will be achieved. Certainly, as I understand it, this is the expectation of Christian theism. But I have already made it clear that issues of this sort lie outside the scope of my philosophical enquiry.

This page intentionally left blank

Bibliography

- ARMSTRONG, D., *A Materialist Theory of the Mind* (London: Routledge & Kegan Paul, 1968).
- AYER, A. J., *Language, Truth, and Logic* (2nd edn., London: Gollancz, 1946).
- BENNETT, J., 'Substance, reality, and primary qualities', *American Philosophical Quarterly*, 2 (1965), 1–17.
- BERKELEY, G., *Three Dialogues between Hylas and Philonous*, in his *Philosophical Works*, ed. M. Ayers (London: Dent, 1975).
- *A Treatise Concerning the Principles of Human Knowledge*, in his *Philosophical Works*, ed. M. Ayers (London: Dent, 1975).
- CHISHOLM, R., *Perceiving* (Ithaca, NY: Cornell University Press, 1957).
- DESCARTES, R., *Meditations on the First Philosophy*, in *The Philosophical Writings of Descartes*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch, ii (Cambridge: Cambridge University Press, 1984).
- *The Passions of the Soul*, in *The Philosophical Writings of Descartes*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch, i. (Cambridge: Cambridge University Press, 1984).
- DUCASSE, C., 'Moore's refutation of idealism', in P. Schilpp (ed.), *The Philosophy of G. E. Moore* (Chicago, Ill.: Northwestern University Press, 1942), 223–51.
- FOSTER, J., *Ayer* (London: Routledge & Kegan Paul, 1985).
- 'Berkeley on the physical world', in J. Foster and H. Robinson (eds.), *Essays on Berkeley* (Oxford: Oxford University Press, 1985), 83–108.
- 'A brief defence of the Cartesian view', in K. Corcoran (ed.), *Soul, Body, and Survival* (New York: Cornell University Press, 2001), 15–29.
- *The Case for Idealism* (London: Routledge & Kegan Paul, 1982).
- 'In defence of phenomenalist idealism', *Philosophy and Phenomenological Research*, 54 (1994), 509–29.
- *The Divine Lawmaker* (Oxford: Oxford University Press, 2004).
- *The Immaterial Self* (London: Routledge, 1991).
- *The Nature of Perception* (Oxford: Oxford University Press, 2000).
- 'The succinct case for idealism', in H. Robinson (ed.), *Objections to Physicalism* (Oxford: Oxford University Press, 1993), 293–313.
- HARRÉ, R., and MADDEN, E., *Causal Powers* (Oxford: Blackwell, 1975).
- HINTON, J., *Experiences* (Oxford: Oxford University Press, 1973).
- HUME, D., *A Treatise of Human Nature*, ed. L. Selby-Bigge, 2nd edn. rev. P. Nidditch (Oxford: Oxford University Press, 1978).
- KRIPKE, S., *Naming and Necessity* (Oxford: Blackwell, 1980).

- LOCKE, J., *An Essay Concerning Human Understanding*, ed. A. Campbell Fraser (New York: Dover, 1959).
- MARTIN, M., 'The limits of self-awareness', *Philosophical Studies*, 120 (2004), 37–89.
- 'The transparency of experience', *Mind and Language*, 17/4 (2002), 376–425.
- McDOWELL, J., 'Criteria, defeasibility, and knowledge', *Proceedings of the British Academy*, 68 (1982), 455–79.
- PITCHER, G., *A Theory of Perception* (Princeton, NJ: Princeton University Press, 1971).
- ROBINSON, H., 'The general form of the argument for Berkeleian idealism', in J. Foster and H. Robinson (eds.), *Essays on Berkeley* (Oxford: Oxford University Press, 1985), 163–86.
- *Perception* (London, Routledge, 1994).
- SNOWDON, P., 'The objects of perceptual experience', *Proceedings of the Aristotelian Society*, supp. 64 (1990), 121–50.
- 'Perception, vision, and causation', *Proceedings of the Aristotelian Society*, 81 (1980–1), 175–92.
- SWINBURNE, R., *The Christian God* (Oxford: Oxford University Press, 1994).
- *The Coherence of Theism* (Oxford: Oxford University Press, 1977).
- *The Existence of God* (2nd edn., Oxford: Oxford University Press, 1991).

Index

- adverbialism, with respect to perception 29
n. 12
- analytical (conceptual) reduction, of the
physical to the non-physical 105–7
- appropriateness requirement 22–6
- argument from hallucination 12–14
- Armstrong, D. 29 n. 11
- Ayer, A. J. 106 n.
- Bennett, J. 50
- Berkeleyan version, of the theistic
approach 222–5, 230–6, 242–3
- Berkeley, G. viii–ix, 40–1, 160 n., 223–5,
226, 230–6
- bivalence 197–8
- Boscovich, R. 66 n.
- canonical idealism
central argument for 123–63
defended against objections 164–241
exposition of 107–22, 123–4, 164–245
- Cartesian view, of the mind 179–80,
181–2, 203–9, 237
- Chisholm, R. 29 n. 12
- conceptual (analytical) reduction, of the
physical to the non-physical 105–7
- constitution, relation of 1–6, 212–14; *and
passim*
- constitutive (metaphysical) reduction, of the
physical to the non-physical 40, 43,
70–1, 87–8, 96–7, 105–7; *and then
passim*
- content property, defined 46; *see also*
intrinsic content
- core physical reality, defined 103
- decompositional view, of
perception 27–36, 36–7, 39–40, 42,
205–6
- Descartes, R. 140–1, 179, 207, 225,
228–9; *see also* Cartesian view
- Ducasse, C. 29 n. 12
- empirical immanence, of world 135–48,
151, 164–5, 167–8, 199, 202,
211–12, 215, 242–3
- empirical projection 137–157, 161–2,
215–6
- Euclidean geometry 73–5, 79, 81, 86–7,
92
- excluded middle, law of 189–93,
196–8
- experiential interpretation 28–9, 109–10,
175
- external reality
and the inscrutability of intrinsic
content 83–101, 119–22,
124–5
its relationship to the physical
reality 39–41, 83–163, 209–30
its role in canonical idealism 115–18,
209–36
- externalism, with respect to forms of
mentality 205–6
- externalist strategy, with respect to the issue
of objectivity 209–36
- Faraday, M. 66 n.
- fundamentalist claim, as a component of
physical realism 43, 97, 100–1,
105–7; *and then passim*
- fundamentalist view, of perception 6–27,
36–7. 39, 205–6
- general theory of relativity 69, 74
- geometrical structure 69–70, 73–80,
81–2, 83–7, 89–93, 99, 100,
119–22, 124, 126–44, 146–7, 148,
150–2, 154, 158, 160
- God
existence of 236–41
nature of vii–viii, 226–30, 231–4,
237
role of in canonical idealism 220–44
- grativity 69, 74, 182

- hallucination 12–14, 16–17, 19, 20, 27
- Harre, R. 66
- Hinton, J. 13 n. 5
- human mind, nature of 28, 104–5, 179–80, 203–9, 237
- Hume, D. 111, 204, 207
- idealism *see* phenomenalist idealism
- immanence, empirical *see* empirical immanence
- inclusive external organization, defined 152
- independence claim, as a component of physical realism 42–3, 100–5; *and then passim*
- indeterminacy, physical 189–98
- individuation, of parcels of matter 57–9
- inscrutability thesis, *see* intrinsic content
- interchange case, specified 127–35
- internalism, with respect to forms of mentality 205–6
- internalist view, in the context of the fundamentalist view of perception 20–6
- interpretation, as an element in perceptual experience 28–9, 109–10, 175
- intrinsic autonomy, of events of mental agency
- intrinsic content 42–101, 119–22, 124–5, 188
- intrinsic property (intrinsic character), defined 47–8
- Judaeo-Christian God, *see* God
- Kant, I 66 n.
- knowledge, of physical world *see* physical knowledge
- Kripke, S. 3
- Locke, J. 31 n. 14, 52–5, 56–60, 64, 85, 92
- logical necessitation 3–6
- McDowell, J. 13 n. 5
- Madden, E. 66
- Martin, M. 13 n. 5
- material objects 43–72; *and passim*
types of property of 44–8
- mechanistic version, of the theistic approach 221–2, 224–5, 230–6, 242–3, 244
- metaphysical (constitutive) reduction, of the physical to the non-physical 40, 43, 70–1, 87–8, 96–7, 105–7; *and then passim*
- mind, nature of 28, 89–90, 104–5, 179–80, 203–9, 237
- mundane framework, defined 186
- nihilism, physical 166–72, 177–8, 239, 243–4
- non-veridicality, of perception 19–21, 22–5, 176
- objectivity of the world 166, 167–8, 178, 187, 199–243
- ontological independence, of world from human mind 1, 40–1, 42–3, 100–5, 107, 123, 124, 125–6, 139–40, 165–6, 177–87, 188–9, 199–201
- opaque conception, defined 54
- orderliness of sensory experience 108–118; *and then passim*
- other minds, knowledge of 170
- perception 1–41, 42–4, 100–1, 124, 173–7, 188–9, 206, 238–41
- perceptual access, to world 1, 30–41, 42–4, 100–1, 174–5, 188–9, 238–41
- perceptual experience, nature of 14–30, 109–10, 175–6
- perceptual mediation 7, 9, 32–6
- phenol-thio-urea 50
- phenomenal content 14–28, 30, 36–7, 38; *see also* sensible appearance
- phenomenalist idealism 40–1, 42–3, 100–1, 107–22; *and then passim*
canonical version of, *see* canonical idealism
- ϕ -terminality, defined 7
- physical indeterminacy 189–98

- physical knowledge 38–9, 42–101,
 119–22, 124–5, 137–44, 166–7,
 169–77, 238–41
 as a basis for knowledge of the experiential
 realm 169–73, 238–9
 and the issue of intrinsic
 content 43–101, 119–22, 124–5,
 188
 and the outlook of common sense 38–9,
 174–7, 238–41
 and radical scepticism 38–9, 43–4,
 140–1, 166–7, 173–5, 177,
 238–41
 physical nihilism 166–72, 177–8, 239,
 243–4
 physical particles, nature of 56–72, 87–8,
 97, 121–2
 physical realism 43, 48, 69, 72, 97, 98, 99,
 100–1, 101–7; *and then passim*
 claims of 43, 100, 101–7
 and the problem of perception 36–41,
 42–3, 100–1, 174–5, 188–9
 and the problems of
 inscrutability 83–101, 119–22,
 124–5
 refutation of 123–59
 physical space 69–70, 72–80, 81–2, 94–9,
 99–100, 126–44, 146–7, 148,
 150–2, 154, 158–9, 160
 physical world
 its empirical immanence 135–48, 151,
 164–5, 167–8, 199, 202, 211–12,
 215, 242–3
 idealist view of 39–41, 42–3, 100–1,
 107–22, 123–5, 159–63,
 164–245
 knowledge of *see* physical knowledge
 and the issue of intrinsic
 content 43–101, 119–22, 124–5,
 188
 its objectivity in relation to the human
 mind 166, 167–8, 178, 187,
 199–243
 perceptual access to 1, 30–41, 42–4,
 100–1, 174–5, 188–9, 238–41
 realist view of 43, 100–1, 101–7,
 123–59, 164–5
 physicalism 28, 89, 104–5, 139, 179–80,
 207–9, 237
 Pitcher, G. 29 n. 11
 powers view, of physical particles 65–72
 presentation (presentational awareness),
 defined 17–18
 presentational view, of ϕ -terminal
 perceiving 17–20, 21, 23, 24, 26
 Priestley, J. 66 n.
 principle of individuation 57–9
 projection, on to our empirical
 viewpoint 137–157, 161–2,
 215–6
 psychophysical causation 102–3, 104,
 179–87, 199–201, 210, 234–5
 realism, physical *see* physical realism
 relativity, general theory of 69, 74
 reduction, of the physical realm 43, 105–7,
 123, 127, 168, 173–4, 187, 205–6
 distinction between conceptual
 (analytical) and constitutive
 (metaphysical) forms of 105–7
see also phenomenalistic idealism;
 canonical idealism
 regularity 110–12, 209–10
 relevant deviance, defined 145–6
 relevant external reality, defined 126–7
 Robinson, H. viii, 13 n. 6
 scepticism, about the physical world 38–9,
 43–4, 140–1, 166–7, 173–5, 177,
 228–9, 238–41
 secondary qualities 48–55, 55–6, 61,
 77–8, 81, 84–5, 92, 120–1, 188
 sense-datum theory 27–30, 31–2, 33, 35,
 110, 175
 sensible appearance 14–27, 38, 48–55,
 84–5, 120–1, 124, 145, 188
 sensory experience
 orderliness of 108–18; *and then passim*
 organization of 112–18; *and then passim*
see also perception; perceptual experience
 sensory organization, defined 112
 Snowdon, P. 13 n. 5
 solidity 56–61, 64
 space 69–70, 72–80, 81–2, 94–9,
 99–100, 126–44, 146–7, 148,
 150–2, 154, 158–9, 160
 Swinburne, R. 237 n. 18

- theism, Judaeo-Christian *see* God
- theistic approach 220–44
- topic-neutrality, of knowledge of the external reality 84–100, 119–21, 124–5
- transcendental framework, defined 186
- transparent conception, defined 54
- vagueness 25–6, 196
- veridicality, of perception 19–21, 22–5, 171, 176
- volitional causation 223–4, 230–1
- Wittgenstein, L. 138
- world *see* physical world
- zero deviance, defined 149