

OXFORD

the possibility of knowledge

Quassim Cassam



THE POSSIBILITY OF KNOWLEDGE

This page intentionally left blank

The Possibility of Knowledge

QUASSIM CASSAM

CLARENDON PRESS · OXFORD

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

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

Oxford New York

Auckland 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

© Quassim Cassam 2007

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

First published 2007

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
Data available

Typeset by Laserwords Private Limited, Chennai, India
Printed in Great Britain
on acid-free paper by
Biddles Ltd., King's Lynn, Norfolk

ISBN 978-0-19-920831-9

1 3 5 7 9 10 8 6 4 2

Preface

This book is about how-possible questions in epistemology. A how-possible question asks how something is, or was, possible. Such questions aren't necessarily philosophical. Students of British politics might wonder how it was possible for John Major to become Prime Minister in 1990 but this is a question for historians and political scientists rather than for philosophers. The how-possible questions that are of philosophical interest are metaphysical, ethical, or epistemological. So, for example, a philosopher might ask how freedom of the will is possible or how evil is possible or how knowledge is possible. The latter is an example of an epistemological how-possible question. The following chapters are about how this question arises and what a good answer to it would look like.

The basic idea is very simple. We start by assuming that knowledge is possible but then come across apparent obstacles to its existence or acquisition. So the question is: how is knowledge, or knowledge of some specific kind, possible given the various factors that make it look impossible? On this account, how-possible questions are obstacle-dependent. Sceptics are people who think that the obstacles to knowledge are insuperable and that knowledge is therefore impossible. If we don't want to end up as sceptics we will have to show that the alleged obstacles are unreal or that they can be overcome. Either way, it is the perceived obstacle that gives the how-possible question its bite.

What would be an example of an obstacle to knowledge? It's easier to get a sense of the problem in relation to particular kinds of knowledge, say knowledge of the external world. This kind of knowledge isn't possible for us unless we have ways of knowing something about the external world, that is, viable sources of knowledge of the world around us. The absence of such sources would therefore be one obstacle. If this is the obstacle that triggers the how-possible question then it looks as though an effective response is going to have to operate at several different levels: at one level, it is going to have to show that we do have at our disposal genuine ways or means of knowing about the external world. Second, any apparent obstacles to the acquisition of knowledge by the proposed means will need to be dealt with. Third, an account

might be given of what makes it possible for us to come to know things about the world around us by these suggested means.

A response to a how-possible question that operates on all these levels is what I call a multi-levels response. I defend this approach to epistemological how-possible questions in Chapter 1. In Chapter 2 I try to explain why transcendental arguments, arguments that focus on necessary conditions of knowledge or experience, don't explain how knowledge is possible. In the remaining chapters, I put the multi-levels approach to work in dealing with three specific how-possible questions: how is perceptual knowledge possible? (Chapters 3 and 4); how is knowledge of other minds possible? (Chapter 5); and how is a priori knowledge possible? (Chapter 6). Perhaps it's worth adding that talk of the different 'levels' of a multi-levels response shouldn't be taken too literally. It's more a matter of a satisfactory response to a how-possible question having to do several different and interconnected things in the course of a single evolving enquiry.

When I first started to work on this topic I assumed that I would have to wade through a large body of literature before trying to come up with my own view. Thankfully I was mistaken. Despite the fact that 'how is knowledge possible?' is often represented as one of the defining questions of epistemology little has been written on the nature of how-possible questions as such. A notable exception is William Dray's *Laws and Explanation in History*, published in 1957. Although Dray focuses on how-possible questions in the philosophy of history his ideas have been taken up and given an epistemological twist by Nozick and Stroud. McDowell also talks about how-possible questions in *Mind and World*. But as far as I know that's pretty much as far as the philosophical literature on how-possible questions goes.

What makes this all the more surprising is that Kant is the patron saint of how-possible questions in philosophy, and he is not usually someone whose views on any given topic fail to get the attention they deserve. The central question of the *Critique of Pure Reason* is 'how is synthetic a priori knowledge possible?' yet few of the many commentaries on the first *Critique* have much to say about Kant's conception of questions of this general form. Instead they tend to assume that he is really interested in necessary conditions of the possibility of experience and that he is somehow in a position to explain how knowledge is possible by identifying such conditions. This is a mistake. Even if, as Kant supposes, experience is the same thing as empirical knowledge, figuring out what is necessary for experience isn't a good way of explaining how empirical

knowledge is possible let alone how synthetic a priori knowledge is possible. There is more on the irrelevance of transcendental arguments in Chapter 2.

It took me a while to see that what is often taken to be the Kantian approach to how-possible questions is mistaken. Indeed, this began as a book on transcendental epistemology, that is, on the project of coming up with and explaining necessary conditions for experience or knowledge in general. The focus shifted as I thought more about the relationship between this project and the project of explaining how knowledge is possible. Explaining how knowledge is possible is fundamentally a matter of figuring out how to get it, and the fact is that there are lots of different ways of getting knowledge of the world around us. Recognizing the existence of a variety of what Alvin Goldman calls 'pathways to knowledge' has a liberating effect and also opens up the possibility of a version of transcendental epistemology that is different from the standard version. Instead of thinking about necessary conditions of the possibility of experience or knowledge 'in general' we can think about necessary conditions for the acquisition of knowledge by specific means. Different means of coming to know things about the external world might have different necessary conditions, and it is the identification of means of knowing that does the crucial work in explaining how knowledge of the external world is possible.

Before I realized that I wanted to write a book on how-possible questions my research was supported by a grant from the Arts and Humanities Research Board under its special leave scheme. I'm grateful to the AHRB for its support. The first time I talked at length about how-possible questions was in 2004 in a graduate class at Northwestern University. Thanks to Northwestern for its hospitality and to those attending the class for some useful feedback. I also profited from discussions of the first three chapters at the Instituto de Investigaciones Filosóficas in Mexico City. In 2006 the first two chapters were given a going over, in the nicest possible way, at the NYU Mind and Language Seminar. I thank Béatrice Longuenesse and Don Garrett for arranging the visit to New York and for their comments. I also thank audiences at the following universities for many other helpful remarks: Edinburgh, Glasgow, Johns Hopkins, Nottingham, Oxford, Sheffield, Stirling, UCL, Utrecht, and Warwick.

Bill Brewer and Hannah Ginsborg kindly agreed to read the penultimate draft for Oxford University Press and suggested many improvements.

This is (I hope) a much better book than it would have been without their comments and those of an anonymous third reader.

John Campbell and I started talking philosophy in Oxford in 1983 and kept this up for over twenty years until his move to Berkeley. My thinking about the topics covered in this book owes a lot to our conversations over lunch in Wadham, New College, and the Duke of Cambridge. More recently, I have been helped by regular discussions with Ciara Fairley and by her incisive comments on several earlier drafts. Thanks are also due to the following for helpful comments or discussions: Elke Brendel, David Charles, Jennifer Church, Mizue Fukumitsu, Rory Madden, Hanna Pickard, Susanna Schellenberg, Declan Smithies, Paul Snowden, Charles Travis, and Tim Williamson. Apologies to anyone who should be on this list but isn't.

It has been a pleasure working with Peter Momtchiloff at OUP. Thanks also to Jackie Pritchard for her copy-editing and to Paul Raeside for the cover photograph. This book is dedicated to my parents, Sultan and Amir Cassam.

Contents

1. The Possibility of Knowledge	1
1.1 How-Possible Questions	1
1.2 Kant's Problem	10
1.3 Perceptual Knowledge	22
1.4 Anti-Minimalism	36
2. Transcendental Arguments	51
2.1 Regressive Transcendental Arguments	51
2.2 The Problem of Generality	62
2.3 Revelation, Validation, Explanation	67
3. Perceptual Knowledge (I): Space	85
3.1 What Makes it Possible?	85
3.2 Spatial Perception	90
3.3 Perceptual Differentiation (I)	93
3.4 Perceptual Differentiation (II)	107
3.5 The Direct Argument and Primary Qualities	116
3.6 Foundations, A Priority, and Minimalism	123
4. Perceptual Knowledge (II): Concepts	129
4.1 The Role of Concepts	129
4.2 Object Perception	132
4.3 Epistemic Perception	140
4.4 Foundations, A Priority, and Minimalism	152
5. Other Minds	155
5.1 The Perceptual Model	155
5.2 Obstacle-Removal	161
5.3 Enabling Conditions	172
5.4 How Many Levels?	183

6. A Priori Knowledge	188
6.1 The Problem	188
6.2 Reflection, Reasoning, and Calculation	196
6.3 Idealism and Explanation	210
6.4 Rationalism	220
<i>Bibliography</i>	225
<i>Index</i>	231

1

The Possibility of Knowledge

1.1 HOW-POSSIBLE QUESTIONS

I'm listening to a baseball game on the radio on the way home from my office. To my surprise, I hear the announcer describing a fielder catching a long fly ball that was just about to hit high up on the fence. I'm puzzled because I happen to know that the fence is twenty feet high. I wonder how, given the height of the fence, it was possible for the fielder to make the catch. I only stop being puzzled when I discover that he used a ladder attached to the scorekeeper's platform to reach the ball.¹ In the meantime, what I want to know is not *whether* it was possible for him to make a catch twenty feet off the ground (since he actually made it) but *how* this was possible. This is an example of a *how-possible* question. To ask a how-possible question is to ask how something which looks impossible given other things that one knows or believes is nevertheless possible.²

How-possible questions matter in philosophy because, as Nozick points out, 'many philosophical problems are ones of understanding how something is or can be possible' (1981: 8). Familiar philosophical how-possible questions include 'how is freedom possible?' and 'how is evil possible?'³ The first of these questions arises because of the tension between the natural assumption that we are capable of acting freely and the equally natural assumption that all actions are causally determined. But if all actions are causally determined doesn't that mean that freedom of action

¹ This is William Dray's example. See Dray 1957: 158. Dray got it from a magazine report on a baseball game in Canada. My account of how-possible questions owes a lot to Dray's seminal discussion, but he focuses on the philosophy of history rather than epistemology.

² As Dray puts it, 'explanation is called for because what happened seemed *impossible* in the circumstances' (1957: 160). Dray's conception of how-possible questions has also been taken up by Nozick and Stroud. See Nozick 1981: 8–11, and Stroud 1984: 144.

³ These are Nozick's examples. See Nozick 1981: 8–9.

is impossible? Similarly, if we take it that there is an omniscient, omnipotent, and benevolent God doesn't that make it impossible to account for the existence of evil in the world? So if we think that some of our actions are actually free, or that evil actually exists, then we ought to be puzzled. How-possible questions are a vivid way of giving expression to this puzzlement. We ask how freedom and evil are possible at the point at which the very existence of freedom and evil begins to strike us as problematic.

On this account, how-possible questions are *obstacle-dependent* questions. We ask how x is possible when there appears to be an obstacle to the existence of x .⁴ We don't ask how x is possible if there is no perceived obstacle or no inclination to suppose that x is possible. So, for example, we don't ask how baseball is possible or how round squares are possible. Where an obstacle-dependent how-possible question does arise there appear to be two basic strategies for dealing with it. The first is to deny the existence of the obstacle which gave rise to the question. This is an *obstacle-dissipating* strategy. So, for example, we can try to explain how freedom is possible by denying that all actions are causally determined, or how evil is possible by denying the existence of God.⁵ A different approach would be to argue that freedom is possible even if all actions are causally determined, or that evil is possible even if God exists. These are *obstacle-overcoming* rather than obstacle-dissipating strategies since they don't straightforwardly deny the existence of the obstacle in either case. What they deny is that the alleged obstacles are insuperable and, in this sense, genuine.⁶

⁴ This is what Nozick is getting at in the following passage: 'the form of [how-possible] questions is: how is one thing possible, given (or supposing) certain other things? Some statements r_1, \dots, r_n are assumed or accepted or taken for granted, and there is a tension between these statements and another statement p ; they appear to exclude p 's holding true. Let us term the r_i *apparent excluders* (of p). Since the statement p also is accepted, we face the question of how p is possible, given its apparent excluders' (1981: 9). Nozick's 'excluders' are my 'obstacles'. Although, as Nozick points out, logical incompatibility is the strongest mode of exclusion, it isn't the kind of incompatibility that is at issue in all how-possible questions. Dray's example clearly illustrates this point. In general, I take obstacle-dependence to be a pragmatic rather than a semantic phenomenon.

⁵ Denying the existence of a putative obstacle doesn't always do the trick. Sometimes, it just creates another obstacle. For example, suppose we deny that all actions are causally determined. Does that mean that actions occur randomly? If so, that is just as problematic for attributions of moral responsibility as the truth of determinism.

⁶ If denying that an obstacle is genuine is the same as denying that it exists then the distinction between the two obstacle-removing strategies will have to be explained slightly differently. In practice, however, the distinction between dissipating and overcoming an obstacle isn't difficult to understand. It corresponds to Nozick's distinction between two ways of dealing with an apparent incompatibility between a statement p and its apparent

What if it turns out that the obstacle which got the discussion going in the first place can neither be dissipated nor overcome? If we don't want to deny that all actions are causally determined but can't think of a plausible way of reconciling freedom and determinism then it seems that we are going to have to accept that freedom is impossible. This looks like a form of scepticism, so the next question is whether we can live with scepticism. Maybe we can. On the other hand, it might turn out that our intellectual commitment to the possibility of acting freely is stronger than our intellectual commitment to the principle that all actions are causally determined, or to the incompatibility of freedom and determinism. In that case, we should look again at whether the alleged obstacle to freedom can be overcome or dissipated. The worst case scenario in philosophical terms is if scepticism about the possibility of freedom is intolerable but the obstacle continues to strike us as insuperable however hard we try to make it disappear or find a way around or over it.

My concern is with epistemological rather than metaphysical, ethical, or theological how-possible questions. Not surprisingly, discussions of how-possible questions in the theory of knowledge have the same basic structure as discussions of how-possible questions in other branches of philosophy. Epistemological how-possible questions start from some cognitive achievement which they assume to be genuine and ask how that achievement is possible. The most general epistemological how-possible question, which is also often represented as one of the defining questions of epistemology, is:

(HP) How is knowledge possible?

As it stands, this question is entirely general in scope. It asks how it is possible for anyone or anything to know anything about anything. Typically, however, epistemological how-possible questions are more specific. For a start, (HP) is usually read as a question about *human* knowledge, as asking how it is possible for *us* and creatures like us to know anything.⁷ But even this looks too general to be manageable. A

excluder. First, 'one of the apparent excluders can be denied, or there can be a denial of their conjunction all together' (1981: 10). This amounts to what I am calling 'obstacle dissipation'. Second, 'each of the apparent excluders can continue to be maintained, while their apparent incompatibility with *p* is removed' (ibid.). To do this is to overcome a supposed obstacle.

⁷ The restriction to human knowledge should come as no surprise if we understand epistemology as 'the philosophical study of human knowledge' and therefore as seeking

better approach is to concentrate on specific types of knowledge. For example, we might ask how knowledge of the world around us, as distinct from self-knowledge or knowledge of logic, is possible. Other equally familiar epistemological how-possible questions concern a priori knowledge and knowledge of other minds. In each case, however, the basic idea is the same. We assume, at least to begin with, that we actually have some knowledge of the kind in question, and then try to account for its possibility.

A priori knowledge and knowledge of other minds will be the focus of later chapters. In this chapter I want to take a closer look at the following question:

(HP_{ew}) How is knowledge of the external world possible?

The ‘external world’ is the world around us, and to have knowledge of the external world is to have knowledge of the existence and nature of objects, processes, events, or states of affairs which exist independently of human thought or perception.⁸ My aim here is not just to outline an answer to (HP_{ew}) but also, at the same time, to develop a general approach to answering epistemological how-possible questions which can be used to explain how other kinds of knowledge are possible. The specific approach to epistemological how-possible questions which I will be defending is what I am going to call a *multi-levels* approach. As we will see, the tackling of obstacles to knowledge, either by dissipating or overcoming them, is one element of this approach but not the only element.

What is the intuitive obstacle to our knowing anything about the world around us which might lead one to ask (HP_{ew})? The fact is that there are many such obstacles, and that there is no hope of our getting to the bottom of every consideration which might conceivably result in someone asking how knowledge of the external world is possible. Obstacles can only be dealt with on a piecemeal basis, as and when they

to understand ‘what human knowledge is and how it comes to be’ (Stroud 2000d: 99). The focus on human knowledge also plays an important part in some accounts of scepticism. In his ‘Defence of Scepticism’, for example, Unger argues specifically for the thesis that ‘every human being knows, at best, hardly anything to be so’ (1971: 198).

⁸ Michael Williams objects that knowledge of the external world isn’t a ‘*kind*’ of knowledge, which we might assess or explain as a whole, and that in this sense there is ‘no such thing as knowledge of the external world’ (1996: p. xii). I don’t find Williams’s arguments for this view convincing but won’t discuss them here. If Williams is right, however, then (HP_{ew}) is not a good question.

are encountered. If this is right then the important thing is to have a good general sense of how to tackle the most commonly encountered obstacles. For illustrative purposes, therefore, I am going to concentrate on one consideration, or set of considerations, which has led some philosophers to wonder how we can know anything about the world around us. Hopefully, it will turn out that other commonly encountered obstacles can be dealt with in the same way.

The particular route to (HP_{ew}) which I have in mind has as its starting point a ‘how question’ rather than a how-possible question. The idea is this: if someone is said to know that *p*, where *p* is a proposition about the external world, we can always ask *how* he knows it. To ask how he knows it is to ask how he *came* to know it; it is a question about the source of his knowledge or his route to the knowledge that *p*.⁹ There need be no implication that the truth of *p* is something that he couldn’t know. This is what distinguishes a simple how-question, such as ‘how does *S* know that *p*?’, from the corresponding how-possible question, ‘how is it *possible* for *S* to know that *p*?’¹⁰ The how-question might be prompted by nothing more than what Austin calls ‘respectful curiosity’ (1979: 78). How-possible questions, in contrast, are usually expressions of more than just respectful curiosity. They are not just questions but challenges.

How, then, do we come to know things about the world around us? What is the source of our knowledge of the external world? The obvious answer is that there are lots of different ways of coming to know things about the external world.¹¹ Looking around, talking to people, reading newspapers, doing Google searches, are all ways of acquiring worldly knowledge, so what is the problem? One problem is that not all sources of knowledge are equal. Some are more basic than others, and the most basic general source of all is sense perception.¹² But why

⁹ As Austin remarks, the question ‘*how* do you know?’ is one to which (some) claims to knowledge are ‘directly exposed’ (1979: 77). See Hampshire 1979 for more on this theme, and for a defence of the idea that there are some knowledge claims in relation to which the how-question would be ‘at least absurd, and perhaps unintelligible’ (1979: 282).

¹⁰ Dray has some useful things to say about the differences between simple how-questions and how-possible questions. See Dray 1957: 166.

¹¹ Cf. Stroud 2000b: 3. The importance of recognizing that there are different ‘pathways’ to knowledge has also been emphasized by Alvin Goldman. See the preface to Goldman 2002.

¹² This is something that Stroud frequently emphasizes, as in this characteristic passage: ‘speaking “anthropologically”, we can say that human beings gain knowledge of the world through sense-perception. And when we look more closely into how

is that a problem? Because, according to one line of thinking, ‘there are certain apparently undeniable facts about sense-perception’ that make it difficult to understand ‘how we could get any knowledge at all of the world around us on the basis of sense-perception’ (Stroud 2000b: 5). If it is difficult to understand how we could get any knowledge of the world around us on the basis of sense perception then it is difficult to understand how we could get *any* knowledge of the world around us. And that is why (HP_{ew}) looks like a good question. On this account, the alleged obstacles to the acquisition of *perceptual* knowledge are among the epistemological obstacles which might reasonably lead us to wonder more generally how *any* knowledge of the external world is possible.

This way of getting to (HP_{ew}) brings out the limitations of what might otherwise seem the best way of tackling this question. Consider again the relationship between ‘how’ and ‘how-possible’ questions. We can explain how S knows that p by figuring out how he came to know that p, but figuring out how he actually *came* to know that p is also a way of explaining how it was *possible* for him to know that p. This suggests that we can explain how knowledge of the external world is possible by identifying ways or means of coming to know something about the world around us.¹³ Let’s call this a *Means Response* to (HP_{ew}). A Means Response to a how-possible question regards the identification of one or more of the means by which something can come about as a means of explaining how it is possible. So, for example, if perceiving is a means of coming to know something about the world around us then it is also a means by which knowledge of the external world is possible.

But this can’t be the end of the story if there are apparently undeniable facts about sense perception which threaten to make the acquisition of perceptual knowledge impossible. Given the obstacles to the acquisition of perceptual knowledge, all that the proposed Means Response to (HP_{ew}) does is to shift the focus of discussion from this question to another how-possible question, namely:

(HP_{pk}) How is perceptual knowledge possible?

sense-perception works, and what exactly we perceive, it can become difficult to see how perceptual knowledge of the world is possible’ (2000e: 123).

¹³ Stroud mentions ‘ways of coming to know something’ (2000b: 3) in one of his accounts of knowledge and scepticism. Peacocke also employs the notion of a ‘way of coming to know that something is the case’ in his account of a priori knowledge. See Peacocke 2000. Ways of coming to know are what Goldman calls ‘pathways to knowledge’. See n. 11 above.

What is the relationship between (HP_{pk}) and (HP_{ew})? Suppose that perceptual knowledge is defined as knowledge of the world that has its source in perception or perceptual experience. To define perceptual knowledge in this way is to define it by reference to both its *source* (perception) and its *subject matter* (the external world). In contrast, the knowledge that is at issue in (HP_{ew}) is identified just by its subject matter. This means that (HP_{ew}) and (HP_{pk}) call for somewhat different responses. It is a substantive claim that *knowledge of the external world* is possible by perceptual means but not that *perceptual knowledge* is possible by perceptual means. If there is such a thing as perceptual knowledge then, by definition, perception is its source.

It doesn't follow, however, that there is no such thing as a Means Response to (HP_{pk}). Suppose that I know that the cup into which I'm pouring coffee is chipped. How do I know this? By seeing that it is chipped. Seeing that the cup is chipped is a particular way of coming to know that the cup is chipped, and therefore a particular means by which perceptual knowledge is possible. Feeling that the cup is chipped is a different way of coming to know that the cup is chipped and therefore a different means by which perceptual knowledge is possible. These are examples of Means Responses to (HP_{pk}). Such responses are possible because 'perception' is a generic source of knowledge, and perceptual knowledge has different specific sources. To describe a piece of knowledge as 'perceptual' is to abstract from differences between different kinds or modes of perception, and that is why there is scope for a Means Response to (HP_{pk}).

Let us now return to the suggestion that it is difficult to understand how perception can provide us with knowledge of the world around us. What is the significance of this suggestion? One might think that its significance for (HP_{ew}) is limited because we could agree that knowledge of the world around us isn't possible by perceptual means without agreeing that it isn't possible by other means.¹⁴ On the other hand, it

¹⁴ The underlying question is this: if knowledge of the external world is attainable by 'a wide variety of sometimes independent and sometimes interconnected pathways' (Goldman 2002: p. xi), and if perception is just one of these pathways, why should we think that knowledge of the external world wouldn't be possible just because perceptual knowledge isn't possible? Why couldn't we rely on other sources to provide us with knowledge of the external world? The answer to this question depends, in part, on why perceptual knowledge is thought to be impossible. Suppose that it is thought to be impossible because we are incapable of perceiving external objects. But if we can't perceive external objects then, on some views of thought, we can't think about them either, and if we can't think about them then we can't have knowledge of them. In this

does seem right that perception is, for us, an absolutely fundamental source of knowledge, and that our epistemic situation would be very different from what we normally take it to be if perceptual knowledge isn't possible. There are many things about the world our knowledge of which is ultimately perceptual, so it is not as if we can easily give up on perceptual knowledge and concentrate on other sources. What this amounts to is the suggestion that if we are seriously interested in answering (HP_{ew}) then we can't just ignore (HP_{pk}).

This brings us back to the idea that there are obstacles to the acquisition or existence of perceptual knowledge. It remains to be seen what these obstacles are, but the implication of the discussion so far is that explaining how perceptual knowledge is possible has basically got to be a matter of identifying more specific perceptual means by which it is possible and dissipating or overcoming the alleged obstacles. Suppose, for example, that one particular obstacle takes the form of an epistemological requirement on the acquisition of knowledge, and that it is hard to see how perceptual knowledge is possible because it is hard to see how we could possibly meet the suggested requirement. In that case, an obstacle-overcoming response to (HP_{pk}) and, by implication, to (HP_{ew}) would be one which shows that we *can* satisfy the requirement. An obstacle-dissipating response, in contrast, would be one which makes it plausible that there is no such requirement. Indeed, on a Moorean view, the very fact that a particular epistemological requirement calls the possibility of perceptual knowledge into question is a good reason for rejecting that requirement.

I will come back to this suggestion later in this chapter. First, there are some other matters to discuss. So far, I have suggested that a good answer to an epistemological how-possible question will have at least two dimensions: it will identify specific means by which a particular kind of knowledge is possible and it will remove obstacles to the acquisition of knowledge by the proposed means. Is that enough? According to a position which I'm going to call *explanatory minimalism*, it is enough: explaining how knowledge of a certain subject matter is possible is simply a matter of identifying means of knowing about that subject matter, and showing that there is nothing that stands in the way of our

way, an alleged obstacle to perceptual knowledge transforms itself into an obstacle to our having *any* knowledge of the external world, even if perception isn't the only pathway to this kind of knowledge. This would be one way of justifying Stroud's insistence on the centrality of perceptual knowledge.

knowing about that subject matter by the proposed means. Once we have done that, we have done everything that can or needs to be done to answer the original how-possible question. In the case of propositions about the external world, therefore, the explanatory minimalist thinks that we have shown how knowledge of such propositions is possible by pointing out that perception is a means of coming to know things about the external world, and showing that there is nothing that stands in the way of our coming to know things about the external world by perceptual means.

Later in this chapter, I will be taking a closer look at explanatory minimalism. I will be defending the suggestion that more can be done to explain how knowledge is possible than it allows. Having identified perception as a means of knowing about the external world, and satisfied ourselves that there is nothing that prevents us from acquiring knowledge of the external world from this source, there is a further question which we might ask. The further question is: what *makes it possible* for us to acquire knowledge of the things around us by means of perception? This is a question about what might be called the *enabling conditions* of perceptual knowledge, the conditions under which it is possible for perception to be a source of knowledge of the things around us. The thought that underpins this question is that there is more to explaining how something is possible than showing that it isn't impossible.¹⁵ What we want is, as it were, a positive explanation of the possibility of perceptual knowledge, that is, an account of what makes perceptual knowledge possible.¹⁶

What we now have, in outline at least, are all the elements of a *multi-levels* account of the possibility of perceptual knowledge or knowledge of the external world. A multi-levels response to a how-possible question operates on three different levels. Level 1 is the level of means, the level at which means of knowing about a certain subject matter are identified. In the case of (HP_{ew}) the proposed means might

¹⁵ Nozick appears to have something similar in mind when he remarks that 'the task of explaining how *p* is possible is not exhausted by the rearguard action of meeting arguments from its apparent excluders' (1981: 11).

¹⁶ What-makes-possible explanations are also the focus of Peacocke's moderate rationalist account of a priori knowledge. There is a priori knowledge only if there are a priori ways of coming to know a proposition, and rationalism is the view that a priori ways of coming to know are grounded in the understanding. This leaves the rationalist with an explanatory task: 'the task is to say what it is about understanding that makes a priori knowledge possible' (Peacocke 2000: 257). There is more on all of this below, in Chapter 6.

be as generic as ‘perception’ or ‘testimony’. In contrast, (HP_{pk}) calls for the specification of more specific means since perceptual knowledge is defined by its generic source as well as its subject matter. In neither case, however, can the how-possible question be satisfactorily answered without saying something about means or ways of knowing. Level 2 is the obstacle-removing level, the level at which obstacles to the acquisition of knowledge by the proposed means are overcome or dissipated. Finally, Level 3 is the level at which enabling conditions for knowing by the proposed means are identified. The minimalist wants to stop at Level 2. The contrasting *anti-minimalist* approach says that we can or should continue to the level of enabling conditions, and that there are worthwhile philosophical questions about perceptual knowledge which are best understood as Level 3 questions.

This is only a sketch of the main elements of a multi-levels response, and it clearly needs a lot of fleshing out. A good way of doing this would be to look at an actual example of a multi-levels response to a how-possible question in the history of philosophy. The particular philosopher who comes to mind in this connection is Kant. He did more than anyone to bring how-possible questions to prominence in philosophy, and his approach to such questions is a multi-levels approach. But his topic, at least in the introduction to the first *Critique*, is not the possibility of perceptual knowledge, or anything as general as ‘our knowledge of the external world’. Instead, his how-possible questions all boil down to the single question:

(HP_{sap}) How is synthetic a priori knowledge possible?

Nevertheless, if it is true that Kant’s response to (HP_{sap}) is a multi-levels response then it would be worth taking a closer look at what he says. As will shortly become apparent, his response to (HP_{sap}) raises some important general questions about the multi-levels approach, and these questions will need to be addressed before we return to the project of developing a multi-levels response to (HP_{pk}) .

1.2 KANT’S PROBLEM

The background to Kant’s interest in (HP_{sap}) is his conviction that pure mathematics and pure physics are bodies of synthetic a priori knowledge. A priori knowledge is knowledge that doesn’t have its source in experience. Synthetic knowledge is non-analytic knowledge,

knowledge that doesn't have its source in 'the analysis of concepts' (B23).¹⁷ Our knowledge of the propositions of pure mathematics and pure physics must be a priori, according to Kant, because these propositions are necessarily true, and because experience can't provide us with knowledge of necessary truths. At the same time, given that the propositions of pure mathematics and pure physics aren't analytically true, and that conceptual analysis can't provide us with knowledge of propositions which aren't analytically true, mathematical and pure scientific knowledge must be synthetic.

Kant remarks that since pure mathematics and pure natural science actually exist, 'it is quite proper to ask *how* they are possible; for that they must be possible is proved by the fact that they exist' (B20). In contrast, metaphysics *ought* to contain synthetic a priori knowledge, but 'we cannot assume metaphysics to be an actual science' (Kant 1977: 275). Thus, as far as metaphysics is concerned, the issue is *whether* it is possible; in the case of pure mathematics and pure natural science we only have to think about *how* they are possible. In the words of the *Prolegomena*:

We have therefore some, at least uncontested, synthetic *a priori* knowledge, and need not ask *whether* it be possible (for it is actual) but *how* it is possible, in order that we may deduce from the principle that makes the given knowledge possible the possibility of all the rest. (1977: 275)

So the idea is that once we understand how uncontested synthetic a priori knowledge is possible, we should be in a better position to establish and explain the existence of more controversial forms of synthetic a priori knowledge, including metaphysical knowledge.

Suppose that Kant is right to assume that pure mathematics and pure physics are bodies of synthetic a priori knowledge. Why, then, is it 'quite proper to ask how they are possible'? Given my account of how-possible questions, this question only arises if synthetic a priori knowledge is in some way problematic, that is, if there is at least an apparent obstacle to its existence. The apparent obstacle is this: according to Kant, experience and conceptual analysis are two basic sources of human knowledge, but we have already seen that neither experience nor conceptual analysis can yield synthetic a priori knowledge. Experience can't yield a priori knowledge and conceptual analysis can't yield synthetic knowledge. That is why it is proper to ask how synthetic a priori knowledge is

¹⁷ All references in this form are to Kant 1932.

possible. What gives this question its philosophical bite is the existence of synthetic a priori knowledge that can't be accounted for by reference to certain presupposed basic sources of knowledge. Let's call this the *problem of sources*.

One response to this problem would be to argue that it's false that neither experience nor conceptual analysis can account for mathematical knowledge. Conceptual analysis can account for it if mathematical truths are analytic rather than synthetic. Alternatively, there is no reason why mathematical knowledge couldn't come from experience if mathematical truths aren't necessary, or if it is false that experience can't provide us with knowledge of a necessary truth. Each of these obstacle-dissipating responses to Kant's question amounts to what might be called a *presupposed sources* solution to the problem of sources. In each case the possibility of mathematical knowledge is accounted for by reference to one of the presupposed sources of knowledge. However, this isn't Kant's own preferred solution. His solution is an *additional sources* solution since it involves the positing of what he calls 'construction in pure intuition' as an additional source of knowledge by reference to which at least the possibility of geometrical knowledge be accounted for.¹⁸ It is 'additional' not in the sense that mathematicians haven't been constructing in pure intuition all along but in the sense that no account was taken of this source of mathematical knowledge in the discussion leading up to the posing of (HP_{sap}).

Viewed in one way, the additional sources solution looks like an obstacle-overcoming rather than an obstacle-dissipating response to (HP_{sap}) since it doesn't dispute the suggestion that neither experience nor conceptual analysis can account for our mathematical knowledge. It endorses this suggestion but argues that mathematical knowledge is still possible. Viewed in another way, however, Kant's solution to the problem of sources looks more dissipationist. What gives rise to this problem in its sharpest form is the assumption that experience and conceptual analysis are our *only* sources of knowledge. Kant challenges this assumption. He implies that if we had acknowledged the existence of non-conceptual yet non-empirical sources of knowledge in the first place we wouldn't have been puzzled by the existence of synthetic a priori knowledge. This suggests that there isn't a genuine obstacle to the existence of this kind of knowledge, only a false assumption about the

¹⁸ Kant defends the view that 'mathematical knowledge is the knowledge gained by reason from the construction of concepts' (A713/B741) in the section of the *Critique of Pure Reason* called 'The Discipline of Pure Reason in its Dogmatic Employment'.

range of epistemic sources that are available to us. Once we reject this assumption the perceived obstacle disappears. So what started out as an obstacle-overcoming exercise has ended up as an exercise in obstacle dissipation.

When Kant asserts that geometrical knowledge is the knowledge gained by reason from the construction of concepts he is reminding his readers that Euclidean geometrical proofs are essentially diagrammatic. To construct a figure in pure intuition is to 'draw' it in the imagination, and Kant's proposal is that the construction of geometrical concepts in pure intuition is a genuinely non-conceptual, non-empirical means of coming to know geometrical truths, and therefore a means of acquiring synthetic a priori knowledge. In my terms, this is a Level 1 response to (HP_{sap}). In keeping with the general idea of a Means Response to a how-possible question, Kant is assuming that the identification of means by which a specific type of knowledge is possible is a means of explaining how knowledge of that type is possible.

Let's take a closer look at this assumption. The first thing to notice is that the means by which something is possible needn't be necessary conditions for its possibility. The precise significance of this point will become clearer as we go along. In the meantime, a simple example might help to clarify the distinction between means and necessary conditions: if someone asks how it is possible to get from London to Paris in less than three hours, it would be an acceptable answer to say that one can do this by catching the Eurostar. Catching the Eurostar is a means of getting from London to Paris in less than three hours but obviously not a necessary condition; one can also fly. By the same token, there is a difference between identifying a means of achieving something and showing that it is the *only* means of achieving it. Catching the Eurostar is a means of getting from London to Paris but not the only means.

The question which this example raises is whether, assuming that one's sole aim is to answer (HP_{sap}), it matters whether construction in intuition is the only means of acquiring synthetic a priori knowledge. It would seem not. If one lands up asking (HP_{sap}) because one thinks that some of our knowledge is synthetic a priori, and because one can't figure out how synthetic a priori knowledge is possible, it is enough that intuitive construction is *a* means of acquiring synthetic a priori knowledge. If it is the only means then that would obviously be interesting but it is hard to see why a good answer to (HP_{sap}) has got to establish unique means. As far as Kant is concerned, constructing concepts in pure intuition is our only route to geometrical knowledge

but not our only route to synthetic a priori knowledge. He thinks that some of our synthetic a priori knowledge is metaphysical rather than mathematical, and that there is no question of construction in pure intuition being the source of our non-mathematical synthetic a priori knowledge. So his basic answer to (HP_{sap}) is that synthetic a priori knowledge is possible by several different means, construction in pure intuition being one of these means.

Even if one sticks to one specific type of synthetic a priori knowledge, say geometrical knowledge, it's not obvious how one would establish unique means. For example, how can we be sure that some form of 'rational insight' couldn't possibly be a source of geometrical knowledge? In reality, however, this question is neither here nor there given that rational insight is not a source of knowledge that is available to us. The lesson is that a good response to a how-possible question needs to identify *practical* rather than *unique* means. For the purposes of explaining how geometrical synthetic a priori knowledge is possible, all that matters is that the construction of concepts is a source of this kind of knowledge, and that it is something that we are actually capable of doing.

At this point we run into the following problem: if the recommended answer to 'how is it possible for us to acquire synthetic a priori geometrical knowledge?' is: 'by constructing concepts in pure intuition', one might reasonably think that this only raises a further and no less pressing how-possible question, namely, 'how is it possible for construction in pure intuition to be a source of synthetic a priori knowledge?' If how-possible questions are obstacle-dependent, then someone who asks this question must have a particular obstacle in mind, and it's not difficult to guess what that obstacle might be. According to Kant, strict universality is a criterion of a priori knowledge, so a priori geometrical knowledge must be of strictly universal propositions. Yet what we construct in pure intuition are individual figures. So if we want to understand *how* construction in pure intuition can be a source of a priori geometrical knowledge, we must explain how, as Kant puts it, the construction of a concept by means of a *single* figure can 'express universal validity for all possible intuitions which fall under the same concept' (A713/B741). Let's call this the *problem of universality* for Kant's theory of geometrical proof.

Kant's solution to this problem is buried in the depths of the Schematism.¹⁹ Roughly, the idea is that geometrical proofs have universal validity

¹⁹ See, especially, A141/B180.

as long as the figures that mathematicians construct are ‘determined by certain universal conditions of construction’ (A714/B742). These universal conditions of construction are the *schemata* of geometrical concepts, that is, rules for constructing them. This is now a Level 2 response to (HP_{sap}). Having identified construction in pure intuition as a means of acquiring synthetic a priori knowledge, and acknowledged an apparent conflict between the singularity of constructed figures and the strict universality of a priori geometrical propositions, Kant draws on the Schematism in an attempt to overcome this apparent obstacle to the acquisition of geometrical knowledge by the proposed means. In essence, his proposal is that it is the fact that construction is a rule-governed activity that makes it possible for geometry to discern ‘the universal in the particular’ (A714/B742).

Since we are at present concerned with the structure of Kant’s account rather than its details, it doesn’t matter very much whether this exercise in obstacle-removal is successful. The important point, as far as Kant is concerned, is that the problem of universality *must* be soluble because the existence of synthetic a priori geometrical knowledge is a given. He is certain that geometry *is* a body of synthetic a priori knowledge, and that construction in pure intuition *is* its source. What we have to do, therefore, is to figure out *how* these things are possible. He thinks that it’s just not an option to end up concluding that we don’t have any synthetic a priori knowledge after all, or that we can’t get it from construction in pure intuition. In both these respects, Kant’s attitude towards (HP_{sap}), at least in the *Prolegomena*, might be described as *Moorean*. What this means is Kant regards the claim that synthetic a priori knowledge exists as non-negotiable, somewhat in the way that Moore regards the existence of perceptual knowledge as non-negotiable.²⁰ For Kant, it makes no more sense to question the existence of synthetic a priori knowledge than to question the existence of pure mathematics.

For the moment, this is as much as I want to say about Kant’s Level 2 response to (HP_{sap}). Let us now consider whether, if Kant’s story so far were convincing, he would need to go any further in relation to (HP_{sap}). The explanatory minimalist thinks not but Kant disagrees. His thought is that even though construction in pure intuition is a bona fide source of synthetic a priori geometrical knowledge, there is more that needs to be said. What still needs to be explained is the capacity of construction in pure intuition to provide us with geometrical knowledge.

²⁰ On the non-negotiability of perceptual knowledge see Moore 1922.

Explaining this isn't just a matter of dealing with actual or apparent obstacles to the acquisition of geometrical knowledge; it's also a matter of explaining what makes it possible for construction in pure intuition to be a source of synthetic a priori geometrical knowledge. What we now have, therefore, is a *what-makes-it-possible* question rather than a how-possible question. How-possible questions are obstacle-dependent but what-makes-it-possible questions are *explanation-seeking*.²¹ What they seek is not a way round some specific obstacle but, as it were, a positive explanation of the possibility of acquiring a certain type of knowledge by certain specified means.

Everything now depends on the nature of the required explanation and the reasons for thinking that any such explanation is either necessary or possible. On the first of these issues, there are at least two different things which might be involved in explaining what makes it possible for construction in pure intuition to be a source of geometrical knowledge. On the one hand, it might involve explaining what makes construction in pure intuition possible. Let's call this a *type A* explanation. In my terms, a type A explanation is one that seeks to explain the possible *occurrence* of a certain cognitive activity. On the other hand, one might think that what needs explaining is not just the possibility of constructing figures in pure intuition but also the fact that doing this is a source of a certain kind of *knowledge*. Let's call this a *type B* explanation; what makes it a type B explanation is that it seeks to explain the *epistemological significance* of a certain cognitive activity.

What would a type A explanation of the possibility of constructing figures in pure intuition look like? To explain what makes this possible would be to identify the conditions under which it is possible. The conditions under which it is possible are its enabling conditions. There are many different kinds of type A enabling condition that one might have in mind at this point. For example, suppose that only thinkers with a certain physiology could construct figures in pure intuition. Possession of the appropriate psychology would constitute one kind of enabling condition, a *physiological* enabling condition. Another line of thinking would be that only thinkers with certain other cognitive capacities, such as the capacity to see or touch, would be able to construct concepts in pure intuition. Such cognitive capacities would constitute *cognitive* enabling conditions for construction in pure intuition.

²¹ See McDowell 1998d: 444 for the notion of an explanation-seeking question.

What are enabling conditions? In essence, they are a sub-class of necessary conditions. For example, the existence of at least one thinker is a necessary condition for the construction of figures in intuition but it isn't what I am calling an enabling condition. The existence of at least one thinker is necessary for there to be any thinking at all, not just necessary for the kind of thinking that is involved in the construction of geometrical proofs. In contrast, enabling conditions are more specific; enabling conditions for the construction of figures in intuition aren't necessary conditions for there to be any cognitive activity whatsoever. Enabling conditions are necessary conditions for achieving something by a particular means. Relatedly, enabling conditions are *background* conditions, which may or may not be causal.²² Being an unmarried man is a necessary condition for being a bachelor but being an unmarried man isn't an enabling condition for being a bachelor. Intuitively, the reason is that being an unmarried man isn't a 'background condition' for being a bachelor. Being an unmarried man doesn't just 'enable' one to be a bachelor, it is what being a bachelor consists in.

Burge makes the notion of a 'background condition' a bit more precise in his paper 'Content Preservation'. One of his examples concerns the role of memory in deduction. The important point is that claims about memory or the past needn't figure as *premisses* in one's reasoning even though any reasoning in time must *rely* on memory. Specifically:

In a deduction, reasoning processes' working properly depends on memory's preserving the results of previous reasoning. But memory's preserving such results does not add to the justificational force of the reasoning. It is rather a background condition for the reasoning's success... Memory failures that cause demonstrations to fail are failures of background conditions necessary to the proper function of reasoning. (1993: 463–4)

What we have here is, in effect, an answer to a what-makes-it-possible question. The question is: 'what makes transitions of reason possible?', and the answer is: preservative memory. Something similar can be said about the role of perception in interlocution, since perception is

²² Enabling conditions figure in Dretske's account of 'epistemic seeing' (seeing that *b* is *P*) and in Searle's account of intentionality. Both writers stress that enabling conditions, as they conceive of them, are background conditions. See Dretske 1969: 82–3 and Searle 1983: 141–59. But both Dretske and Searle take it that enabling conditions are causally necessary conditions. See, for example, Searle 1983: 157–8. This isn't built into my account of enabling conditions, which is not to deny that some enabling conditions are causal.

a background condition necessary for the acquisition of beliefs from others:

In interlocution, perception of utterances makes possible the passage of propositional content from one mind to another rather as purely preservative memory makes possible the preservation of propositional content from one time to another. Memory and perception of utterances function similarly, in reasoning and communication respectively. Their correct functioning is necessary for the enterprises they serve. (Burge 1993: 481)

But the role of memory and perception in these enterprises is only to 'preserve and enable'. In my terms, memory is an enabling condition for reasoning and that perception is an enabling condition for interlocution.

With this account of the notion of enabling conditions in place, we can now briefly examine Kant's type B explanation of what makes it possible for the construction of figures in pure intuition to be a source of geometrical knowledge. The proposed explanation is metaphysical rather than one in terms of cognitive capacities. Geometry is supposed to tell us something about the nature of physical space, so a type B explanation will be an account of what makes it possible for the construction of figures in pure intuition to yield knowledge of the geometry of physical space. According to Kant, what makes this possible is the fact that space itself is an '*a priori* intuition' that 'has its seat in the subject only' (B41). The proposal, in other words, is that the transcendental ideality of space is a background enabling condition for the acquisition of geometrical knowledge by means of construction in intuition. That's why, if we want to explain how geometrical knowledge is possible, we must be transcendental idealists.

How convincing is this argument? We might fail to be convinced by it because we reject its starting point, because we think that Kant was wrong to regard geometry as a body of synthetic *a priori* knowledge or as providing us with *a priori* knowledge of physical space. Or we might think that his solution to the problem of universality fails for one reason or another. Finally, we might reject the last step of the argument on the grounds that it fails to establish that it would be impossible to account for synthetic *a priori* knowledge of physical space outside an idealist framework. We might wonder, for example, why it wouldn't be possible for us to have *innate* synthetic *a priori* knowledge of the geometry of physical space even if space isn't ideal. If this is a genuine possibility then the argument for idealism doesn't go through.

We don't need to go into these issues here. Instead, let's focus on the relevance of Kant's position for the dispute between minimalism and anti-minimalism. Suppose, then, that what is at issue is the acquisition of knowledge of kind K by means M. Minimalists shouldn't deny that there might be causal enabling conditions for the acquisition of K by M. What they should say is that the uncovering of such conditions is a matter for empirical science rather than armchair philosophy.²³ Minimalism is therefore the view that distinctively *philosophical* explanations of the possibility of knowledge can't go beyond Level 2.²⁴ In contrast, *moderate* anti-minimalism is not sceptical about the existence of enabling conditions for the acquisition of K by means of M which philosophical reflection can bring to light. It agrees, therefore, that philosophical Level 3 explanations are possible. It denies, however, that such explanations are necessary for the purpose of explaining how knowledge of K kind is possible. Finally, *extreme* anti-minimalism is the view that philosophical Level 3 explanations are both possible and necessary. According to the extreme anti-minimalist, if there are Level 3 conditions for the acquisition of K by M that philosophy, and perhaps only philosophy, can bring to light, then a philosophical explanation of how knowledge of kind K is possible by means M will be incomplete unless it says what these conditions are.

In these terms, Kant's account of what makes it possible for the construction of figures in pure intuition to be a means of acquiring geometrical knowledge looks like an implicit argument for extreme anti-minimalism. The argument goes like this: it is a priori philosophy rather than empirical science that tells us that the ideality of space is what makes it possible for construction in pure intuition to be a source of geometrical knowledge. So it is false that a philosophical account of what makes geometrical knowledge possible *can't* go beyond Level 2; Kant's is a philosophical account of what makes geometrical

²³ This would presumably be Quine's view. See Quine 1969. What is effectively Quine's approach to the study of enabling conditions is helpfully summarized in the following terms by Stroud: 'It is not that we have any difficulty understanding the general idea of a study of the conditions of human knowledge—an investigation of those characteristics of human organisms that make it possible for them to come to know things about what is going on around them. But the best way to carry out such a study would seem to be by observing human beings and trying to understand how they work. It would be an empirical investigation—which is not to say that it would be easy to carry out' (1984: 160).

²⁴ This assumes, of course, that distinctively philosophical explanations are armchair explanations, and that philosophy is therefore discontinuous with empirical science.

knowledge possible that *does* go beyond Level 2. It is also false that we can adequately explain how geometrical knowledge is possible without saying anything about the metaphysical conditions under which this is possible. If it really only makes sense to suppose that construction in pure intuition is a source of geometrical knowledge on the assumption space is ideal, how can reference to this fact possibly be an optional extra in a philosophical explanation of the possibility of geometrical knowledge? This, then, is the basis of Kant's extreme anti-minimalism. It is the idea that it is essential for a philosophically satisfying response to (HP_{sap}) to reach all the way down to the level of enabling conditions.

From a multi-levels perspective the problem with this argument is that it calls into question the distinction between Level 2 and Level 3. For surely Kant only thinks that the ideality of space is a necessary condition for the acquisition of geometrical knowledge by means of construction in intuition because he takes it that the mind-*independence* of space would constitute a kind of obstacle to our acquiring knowledge of its geometrical properties by these means. But this looks like a reason for locating his argument for the ideality of space at Level 2 rather than at Level 3. More generally, if C is an enabling condition for the acquisition of knowledge of kind K by means M then the non-fulfilment of C can't fail to represent an obstacle to the acquisition of K by M. But if the denial of a putative enabling condition is always an obstacle to knowledge doesn't it follow that the suggested distinction between Level 2 and Level 3 explanations is spurious?²⁵ In that case, the only sense in which a philosophically satisfying response to (HP_{sap}) must reach all the way down to the level of enabling conditions is that such an account can't afford to ignore obvious obstacles to the acquisition of synthetic a priori knowledge by means of construction in pure intuition.

What is right about this is that when an enabling condition C for the acquisition of K by M isn't fulfilled, the very fact that it isn't fulfilled *becomes* an obstacle to the acquisition of K by M. It doesn't follow from this that the point of identifying C must be to defuse some intuitive, pre-existing obstacle to the acquisition of K by M. Suppose, for example, that it is true that in order to be able to construct concepts in intuition one must have the capacity to see or touch. There is no intuitive obstacle to the acquisition of geometrical knowledge by means of construction in intuition that one would be overcoming or dissipating by identifying such perceptual capacities as background

²⁵ Thanks to Bill Brewer and Tim Williamson for raising this question.

necessary conditions for construction in intuition. If one *lacked* the relevant perceptual capacities that would make it impossible for one to acquire geometrical knowledge by constructing concepts in intuition, but that doesn't make the mere *identification* of enabling conditions an exercise in obstacle-removal. Identifying what makes it possible for construction in pure intuition to be a source of geometrical knowledge is not primarily a matter of responding to independent reasons for thinking that it couldn't be.

This suggests that there is a distinction between Level 2 and Level 3 explanations of the possibility of knowledge even if, as Kant's argument for the ideality of space illustrates, it is sometimes unclear whether a particular explanation belongs at Level 2 or at Level 3. Once it is agreed, however, that there are some explanations of the possibility of knowledge that clearly belong at Level 3 rather than at Level 2 then moderate anti-minimalism begins to look like a serious option. For example, it is not obvious that someone who thinks that one would not be able to construct concepts in pure intuition if one lacked certain perceptual capacities is committed to thinking that an explanation of the possibility of geometrical knowledge which fails to mention this fact is necessarily inadequate. Given the obstacle-dependence of how-possible questions, an inadequate explanation of how knowledge of geometry is possible is not one that fails to identify enabling conditions. An inadequate explanation is one that fails to come up with viable means of acquiring geometrical knowledge and which therefore fails to solve the problem of sources.

The point of looking at Kant's response to (HP_{sap}) was to flesh out the idea of a multi-levels response to an epistemological how-possible question. Kant's multi-levels response to (HP_{sap}) is illuminating because it shows how obstacle-overcoming can sometimes shade off into obstacle dissipation. It doesn't show that there is no distinction between Level 2 and Level 3 explanations of the possibility of knowledge but it does suggest that this distinction is not always as sharp as my initial discussion might have suggested. It brings into focus the differences between different kinds of enabling condition, and between different forms of anti-minimalism. It brings out the force of extreme anti-minimalism but fails to establish that, at least in relation to (HP_{sap}), moderate anti-minimalism isn't in the running.

Bearing all of this in mind, let's now go back to (HP_{pk}), and to the project of filling in a multi-levels response to this question. We have already seen how (HP_{ew}), which is where we started, can lead to (HP_{pk}). We have

the outlines of a multi-levels response to (HP_{pk}) but very little sense of how to fill in the details. If there are obstacles to the acquisition of knowledge of the external world by perceptual means then we haven't explained how knowledge of the external world is possible unless we have tackled these obstacles. Having done that, we can then consider what makes it possible to acquire perceptual knowledge, and whether we should be minimalists, moderate anti-minimalists, or extreme anti-minimalists in relation to (HP_{pk}) and, by implication, in relation to (HP_{ew}).

1.3 PERCEPTUAL KNOWLEDGE

The first problem we encounter when we try to apply Kant's multiple-levels account of (HP_{sap}) to (HP_{pk}) is this:²⁶ the possibility of synthetic a priori knowledge needs explaining because neither experience nor conceptual analysis can be its source. In contrast, it is not the case that experience can't be the source of perceptual knowledge, so it seems that (HP_{pk}) is going to have to be motivated by something other than the problem of sources. Indeed, Kant is an example of a philosopher who appears to think that (HP_{pk}) lacks *any* respectable motivation. In his terms, perceptual knowledge would be synthetic a posteriori knowledge, but 'the possibility of synthetic *a posteriori* judgements, of those which are gathered from experience ... requires no special explanation; for experience is nothing but a continual joining together (synthesis) of perceptions' (1977: 275). If the possibility of synthetic a posteriori knowledge requires no special explanation, then (HP_{pk}) simply doesn't arise; there is no obstacle for it to trade on.

The point generalizes. Empiricists characterize a priori knowledge 'negatively', as knowledge that doesn't come from experience or that isn't justified by experience. 'Empirical' knowledge is characterized 'positively' as knowledge that comes from experience or that is justified by experience. If what we know about a priori knowledge is that it doesn't come from experience, then an obvious question is:

(HP_{apk}) How is a priori knowledge possible?

Answering this question will be a matter of identifying sources of a priori knowledge, or the means by which it is possible, given that experience

²⁶ This needs to be qualified. There are, of course, conceptions of experience on which experience can't be a source of perceptual knowledge but it is neither necessary nor plausible to think of experience in this way.

isn't its source. What we need, in other words, is at least a Level 1 explanation of the possibility of a priori knowledge. But if empirical knowledge is defined as knowledge that originates in experience, then we already know how it is possible; it's possible by means of experience. The means by which we acquire empirical knowledge is already built into our conception of what makes it *empirical* knowledge, just as the means by which we acquire perceptual knowledge is built into our conception of what makes it *perceptual* knowledge. In neither case, therefore, is there any need to look for an explanation of its possibility.

This attempt to bypass (HP_{pk}) fails. For a start, it's easy to imagine a priori knowledge being given the positive and empirical knowledge the negative characterization. A rationalist might stipulate that a priori knowledge is knowledge that is grounded in reason or rational intuition and that empirical knowledge is knowledge that is not so grounded. From this perspective, the pressing question would be:

(HP_{ek}) How is empirical knowledge possible?

This, rather than (HP_{apk}), would now be the pressing question because we would no longer be building the source of empirical knowledge into our conception of what makes it *empirical* knowledge. The rationalist would then regard it as a substantive claim that perception is a potential source of empirical knowledge, in the same way that the empiricist regards it as a substantive claim that rational intuition is a potential source of a priori knowledge.

This is just a way of making the obvious point that what strikes one as a worthwhile how-possible question is bound to be influenced by one's background assumptions. Different assumptions throw up different potential obstacles to knowledge and different how-possible questions. Still, it's important not to exaggerate the significance of this consideration. For even if one defines a priori knowledge as knowledge that is grounded in reason, this doesn't mean that (HP_{apk}) can't or doesn't arise. Someone who defines a priori knowledge in this way still needs to explain *how* reason can be a source of a priori knowledge.²⁷ In explaining how this can be so, one might find oneself running into actual or potential obstacles to the acquisition of a priori knowledge from this source, and this would leave (HP_{apk}) with considerable philosophical bite. By the same token, the fact that empirical or perceptual knowledge is characterized as having its source in experience or perception will still

²⁷ Peacocke 2000 is helpful on this point.

leave (HP_{ek}) and (HP_{pk}) with considerable philosophical bite if there are actual or potential obstacles to the acquisition of knowledge from experience or perception. Either way, defining a type of knowledge by reference to its source needn't prevent one from asking a how-possible question about that type of knowledge; in asking the question, one would be asking *how* the proposed source can be a source of knowledge of that type.

This attempt to rehabilitate (HP_{pk}) from an empiricist perspective assumes that there are actual or potential obstacles to our coming to know things about the external world by means of perception. What are these obstacles? Many different answers to this question have been proposed but for present purposes I'm going to focus on Stroud's account of one alleged obstacle. Although I will be suggesting that this obstacle isn't genuine, seeing why not will be the first step along the way to the overall approach to (HP_{pk}) that I will be recommending. In essence, I want to try to make it plausible that there is a viable alternative to explanatory minimalism, and I will do this by developing and defending a moderately anti-minimalist multiple levels response to Stroud's version of (HP_{pk}).

As we have already seen, Stroud's basic thought is that there are certain apparently undeniable facts about sense perception that make it difficult to see exactly how sense perception works to give us knowledge of the world. Specifically, the difficulty is that:

it seems at least possible to perceive what we do without thereby knowing something about the things around us. There have been many versions of that fundamental idea. But whether it is expressed in terms of 'ideas' or 'experiences' or 'sense data' or 'appearances' or 'takings' or 'sensory stimulations', or whatever it might be, the basic idea could be put by saying our knowledge of the world is 'underdetermined' by whatever it is that we get through that source of knowledge known as 'the senses' or 'experience'. Given the events or experiences or whatever they might be that serve as the sensory 'basis' of our knowledge, it does not follow that something we believe about the world around us is true. The problem is then to explain how we nevertheless know that what we believe about the world is in fact true. Given the apparent 'obstacle', how is our knowledge possible? (Stroud 2000b: 5–6)

This passage is in keeping with the idea that (HP_{pk}) is an obstacle-dependent question. Like Kant's worry about synthetic a priori knowledge, Stroud's worry about perceptual knowledge is that it can't easily be accounted for by reference to our presupposed cognitive resources. These resources include perception, but if several different possibilities

are compatible with our perceiving what we do then it is hard to see how perception could be a source of knowledge of the things around us. In other words, it is hard to see how perceptual knowledge is possible, and that is why (HP_{pk}) is a genuine question despite the fact that perceptual knowledge is defined as knowledge that we get from perception.

There is, however, one crucial difference between Kant's question about synthetic a priori knowledge and Stroud's question about perceptual knowledge. The difference is that the geometrical version of (HP_{sap}) can be answered by positing construction in pure intuition as an 'additional' source of knowledge, whereas it's obviously not going to be acceptable to respond to (HP_{pk}) by positing additional, non-perceptual sources of perceptual knowledge. It's not as if it would make sense to agree that sense perception can't give us knowledge of the world and then argue that this doesn't matter because such knowledge is possible by some other means. Trivially, perceptual knowledge is possible only if *perception* can be a source of knowledge, and it is at least arguable that knowledge of the external world is only possible if perceptual knowledge is possible. So what needs to be made plausible is that perception can provide us with knowledge of the world. This would be a 'presupposed sources' rather than an 'additional sources' response to (HP_{pk}), one which explains how perceptual knowledge is possible by reference to its presupposed, canonical source.

How can perception be a source of knowledge? The obstacle-generating principle that underpins (HP_{pk}) in Stroud's discussion is the principle that:

(U) Our knowledge of the world is underdetermined by what we get through the senses.

If this is the obstacle that stands in the way of the acquisition of perceptual knowledge, then a Level 2 response to (HP_{pk}) can take one of two forms. One possibility would be to accept (U) but argue that it doesn't prevent perception from being a means of coming to know things about the world around us. This would be an *obstacle-overcoming* response to Stroud's version of (HP_{pk}). The other possibility would be to dispute (U). If this principle is incorrect, then it can't amount to a genuine obstacle to perceptual knowledge. This would be an *obstacle-dissipating* response to (HP_{pk}). While there might be genuine obstacles to perceptual knowledge, the obstacle-dissipating response to (HP_{pk}) suggests that (U) isn't one of them.

In order to decide between these responses to (HP_{pk}) , more needs to be said about (U) and about the sense in which it calls into question the possibility of perceptual knowledge. The point of (U) is to suggest that what we get through the senses is not knowledge of the world but knowledge of something that is *epistemically prior* to knowledge of the world. Specifically, what we get through the senses is only knowledge of the character of our sensory experiences themselves or knowledge of how things seem to us to be. To say that things of one sort are epistemically prior to things of another sort is to say that ‘things of the first sort are knowable without things of the second sort being known, but not *vice versa*’ (Stroud 1984: 141). But if knowledge of how things seem to us to be is epistemically prior to knowledge of the world, then the truth of propositions about the external world would need to be *inferred* from what we get through the senses. Yet inferences from the character of our sensory experiences can’t provide us knowledge of external reality.²⁸ To regard perceptual knowledge as inferential is therefore to call its very possibility into question. So if (U) is a genuine obstacle, then it is one that cannot be overcome; unlike the obstacle to catching a fly ball twenty feet off the ground, there is no ladder that one can use to climb over it.

If the prospects for an obstacle-overcoming response to (HP_{pk}) are really as dim as this argument suggests then an obstacle-dissipating response to (HP_{pk}) must be the way to go. An obstacle-dissipating response to (HP_{pk}) would be one that disputes the existence of the alleged obstacle. Given that ‘the apparent obstacle to our knowledge comes from the doctrine of the epistemic priority of sensory experiences over independently existing objects’ (Stroud 1984: 143), and that (U) is simply an expression of this alleged epistemic priority, the obvious obstacle-dissipating response to (HP_{pk}) would be one that calls this doctrine into question. The claim that needs to be made out, therefore, is that our knowledge of the external world is *not* underdetermined by what we get through the senses, and that this is so because what we get through the senses is *not* epistemically prior to knowledge of the world.

To see what such an obstacle-dissipating response to (HP_{pk}) might look like in practice, consider the following example: as I’m pouring myself a cup of coffee I see that the cup is chipped. If I see that the cup is chipped, I know that the cup is chipped. In general, ‘S sees that p’ entails

²⁸ This claim is controversial but let’s not worry about that here.

'S knows that p'.²⁹ But to know that the cup is chipped is to know something about the external world. So if what I see is *that the cup is chipped*, then it is not possible for me to perceive what I perceive in this case without thereby knowing something about the external world.³⁰ Seeing that the cup is chipped is precisely a *means* of knowing, or of coming to know, something about the external world. The knowledge that it yields isn't 'epistemically prior' to the knowledge that the cup is chipped; it *is* the knowledge that the cup is chipped. It is false, therefore, that our knowledge of the world is always underdetermined by what we get through the senses.³¹

Seeing that something is the case is a form of what Dretske calls *epistemic seeing*.³² In epistemic seeing one sees *that* something is the case, so epistemic seeing is propositional. One can't see that A unless one grasps the proposition A, and one can't grasp this proposition unless one grasps its constituent concepts.³³ Epistemic seeing is therefore conceptual; I can't see that the cup is chipped unless I have the concepts *cup* and *chipped*. Finally, a situation in which one sees that A is 'a type of situation which represents *the acquisition of knowledge by visual means*' (Dretske 1969: 80). So the sense in which seeing that p is *epistemic* is that it embodies or involves an epistemic achievement, the acquisition of knowledge. The contrast is with non-epistemic seeing or 'simple' seeing. One can simply see the chipped cup without thereby knowing or believing that it is chipped. In order to see the chipped cup, one doesn't need to have concepts like *cup* or *chipped*, and one doesn't need to grasp any proposition.³⁴

Seeing that the cup is chipped is obviously not the only way of coming to know that it is chipped. One can also come to know that it is chipped

²⁹ This point has been emphasized by, among others, Dretske, Williamson, and Stroud. See Dretske 1969: 124, Williamson 2000: 37, and Stroud 2004: 167.

³⁰ I am taking it, in other words, that what I perceive in this case is nothing short of the fact that the cup is chipped.

³¹ This is the gist of what Stroud calls the 'most straightforward answer' to (HP_{pk}) and, by implication, to (U). This answer says that 'one *sees* that there is a table in the room, not that one infers that there is a table in the room from something else. And to see that p is to know that p. Whoever sees that p thereby knows that p. Whoever sees that there is a table in the room knows that there are external things' (2004: 167).

³² Dretske gives a detailed account of epistemic seeing in his 1969: 78–139.

³³ Williamson also makes this point. To quote his example, 'a normal observer in normal conditions who has no concept of chess can see a situation in which Olga is playing chess, by looking in the right direction, but cannot see *that* Olga is playing chess, because he does not know what he sees to be a situation in which Olga is playing chess' (2000: 38).

³⁴ See Dretske 2000 for an account of simple seeing.

by running one's finger along its rim, or by being told that it is chipped. Seeing that A, feeling that A, hearing that A, and so on are different forms of what might be called *epistemic perception*, and it seems that we can now explain how perceptual knowledge is possible by drawing on the possibility of epistemic perception. In so far as seeing that A is a way of coming to know that A, perceptual knowledge of the external world is possible by *means* of epistemic seeing or, more generally, by means of epistemic perception. This is effectively a Level 1 or a Means Response to (HP_{pk}). And if seeing that A is a way of coming to know that A by visual means, then there is nothing that stands in the way of our coming to know things about the external world by means of the senses. This is now a Level 2 response to (HP_{pk}).

In so far as the appeal to epistemic perception is intended as a response to Level 2 (HP_{pk}), is it successful? Does it really dissipate the alleged obstacle to our coming to know things about the external world by means of the senses? On the face of it, there are several reasons why one might fail to be convinced by this attempt at obstacle dissipation. For example:

- (a) One might think that it is just false that 'S sees that the cup is chipped' entails 'S knows that the cup is chipped'. Thus, even if what S sees is that the cup is chipped, what he knows by visual means still falls short of knowledge of the world.
- (b) One might deny that it is possible for S to see that the cup is chipped; if S could really *see* that the cup is chipped, there would be nothing which stands in the way of S's knowing that the cup is chipped, but this observation leaves (U) and the doctrine of epistemic priority untouched if there is no such thing as epistemic seeing.
- (c) One might agree that S can see that the cup is chipped and that 'S sees that the cup is chipped' entails 'S knows that the cup is chipped', but still deny that this constitutes an adequate answer to (HP_{pk}); the worry here is that we haven't explained how knowledge of the world is possible just by identifying various different means by which it is possible, such as epistemic perception.

Let's consider these concerns in turn. With regard to (a), it is easy to understand why the entailment from 'S sees that the cup is chipped' to 'S knows that the cup is chipped' might appear problematic. The standard example in this connection is that of a subject who sees that

the cup is chipped but who still doesn't know that the cup is chipped because he mistakenly believes that his senses are malfunctioning.³⁵ If S believes that his senses are malfunctioning then he might very well not believe that the cup is chipped, and if he doesn't believe that the cup is chipped then he doesn't know that it is chipped. But he still sees that the cup is chipped. That is why seeing that the cup is chipped doesn't entail knowing that it is chipped.

Those who think that believing isn't a condition for knowing needn't be worried by such examples because they can describe them as ones in which S sees and thereby knows that the cup is chipped without believing that the cup is chipped.³⁶ But there are also less radical ways of defending the view that seeing that the cup is chipped is a way of knowing that it is chipped. For if S mistakenly believes that his senses are malfunctioning then it is arguable that he fails to see that the cup is chipped. He fails to see that the cup is chipped because he fails to satisfy an intuitive condition on seeing that the cup is chipped: he does not believe the conditions under which he sees the cup are such that the cup would not look the way it looks to him now unless it was chipped.³⁷ So we don't have a case in which S sees but does not know that the cup is chipped. What we have is a case in which S neither sees nor knows that the cup is chipped.

Perhaps, in that case, the problem is, as (b) suggests, that S cannot see that the cup is chipped. One possibility is that there are contingent obstacles which stand in the way of S's seeing any such thing. For example, S might be blind or it might appear to be too dark for S to see anything. If, in such circumstances, S were to assert that the cup is chipped, the question 'how do you know?' would have an obvious point. As Austin points out, 'how' questions are often asked *pointedly*,

³⁵ This is McDowell's example. See his 1998c: 390 n. 37. He implies that in such cases the subject is in a position to know that the cup is chipped without actually knowing that the cup is chipped. But if one can see that the cup is chipped without actually knowing that it is chipped then it is hard to see how seeing that *p* can be a way of knowing that *p*.

³⁶ This is Williamson's view. He claims that cases such as the one I have described 'put more pressure on the link between knowing and believing or having justification than they do on the link between perceiving or remembering and knowing' (2000: 38).

³⁷ This is one of Dretske's conditions on epistemic seeing. A subject S sees that *b* is P in a primary epistemic way only if 'the conditions under which S sees_n *b* are such that *b* would not look, L, the way it looks now to S unless it was P' (1969: 82). The subscript makes it clear that non-epistemic or simple seeing is a component of epistemic seeing. See below, 1.4, for more on the distinction between primary and secondary epistemic seeing.

the implication being that the person to whom the question is directed doesn't really know what he claims to know because he is not in a position to know.³⁸ Yet the assumption that S isn't in a position to know might be mistaken. Maybe S isn't blind; perhaps it's just light enough for S to see.

Presumably, those who say that it is not possible for S to see that the cup is chipped do not say this on account of the existence of such contingent obstacles. The obstacles to epistemic seeing which they have in mind are *epistemological*, and they are allegedly ones which not only prevent S from seeing that the cup is chipped but which also prevent any other person from simply perceiving that something external is the case and thereby knowing that it is the case. For example, it might be claimed that S cannot correctly be said to see that the cup is chipped unless he can eliminate the possibility that he is dreaming or hallucinating, and that these are not possibilities which we are ever in a position to eliminate. On this account, the appeal to the possibility of epistemic perception fails to dissipate the ultimate obstacle to perceptual knowledge since this obstacle is also an obstacle to epistemic perception. Knowing that one is not dreaming at the relevant time is a necessary condition for knowing about the world by means of the senses, and the doctrine of epistemic priority is simply a consequence or reflection of the fact that this is not an epistemological requirement that can be met.³⁹

Faced with this argument, there are basically two ways of pursuing a Level 2 response to (HP_{pk}). One would be to argue that the requirement which calls the possibility of epistemic perception into question is spurious; for S to see that the cup is chipped it must be *true* that S isn't dreaming or hallucinating, but this doesn't mean that S cannot correctly be said to see that the cup is chipped unless he can eliminate these possibilities. The elimination of these possibilities has a bearing on whether S knows that he sees that the cup is chipped, not on whether he actually sees and thereby knows that the cup is chipped. The other approach would be to accept the epistemological requirement but argue that it can be met because one can know that one is not dreaming. According to McDowell, for example, 'one's knowledge that one is

³⁸ See Austin 1979: 78 for an account of the pointed use of 'how' questions.

³⁹ This line of argument is set out in Stroud 1984 and 2000e. Take a case in which any one of us would think that there is a fire in the fireplace right before us, and that we know that it is there because we see that it is there. The introduction of 'alternative, uneliminated possibilities' leads us to conclude that 'whatever we see to be so in that case, we do not simply see *that there is a fire there*' (Stroud 2000e: 131).

not dreaming in the relevant sort of situation owes its credentials as knowledge to the fact that one's senses are yielding one knowledge of the environment—something that does not happen when one is dreaming' (1998b: 238). So if one knows by means of one's senses that the cup is chipped, then one also knows that one is not dreaming.

It might seem that the second of these two approaches fits better with the idea of an obstacle-overcoming response to (HP_{pk}) than with that of an obstacle-dissipating response. Doesn't acceptance of an obstacle-generating epistemological requirement only leave one with the option of trying to overcome it? And doesn't the suggestion that the requirement can be met just amount to the suggestion that the obstacle can in fact be overcome? In fact, this isn't quite right. I have been assuming that an obstacle-generating epistemological requirement is just the requirement that:

(D1) In order to know anything about the world by means of the senses one must know that one is not dreaming.

But even if this is a requirement that can be met in the way that McDowell suggests, there is another, more demanding requirement that can't be met in this way. This is the requirement that:

(D2) In order to know anything about the world by means of the senses one must know that one is not dreaming *independently* of knowing whatever it is that one takes oneself to know about the world by means of the senses.⁴⁰

While it might be acceptable to claim conformity to (D1) on the basis that one's senses are in fact giving one knowledge of the environment, it is obviously unacceptable to claim conformity to (D2) on this basis. To attempt to rule out the possibility that one is dreaming on this basis would not be to provide independent grounds for thinking that one is not dreaming.

Unlike (D1), therefore (D2) represents an obstacle to the acquisition of knowledge of the world which can't be met or overcome. This strengthens the case for giving up on the idea of a viable obstacle-overcoming response to (HP_{pk}) . If the obstacle is (D2) rather than (D1), only an obstacle-dissipating response has any chance of succeeding. The position so far is that there is a certain epistemological requirement which can't be met on its own terms and which, if it were genuine,

⁴⁰ Cf. McDowell 1998b: 238–9.

would make the acquisition of knowledge of the world by means of the senses impossible. So what needs to be done to keep alive the possibility of perceptual knowledge is to show that (D2) isn't a genuine requirement. By showing that (D2) doesn't state a necessary condition for the acquisition of knowledge by means of the senses, one would be dissipating rather than overcoming an obstacle to epistemic perception. Having dissipated the alleged obstacle, one could then reasonably continue to insist that seeing that the cup is chipped is a viable means of coming to know that the cup is chipped.

But how does one decide whether a potentially obstacle-generating epistemological requirement such as (D2) is genuine? In deciding whether to accept a particular requirement one factor to take into account is whether it strikes one as intuitively plausible. A related consideration is whether imposition of the requirement would have acceptable consequences. The consequences that are at issue here are epistemological. Epistemological requirements have epistemological consequences, and acceptable epistemological requirements mustn't have unacceptable epistemological consequences. The problematic cases are ones in which a requirement that initially strikes us as plausible turns out to have unacceptable consequences. In such cases, the fact that the requirement has unacceptable consequences needn't deprive it of whatever plausibility it seemed to have before we registered its consequences, so we have either got to live with the consequences or reject a requirement that continues to strike us as plausible.

We now have to explain what it would be for the epistemological consequences of a principle like (D2) to be unacceptable. Here is one proposal: in any serious investigation of the conditions of knowledge, we start off with the idea that there are certain things that we know, or certain kinds of knowledge that we actually have. We regard some of the knowledge that we take ourselves to have as negotiable and some as more or less non-negotiable. An unacceptable consequence of an epistemological principle would be the undermining of knowledge in the latter category. For example, even if we take ourselves to have some insight into our own motives and desires, it doesn't necessarily count against a theory of mind or knowledge that it implies that many of our beliefs about our deepest motives and desires don't amount to knowledge. Self-knowledge is, to this extent, negotiable. In contrast, there are certain basic forms of perceptual knowledge that are non-negotiable. An example of non-negotiable perceptual knowledge might be my knowledge that the cup into which I am pouring coffee is

chipped. The presumption is that I do know this by means of the senses, so it would be an unacceptable consequence of a principle like (D2) if it implies that this is something I couldn't know; the very fact that it carries this implication would be a strike against it.⁴¹

Clearly, there is a difference between a piece of knowledge being more or less non-negotiable and its being absolutely non-negotiable. Presumptions to knowledge, even presumptions to basic perceptual knowledge, are, in principle, capable of being defeated, but we shouldn't underestimate what it would take to defeat such a presumption. What we would need is an epistemological principle that is at odds with the presumption *and* that has such overwhelming independent plausibility that even the fact that it threatens to undermine what we previously took to be non-negotiable knowledge doesn't warrant the principle's rejection. This is not how things stand with (D2). Whether or not it is overwhelmingly plausible that in order to know anything about the world by means of the senses one must know that one is not dreaming, it is not overwhelmingly plausible that 'the epistemic status of the thought that one is not dreaming must be established independently of the epistemic status of whatever putative perceptual knowledge of the environment is in question' (McDowell 1998b: 225). In other words, (D2) is negotiable to an extent that the perceptual knowledge whose possibility it calls into question is not, and the fact that the truth of (D2) would make the acquisition of knowledge of the environment by means of the senses impossible is itself a good reason for refusing to endorse this principle. Given that denial of basic perceptual knowledge is an unacceptable consequence of (D2), the right thing to think is that (D2) doesn't state a genuine requirement on perceptual knowledge.

This attempt at obstacle dissipation might be described as Moorean. In particular, it is reminiscent of what has been described as Moore's 'argument from differential certainty'.⁴² According to this argument, we are much more certain of our basic perceptual knowledge than we are of the premisses of any sceptical argument for the impossibility of perceptual knowledge. As Moore recognizes, however, the question which this argument raises is whether the certainty to which it appeals is objective or merely subjective. What it seems to require is the objective

⁴¹ Moore argues in this way in some of his writings on scepticism. See, for example, 1953: 121–2.

⁴² This is what Baldwin calls Moore's argument. See Baldwin 1990: 269–74 for a critical discussion of this argument.

certainty or indubitability of one's perceptual knowledge. Yet this is not something that Moore establishes. This shows that the argument from differential certainty has its limitations as an anti-sceptical argument, but there is another way of looking at it. Instead of seeing it as an attempt to refute scepticism, we might simply see it drawing attention to the fact that the epistemological principles we endorse are answerable to our epistemological verdicts in particular cases, and that conflicts between particular verdicts and general principles needn't always be decided in favour of the latter. If we are so strongly committed to a particular verdict as to regard it as non-negotiable, we can assess a general principle by testing its compatibility with that verdict. What is at issue here is *commitment* rather than certainty, and the question is whether our commitment to thinking that we have some basic perceptual knowledge is, as I have been arguing, stronger than our commitment to (D2).

The sceptic's worry is, of course, that our commitments to particular verdicts might be irrational or unfounded, and that we might therefore be wrong to regard our possession of basic perceptual knowledge as non-negotiable. Since there is no general guarantee that our epistemological commitments are well founded, it isn't going to be possible to prove to the sceptic's satisfaction that we are actually in possession of any perceptual knowledge. It's fortunate, therefore, that the obstacle-dissipating response to (HP_{pk}) isn't attempting to prove any such thing. The object of the exercise is simply to explain how perceptual knowledge is possible, *given* that it is possible. By showing that we needn't be committed to obstacle-generating principles such as (D2), we are doing all that needs to be done at Level 2 to secure the possibility of perceptual knowledge. In so far as we have made it plausible that (D2) doesn't state an inescapable requirement on perceptual knowledge, we have left it open that epistemic perception is a possible means of coming to know things about the world around us. We have thereby rebutted the suggestion in (b) that S can't really see that the cup is chipped or know that it is chipped by seeing that it is.

That leaves (c), according to which we still haven't fully explained how perceptual knowledge is possible even after we have both identified epistemic perception as a means of coming to know about the external world and shown that there is nothing that stands in the way of epistemic perception. So what more is required? Just as in the geometrical case the anti-minimalist thinks that what is required is an explanation of what makes it possible for construction in pure intuition to occur and to be a source of synthetic a priori knowledge, so the parallel suggestion

in relation to (HP_{pk}) is that the proposed Means Response to this how-possible question needs to be supplemented by answers to two what-makes-it-possible questions. The first is: what makes it possible to perceive that something is the case? The second is: what makes it possible for perceiving that something is the case to be a means of coming to know that it is the case? These are questions about enabling conditions. An answer to the first question will need to identify type A enabling conditions, that is, the background necessary conditions for the occurrence of epistemic perception. An answer to the second question will need to identify type B enabling conditions, that is, the background necessary conditions for epistemic perceiving to be a source of knowledge. The identification of type A and type B enabling conditions is a Level 3 explanation of the possibility of perceptual knowledge, and the point of (c) is to suggest that an adequate answer to (HP_{pk}) must reach all the way down to this level.

This is effectively an argument for extreme explanatory anti-minimalism in relation to (HP_{pk}), so the issue is whether this form of anti-minimalism is warranted. We need to consider whether Level 3 explanations of the possibility of perceptual knowledge are possible and, if so, whether they are necessary. If they aren't possible then this would count against (c) and in favour of explanatory minimalism. If they are possible but not necessary this would count against (c) and in favour of moderate anti-minimalism. This is the approach for which I want to argue. In opposition to minimalism, this approach allows for the possibility of Level 3 explanations of perceptual knowledge and does not dispute the legitimacy of asking what makes it possible for epistemic perception to occur and to be a source of knowledge. On the other hand, there is no clear sense in which explanations of the possibility of perceptual knowledge that stop at Level 2 are 'incomplete', so there isn't a case for extreme anti-minimalism.

To get a sense of the force of anti-minimalism and of what (c) represents as the limitations of minimalism, it's essential that some of the details of an anti-minimalist account of epistemic perception are filled in. It isn't enough for the anti-minimalist to insist that there are legitimate what-makes-it-possible questions about perceptual knowledge that a philosophically satisfying answer to (HP_{pk}) can or should address. Some indication also needs to be given of what a good answer to these questions might look like in practice, and of how 'philosophical' Level 3 explanations of the possibility of perceptual knowledge differ from other Level 3 explanations. The next challenge, therefore, is to give a sketch

of some of the type A or type B enabling conditions that might figure in a fully-fledged multiple levels response to (HP_{pk}). Having done this, we will be in a better position to assess the relative merits of moderate and extreme anti-minimalism.

1.4 ANTI-MINIMALISM

What makes it possible to see that the cup is chipped? In the primary sense of seeing that the cup is chipped, this is not something that one could see without seeing the cup. So as long as we stick with the primary use of constructions of the form 'S sees that b is P', we now have the proposal that S can't see that b is P without seeing b. Obviously, not all uses of the 'sees that' construction work like this; one doesn't see that the cup is missing by seeing the cup. This points to the need for a contrast between *primary* and *secondary* epistemic seeing, that is, a contrast between 'the cases where we see that b is P *by seeing b* itself, and the cases where we see that b is P *without seeing b*' (Dretske 1969: 79–80).⁴³ Primary epistemic seeing is the more fundamental or basic form of epistemic seeing since without it there would be no secondary epistemic seeing either, and no possibility of coming to know that b is P by seeing that b is P. For this reason I'm going to concentrate on the following question: what are the enabling conditions for primary epistemic seeing or for primary epistemic perceiving?

Suppose, then, that I see that the cup is chipped by seeing the cup. If there are enabling conditions for seeing the cup then they are also going to come out as enabling conditions for seeing that the cup is chipped. But cups are objects, and seeing the cup is therefore an example of (visual) *object perception*. Indeed, cups are not just objects but specifically *material* objects, in other words, bounded, three-dimensional space-occupiers. So as long as we are thinking of a case of primary epistemic seeing, the enabling conditions for seeing that the cup is chipped will include any background necessary conditions for seeing material objects. More

⁴³ As Dretske points out, this account of the distinction between primary and secondary epistemic seeing isn't quite accurate because there are cases in which one sees that *b* is P 'in virtue of the way other objects look or behave when *b* is P' (1969: 153). These are cases of secondary epistemic seeing even if they are ones in which one sees *b* itself. The reason is that seeing *b* itself is incidental. For example, when I insert a toothpick into the middle of a cake I see the cake but it is the way the *toothpick* looks that enables me to see that the cake is done.

generally, given that *seeing* a material object is not the only way of *perceiving* one, we now have the proposal that one way, though not the only way, of figuring out what makes primary epistemic perceiving possible is to figure out what makes the perception of material objects possible.

This proposal is along the right lines, but it needs to be qualified in the following respects: to begin with, while it is true that some epistemic seeing involves the seeing of material objects, it is false that in every case in which one sees that *b* is *P* by seeing *b* itself *b* itself is a material object. At a baseball game one sees that the shadow of a low-flying plane is moving rapidly across the stadium, and one sees that the shadow is moving rapidly across the stadium by seeing the shadow. Yet shadows aren't material objects, even if they are objects in some looser sense of 'object'. So not every case of object perception is a case in which what is perceived is a material object.

A further complication is that there are plenty of cases in which one sees that *b* is *P* by seeing *b* itself but in which *b* itself isn't an object at all, not even an object in a loose sense of 'object'. For example, *b* might be an event; one sees that a game of baseball is in progress by seeing the game but games are temporally extended events rather than objects. What makes it possible to see that *b* is *P* by seeing *b* itself will therefore include what makes it possible to see *b* itself for many different types of '*b*'. This could make things very complicated but it is important at this stage to keep things as simple as possible. That is why it makes sense to concentrate, at least to begin with, on the simplest possible case of primary epistemic seeing, the case in which one sees that *b* is *P* by seeing *b* itself, and in which *b* itself is a material object. This isn't just a case of primary epistemic seeing but a case of what I'm going to call *basic* primary epistemic seeing. Seeing that the cup is chipped is a case of this kind. Later, I will consider the consequence of lifting the restriction to material objects as well as the consequences of lifting the restriction to visual perception.

We are now looking for an account of what makes basic epistemic seeing possible, and the present suggestion is that the enabling conditions for seeing that *b* is *P* will include the enabling conditions for seeing *b* itself. The enabling conditions for seeing *b* itself will be type *A* enabling conditions for seeing that *b* is *P*, that is, conditions which must be fulfilled for primary epistemic seeing to occur. What are the enabling conditions for seeing *b* itself in the case in which *b* itself is a material object? Some of these conditions have to do with the physical

environment. For example, it must be light enough for *b* to be seen and there mustn't be anything blocking the perceiver's view of *b*. Others have to do with the workings of the perceiver's cognitive apparatus. Unless one's eyes and brain are functioning properly, one wouldn't be able to see *b* or, for that matter, anything else. So the enabling conditions for the perception of objects by sight include physiological conditions as well as environmental conditions.

Physiological and environmental enabling conditions are causally necessary conditions. From a naturalistic perspective, the enabling conditions for object perception are always causally necessary conditions, and it is reasonable to suppose that there are many such conditions. It also seems a reasonable assumption that causally necessary conditions can only be discovered empirically. The implication is that the project of explaining what makes it possible to see a material object, and therefore the project of uncovering enabling conditions of basic primary epistemic seeing, can't be completed by armchair philosophy. On the contrary, it now appears that it is the business of empirical science to reveal the background necessary conditions for seeing that *b* is *P* by seeing *b* itself.

On this account, there is nothing wrong with the project of explaining what makes primary epistemic seeing possible as long as this project is conceived of naturalistically. But once it is conceived of in this way, it's no longer obvious that philosophy has much to contribute to it. In a way, therefore, this can be seen as a vindication of a kind of explanatory minimalism; the thought is that *philosophical* explanation comes to an end at Level 2, and that any remaining questions about the means by which it is possible for us to come to know things about the external world are scientific questions. Philosophy can get one as far as the idea that a Level 3 explanation of what makes it possible to see that *b* is *P* will need to incorporate an explanation of what makes it possible to see *b* but it can't explain what makes it possible to see *b*.

This alliance between naturalism and minimalism shows that a defensible anti-minimalism will need to do more than insist on the necessity or possibility of explaining what makes perceptual knowledge possible. I'm taking it that an anti-minimalist is someone who thinks that distinctively philosophical explanations of the possibility of perceptual knowledge are necessary or at least possible. As far as basic epistemic seeing is concerned, therefore, the issue isn't whether there are questions about its causal enabling conditions that can only be answered empirically. That is not in dispute. The issue is whether, as the anti-minimalist insists, there are other questions about what makes this

kind of seeing possible that can, or can only, be answered by means of a priori philosophical reflection. Minimalism is only in trouble if there are such questions.

Suppose, then, that we say that enabling conditions that can *only* be established by some form of a priori reflection are *strongly* a priori conditions. *Weakly* a priori conditions are ones that *can* be established without empirical investigation. In other words, it isn't written into the very idea of a weakly a priori enabling condition that such conditions can't also be established by empirical means. In these terms, minimalism can be understood as denying that there are any strongly *or* weakly a priori enabling conditions for this kind of knowledge. That is why, on the assumption that what is distinctive of philosophical explanation is that it is non-empirical, minimalism thinks that all such explanation comes to an end at Level 2. If this is right then there is no need for anti-minimalists to demonstrate that there are strongly a priori enabling conditions for perceptual knowledge. In order to undermine minimalism they only need to make it plausible that there are enabling conditions that *can* be established without any empirical investigation.

The particular version of moderate anti-minimalism that I want to explore is Kantian in inspiration. To get a flavour of it, let's go back to the example of seeing that a particular cup is chipped by seeing the cup itself. Usually when one sees a cup one doesn't just see the cup. Typically, the cup is one among a range of things that one also sees, and seeing the cup involves being able to differentiate or distinguish it from these other things. If there are other cups in one's field of vision, one must be able to differentiate one cup from another. If one is holding the cup, one must be able to differentiate it from one's hand, and so on. The required differentiation is visual, and anything that is a background necessary condition for visual differentiation is also going to come out as a background necessary condition for the perception of objects by sight. So if there are enabling conditions for the visual differentiation of objects that can be established by armchair reflection, without any empirical investigation, then this would count against what I have been calling explanatory minimalism.

In arguing in this way, one is not committed to thinking that it isn't possible to see an object without differentiating it from its immediate surroundings. Imagine seeing an object *b* which is in contact with another similar object *c*. When one sees *b* one fails to see it as distinct from *c*; in this sense one fails to differentiate *b* from its surroundings but this leaves open the possibility that what one sees in this case is *b*.

At the very least, further work needs to be done to make it plausible that a failure visually to differentiate *b* from *c* amounts to a failure to see *b*.⁴⁴ It's clear, however, that paradigmatic cases of object perception are ones in which what one perceives is differentiated from its surroundings. For example, when one is pouring coffee into a cup it is perceptually manifest to one where the cup ends and the rest of the world begins; that is why the coffee ends up in the cup. So does a priori philosophy have anything useful to say about the background conditions under which this kind of differentiating object perception is possible?

Kant's proposal is that the perception of space is a background necessary condition for visual object perception. The 'perception of space' is a cognitive capacity, the capacity to perceive spatial properties such as shape and location. Kant's idea is that possession of this cognitive capacity is an enabling condition for the perception of objects by sight because it is an enabling condition for visual differentiation. If this claim is correct, and if seeing that *b* is *P* involves visually differentiating *b* from its surroundings, we can conclude that the perception of space is a type A cognitive enabling condition for seeing that *b* is *P*. By making it possible for one to see *b*, the perception of space makes it possible for one to see that *b* is *P*. Kant also thinks that the link between the perception of space and visual differentiation can be established non-empirically. That is why, on his view, the perception of space is not just an enabling condition but an a priori enabling condition for differentiating visual perception. The identification of this condition therefore serves as an example of what philosophy, as distinct from empirical science, can achieve in this area.

Why should one think that the perception of space is a background necessary condition for differentiating visual perception? One thought is that in order to see an object *b* as differentiated from another object *c* one must see *b* and *c* as being in different *places*.⁴⁵ The perception of place, and therefore of space, serves as the *means* by which *b* is differentiated from other things in its environment. But one couldn't see where *b* is unless one has the capacity to perceive spatially. So possession of this capacity comes out as a background necessary condition for differentiating perception as long as it is also true that the perception of place is not just *a* means but the *only* means of differentiating *b* from its

⁴⁴ See Dretske 1969: 18–32 and Chapter 3 below for further discussion of this issue.

⁴⁵ Allison defends something along these lines in his 1983: 83. There is much more on this issue in Warren 1998 and Chapter 3 below.

environment. If differentiating perception needn't be a form of spatial perception, then one would need to find other reasons for thinking that the perception of space is a cognitive enabling condition for seeing that *b* is *P*.

Is Kant's conception of the role of spatial perception in primary epistemic seeing defensible? I will have much more to say about this in Chapter 3. For example, one might wonder whether the perception of place can really be as important for visual differentiation as the thesis suggests. When I pour myself a cup of coffee, I certainly see the cup as distinct from my hand, yet it is implausible that I see them as distinct things by seeing them in different places. I see them as distinct because they *look* distinct, and looking distinct in this sense seems to have little to do with the perception of a difference in location. When senses other than sight are considered, the problem is even more acute. One can hear two people arguing in the next room as distinct from each other just on the basis of their voices. The perception of a difference in location has little to do with it, since I might fail to hear them as being in different places in the next room.

These observations suggest that Kant's thesis faces some formidable challenges. It is worth emphasizing, however, that it could be true that one can't see an object without perceiving any of its spatial properties even if it is false that the perception of space is the key to visual *differentiation*. Where the object is a material object, one might think that one couldn't see it *at all* without seeing it as shaped, located, or extended in space. So the perception of space could still be a background enabling condition for the perception of *b* itself regardless of whether it is the key to differentiating *b* from other objects. If this is right, there would again be a question about the role of space in the perception of objects by non-visual modes of perception, and about the extent to which it is an empirical question whether one can see a material object without seeing any of its spatial properties. For the moment, we can just bracket these questions. Right now all I am trying to do is to give an illustration of what might count as a type A enabling condition for primary epistemic seeing, on the assumption that primary epistemic seeing is a means of coming to know things about the world around us.

Are there any other enabling conditions for primary epistemic seeing which have not yet been mentioned? When I introduced the notion of epistemic seeing I said that this form of seeing is conceptual; for example, in order to see that the cup is chipped one must have the concepts *cup* and *chipped*. These are examples of empirical concepts, that

is, concepts that have their source in experience or that can be derived from experience. Hence, if there are background necessary conditions for the possession or acquisition of empirical concepts, one would be entitled to regard them as enabling conditions for seeing that the cup is chipped. More generally, given that concepts are necessary for epistemic seeing, one would expect a Level 3 explanation of what makes it possible for one to see that *b* is *P* to incorporate an account of what makes it possible for one to have concepts like *b* and *P*.

What are the background conditions under which the possession or acquisition of empirical concepts is possible? Again, the naturalist or minimalist claims that this is best understood as an empirical question about causal enabling conditions. Such conditions might be physiological or biological in nature, and it is not for armchair philosophy to tell us what they are. In contrast, the anti-minimalist denies that Level 3 questions about the concepts that figure in epistemic seeing are necessarily scientific or empirical, even though some of them undoubtedly are. As far as the anti-minimalist is concerned, there are enabling conditions for the possession and acquisition of empirical concepts that can be discovered by armchair philosophical reflection and that are therefore at least weakly if not also strongly a priori. Again, the intended upshot is that philosophical explanation needn't, and perhaps shouldn't, come to an end at Level 2.

One distinctively philosophical Level 3 proposal is that in order to have concepts like *cup* and *chipped* one must have lots of other concepts. Possession of a network of interrelated concepts is a background necessary condition for the possession of individual concepts like *cup* and *chipped*. While it's not uncontroversial whether this holistic constraint on concept possession is correct, the anti-minimalist's idea is that its correctness or otherwise can't be settled empirically; only a philosophical theory of concepts can do that. And the same goes for another putative enabling condition for concept possession. On a linguistic conception of concepts, only a creature with a language can have concepts. Concepts must be exercisable in judgements, and if the capacity to judge depends on language then so do the concepts that figure in one's judgements. If this is right, mastery of a language is required in order to see that the cup is chipped.

I'm not going to have very much to say about these alleged holistic and linguistic constraints on concept possession. Instead, I'm going to concentrate on Kant's conception of what makes it possible for one to acquire and possess concepts like *cup* and *chipped*. It's worth noticing

that the holistic and linguistic constraints aren't just constraints on empirical concepts; for example, if it is in the nature of concepts to presuppose mastery of a language, then this is presumably going to be true of all concepts, empirical or otherwise. The Kantian proposal which I will be discussing makes more of the distinction between empirical and non-empirical concepts. In so far as empirical concepts are ones that have their source in experience or that can be derived from experience, non-empirical or a priori concepts are ones that don't have their source in experience or that can't be derived from experience. Kant's inventory of a priori concepts includes the so-called 'categories' or 'pure concepts of understanding'. His thesis is that the categorial thinking, thinking by means of concepts like *substance*, *unity*, *plurality*, and *causality*, is an a priori enabling condition for the possession and acquisition of empirical concepts like *cup* and *chipped*.⁴⁶ This kind of thinking makes it possible for one to have and acquire such concepts, and therefore also makes epistemic seeing possible. Without the categories in the background, one couldn't have or acquire concepts like *cup* and *chipped*, and without these concepts one couldn't see that the cup is chipped.

As we will see in Chapter 4, this Kantian explanation of the possibility of empirical concepts and, by extension, of the possibility of epistemic seeing faces a range of challenges and objections that are no less formidable than those facing Kant's Level 3 explanation of the possibility of object perception. The main worry isn't that Kant is attempting to establish by means of a priori reflection claims that ought to be established empirically. The worry is rather that it's far from obvious that Kant's claims about the link between empirical concepts and categorial thinking are actually correct. In the end, I will argue that only very watered down versions of these claims have any chance of being defensible. They are nevertheless worth discussing both because of their historical interest and because they provide an excellent illustration of what happens when anti-minimalism is taken too far. Kant's thinking about these matters is underpinned by an extreme form of anti-minimalism, and reflecting on the failings of this approach will make the virtues of the moderate anti-minimalism that I want to defend very much clearer than they would otherwise be.

So much for type A enabling conditions of epistemic seeing. What is there to say about its type B enabling conditions? An account of these conditions would need to explain what makes it possible for seeing that

⁴⁶ Longuenesse attributes something along these lines to Kant. See Longuenesse 1998.

something is the case to be a way of coming to know that it is the case. The minimalist says that all that can be done in this connection is to point out that 'S sees that b is P' entails 'S knows that b is P'. There is nothing that 'makes it possible' for 'S sees that b is P' to entail 'S knows that b is P', so it doesn't make sense to look for type B enabling conditions for epistemic seeing. If explaining how epistemic seeing can be a source of knowledge is a matter of explaining how 'S sees that b is P' can entail 'S knows that b is P', then minimalism denies that any such explanation is either necessary or possible.

This form of minimalism is related to a more general form of minimalism in epistemology. In *Knowledge and its Limits*, for example, Williamson argues that the project of trying to fix non-circular necessary and sufficient conditions for propositional knowledge is doomed, and that the concept *knows* cannot be analysed into more basic concepts.⁴⁷ Instead he gives a modest positive account of this concept according to which 'if one knows that A, then there is a specific way in which one knows; one can see or remember ... that A' (2000: 34). 'Sees that ...', 'remembers that ...', and 'knows that ...' are all examples of *factive mental state operators* (FMSOs). If ϕ is an FMSO, 'S ϕ s that A' entails 'S grasps the proposition that A', and the inference from 'S ϕ s that A' to 'A' is deductively valid. In addition, ϕ is semantically unanalysable. In these terms, the proposal is that if ϕ is any FMSO, then 'S ϕ s that A' entails 'S knows that A. If you see that it is raining, then you know that it is raining. If you remember that it was raining, then you know that it was raining' (2000: 37). The implication of this account of the sense in which 'seeing that A is a way of knowing that A' (2000: 38) is that the concept *knows* can effectively be characterized as the determinable of which such specific ways of knowing are the determinations.

Part of what makes this a form of epistemological minimalism is that it offers no analysis or explanation of the link between seeing that A and knowing that A. In contrast, anti-minimalism does seek to explain this link. It points out, for example, that seeing that A isn't a way of knowing that A in the sense in which remembering that A is a way of knowing that A. Seeing that A is a way of *coming to know* that A, a way of *acquiring* this knowledge, whereas remembering that A is a way of *retaining* the knowledge that A. To explain the link between seeing that A and knowing that A would therefore be to explain how

⁴⁷ See Williamson 2000: 27–33.

seeing that A can be a way of coming to know that A. This is where type B enabling conditions come into the picture. The anti-minimalist's proposal is that we understand how epistemic or 'factive' seeing can be a source of knowledge by identifying the background conditions that must be met for seeing that something is the case to be a way of coming to know that it is the case. Once we have identified these conditions, we might also find ourselves being able to say something illuminating about why 'S sees that A' entails 'S knows that A'.

What are the background necessary conditions for seeing that A to be a way of coming to know that A? Are these conditions any different from the enabling conditions for seeing that A? We can start to make some progress with these questions by drawing on Dretske's account of epistemic seeing in *Seeing and Knowing*. Dretske describes himself as aiming to provide 'an analytic description of those states of affairs which are described by statements of the form "S sees that b is P" *in so far as they tell us how S knows that b is P*' (1969: 81). The account proceeds by specifying necessary and sufficient conditions for epistemic seeing that are adequate to the epistemic implications of this mode of perception. Thus, for S to see that b is P it is necessary and sufficient that (i) b is P, (ii) S sees b, (iii) the conditions under which S sees b are such that b would not look the way it looks now to S unless it was P, and (iv) S, believing the conditions are as described in (iii), takes b to be P. When conditions (i) to (iv) are fulfilled, S has in the way that b looks to him a conclusive reason for believing that b is P, and 'it is the conclusiveness of this reason which supports the entailment between "S sees that b is P" and "S knows that b is P"' (1969: 124).

This is in a way an anti-minimalist account of epistemic seeing but work needs to be done to transform it into a full-blown anti-minimalist account. It is anti-minimalist to the extent that it tries to analyse the claim that S sees that b is P and thereby to explain the entailment between 'S sees that b is P' and 'S knows that b is P'. It also employs the notion of a background enabling condition since these are the conditions that are referred to in conditions (iii) and (iv). On the other hand, Dretske makes it clear that his enabling conditions are causal enabling conditions.⁴⁸ To get to the idea of a priori enabling conditions for the acquisition of knowledge by visual means, we don't need to take a stand on the issue of whether Dretske's conditions are genuinely sufficient. For as long as it's plausible that S must see b in order to know

⁴⁸ See Dretske 1969: 82–4.

by visual means that *b* is *P*, we can proceed to identify a priori enabling conditions for epistemic seeing on this basis.

Suppose that *b* itself is an object. In that case, coming to know that *b* is *P* by seeing that *b* is *P* requires a capacity for object perception. So if there are a priori enabling conditions of object perception they will also be a priori enabling conditions for knowing that *b* is *P* by seeing that *b* is *P*. We have already briefly examined the proposal that a capacity for spatial perception is an a priori enabling condition for the perception of objects, and therefore a type A enabling condition for epistemic seeing. But anything that is an enabling condition for the perception of objects is also going to be a type B enabling condition for epistemic seeing, given that object perception is involved in the acquisition of knowledge by primary epistemic seeing. What we now have, therefore, is the possibility that the perception of space is both a type A and a type B enabling condition for epistemic seeing.

This overlap between type A and type B enabling conditions should come as no surprise. Since it is in the nature of epistemic seeing that it has epistemic implications, one would expect the background conditions under which it is possible for epistemic seeing to occur to be closely related to, if not identical with, the conditions under which it is possible for epistemic seeing to be a source of knowledge. In other words, what makes it possible for there to be such a thing as epistemic seeing can't be sharply distinguished from what makes it possible for there to be such a thing as knowing that something is the case by seeing that it is the case. There remains a notional difference between type A and type B enabling conditions, but a notional difference, a distinction at the level of sense, needn't amount to a real difference, a distinction at the level of reference.⁴⁹

We are now in a position to consider the extent to which Kant's claims about the enabling conditions of perceptual knowledge vindicate anti-minimalism. I have said that an anti-minimalist is someone who thinks that distinctively philosophical Level 3 explanations of the possibility of perceptual knowledge are necessary or at least possible. One way of showing that such explanations are possible would be to produce one. This is what Kant does or at least purports to do. If the perception of space and the categories are enabling conditions for epistemic perception that can be established non-empirically, then minimalism is wrong to

⁴⁹ This is why, in what follows, I will sometimes ignore the distinction between type A and type B conditions.

claim that distinctively philosophical Level 3 explanations of what makes perceptual knowledge possible can't be given. For minimalism to be in the running, it would have to be the case that Kant fails to identify Level 3 conditions for perceptual knowledge. This could either be because spatial perception and the categories aren't background necessary conditions for the acquisition of perceptual knowledge or because they aren't background necessary conditions for the acquisition of perceptual knowledge that can be established non-empirically. I will consider these possibilities in detail in later chapters, though it's worth pointing out that the failure of Kant's Level 3 explanations wouldn't necessarily mean that such explanations can't be given.

Suppose, then, that we are persuaded on the basis of Kant's discussion that in principle Level 3 explanations of the possibility of perceptual knowledge *can* be given. This would be a problem for minimalism, but where does it leave the debate between moderate and extreme anti-minimalism? Extreme anti-minimalists think that in the absence of a Level 3 explanation we can't reasonably claim to have explained how perceptual knowledge is possible, and that this is the sense in which Level 2 explanations aren't good enough. In contrast, moderate anti-minimalists insist that Level 3 explanations are possible while denying they are necessary. Their point is that once we have reached Level 2 and identified epistemic seeing as a means of coming to know things about the world around us we have *already* done everything that needs to be done to explain how perceptual knowledge is possible; we could go further but we don't need to.

How is this dispute to be resolved? Consider this analogy: I ask how it is possible to get from London to Paris in less than three hours and the answer I get is that it's possible to do this by catching the Eurostar. Should I be satisfied by this answer? An extreme anti-minimalist in this context is someone who thinks that I shouldn't be satisfied and that more can and needs to be done to answer the how-possible question. Catching the Eurostar is a means of getting from London to Paris, but what makes it possible to get from London to Paris by train? This is the further what-makes-it-possible question to which extreme anti-minimalism demands an answer. The answer, or at least an answer, to this question is that the existence of the Channel Tunnel is what makes it possible to get from London to Paris by train. In the absence of the Channel Tunnel, or some such link between England and France, going by train would not be a means of reaching Paris from London, so the existence of such a link is an enabling condition for getting from

London to Paris by train. According to the extreme anti-minimalist, I haven't fully understood how it is possible to reach Paris from London in less than three hours unless I recognize this enabling condition, just as I haven't fully understood how it is possible to arrive at synthetic a priori geometrical knowledge by constructing figures in pure intuition unless I recognize the ideality of space.

This is what moderate anti-minimalists find implausible. Their thought is that there is no obvious sense in which a failure to say anything about the background necessary conditions for crossing the English Channel by train constitutes a failure to give a 'complete' answer to the how-possible question. There are lots of factors that make it possible to reach Paris from London by train—the existence of the Tunnel, the existence of trains, and so on—and some of these factors can be established without empirical investigation while others can only be established empirically. Presumably I can know a priori that the existence of trains is necessary for getting anywhere by train but I can't know a priori that the existence of a cross-channel tunnel is an enabling condition for getting from London to Paris by train. Yet there is no *need* to go into any of this if all one wants to know is how it is possible to get from London to Paris in less than three hours. As far as answering the how-possible question is concerned the what-makes-it-possible question is an optional extra even though it's one to which answers *can* be given if someone insists on asking it. If I know that it's possible to get from London to Paris in less than three hours by catching the Eurostar then I know how it's possible to get from London to Paris in less than three hours.

Moderate anti-minimalism's take on (HP_{pk}) is similar to its take on the Eurostar. Again the idea is that once I understand that it's possible to know that the cup is chipped by seeing that it is chipped I understand how this particular piece of knowledge is possible. Just as nothing needs to be said about the existence of the Channel Tunnel in order to explain how it's possible to get from London to Paris in less than three hours, so nothing needs to be said about the enabling conditions of epistemic perception in order to explain how perceptual knowledge is possible. The questions that get addressed at Level 3 of the multiple levels response are perfectly legitimate and interesting but it would be wrong to claim that Level 2 explanations are, in any straightforward sense, incomplete.

None of this amounts to a knockdown argument against extreme anti-minimalism but it does put considerable pressure on defenders of this position to explain why we should be any less satisfied with

a Level 2 response to (HP_{pk}) than with a Level 2 response to many other how-possible questions, including the one about getting to Paris from London. While it's possible to get oneself into a frame of mind in which questions about enabling conditions can seem pressing it's also quite easy to see why someone might think that this question has already been answered at Level 2. If, as I have claimed, (HP_{pk}) is an obstacle-dependent question the key to answering it is to remove the various obstacles that have been alleged to stand in the way of the acquisition of perceptual knowledge. Since this is what happens at Level 2 there is no obvious need to go any further. From this standpoint, the moderate anti-minimalist's insistence that (HP_{pk}) can but needn't evolve into a what-makes-it-possible question appears to be entirely justified. The onus is on extreme anti-minimalism to make it plausible that we haven't explained how perceptual knowledge is possible unless we have explained what makes it possible, and we have so far failed to find any decisive arguments in favour of this approach.

These considerations also have a bearing on the worry that Level 3 explanations are fundamentally no different from Level 2 explanations. The worry was that the denial of a putative enabling condition is always an obstacle to knowledge, and that what goes on at Level 3 of a multi-levels response to a how-possible question is therefore just as much an exercise in obstacle-removal as what goes on at Level 2. My initial response to this worry was to argue that the point of identifying enabling conditions C for the acquisition of knowledge of kind K by means M needn't be to address any intuitive obstacle to the acquisition of K by M. If C is not fulfilled then that *becomes* an obstacle to the acquisition of K by means of M but it doesn't follow that Level 3 explanations are Level 2 explanations by another name. The Kantian account of the role of spatial perception and categorial thinking helps to make this point. There is no intuitive obstacle to epistemic seeing, to knowing that the cup is chipped by seeing that it is chipped, that is overcome or dissipated by the observation that the perception of space and categorial thinking are background necessary conditions for epistemic seeing. Without these cognitive capacities one wouldn't be able to see that the cup is chipped but what Kant is trying to do by talking about what makes perceptual knowledge possible is not to show that perceptual knowledge is not impossible; what he is after is a better understanding of the cognitive foundations of this kind of knowledge.

The lesson, once again, is that the only version of anti-minimalism to which we should be willing to commit ourselves in relation to (HP_{pk})

is moderate anti-minimalism. So the issue is not whether it's possible to answer (HP_{pk}) without identifying a priori enabling conditions for perceptual knowledge but whether perceptual knowledge has any enabling conditions that can be established non-empirically. What we need to consider, therefore, is whether spatial perception and categorial thinking are enabling conditions for the acquisition of perceptual knowledge and, if so, whether they are a priori enabling conditions. As we will see in Chapters 3 and 4, the answer to both of these questions is 'yes'. It isn't true, therefore, that armchair philosophy *can't* tell us anything about what makes perceptual knowledge possible.

But before looking in more detail at the role of space and categorial thinking in the acquisition of perceptual knowledge there is another methodological issue that needs to be addressed. I have written at length about how-possible questions in epistemology but have so far said nothing about so-called 'transcendental arguments', that is, arguments which specify necessary conditions of the possibility of thought, experience, or knowledge. Yet it has often been suggested that transcendental arguments are the best way of answering epistemological how-possible questions. So the next chapter is about transcendental arguments. We need to consider whether such arguments are any good, whether they have any bearing on questions like (HP_{pk}), and, if so, how they relate to the multi-levels approach to this question that I have been recommending.

2

Transcendental Arguments

2.1 REGRESSIVE TRANSCENDENTAL ARGUMENTS

An epistemological how-possible question asks how knowledge of some specific kind is possible. Such questions are obstacle-dependent since they are motivated by the thought that there are actual or apparent obstacles to the existence of whatever kind of knowledge is in question. One such question is:

(HP_{pk}) How is perceptual knowledge possible?

Then there is Kant's favourite how-possible question:

(HP_{sap}) How is synthetic a priori knowledge possible?

In the last chapter, I defended the view that epistemological how-possible questions call for a multi-levels response. A multi-levels response operates at three levels. Level 1 identifies *means* of acquiring the allegedly problematic knowledge. Level 2 is the *obstacle-removing* level, the level at which obstacles to the acquisition of knowledge by the proposed means are overcome or dissipated. Finally, Level 3 seeks to identify *necessary background conditions* for the acquisition of the relevant knowledge by the proposed means.

In this chapter, I want to examine the proposal that the best way to answer an epistemological how-possible question is by means of a transcendental argument.¹ Although it might appear that transcendental arguments are closely related to the multi-levels response to how-possible questions, I'm going to argue that they are different from each other and that the latter response is better. It remains to be seen whether

¹ See Hatfield 1990: 79 and Collins 1999: 91 for versions of this proposal. Although this isn't always made explicit by writers on transcendental arguments many of them appear to take it for granted that epistemological how-possible questions can be satisfactorily answered by arguments of this form.

transcendental arguments have any legitimate role in epistemology, but the point I want to make here is that it's a mistake to think that their role is to explain how knowledge is possible; on my account, transcendental arguments aren't necessary if the object of the exercise is to answer an epistemological how-possible question, and they aren't sufficient either.

Transcendental arguments set out to uncover necessary conditions for experience. The necessary conditions they set out to uncover are non-empirical or a priori conditions rather than causally necessary conditions. Since Kant is the patron saint of transcendental arguments and of how-possible questions in epistemology, it's tempting to think that the two must be connected in some way. In fact, it's far from obvious what transcendental arguments have to do with how-possible questions. Assuming that experience is one thing and knowledge another, how does the identification of what is necessary for *experience* help us to understand how *knowledge* is possible? Even if we focus on conditions of experience that can or must be established by some form of non-empirical philosophical reflection, it's still not clear how the identification of such conditions can be the key to explaining how knowledge, or knowledge of some specific kind, is possible.

This argument assumes that experience is not itself a form of knowledge. That is why the focus on conditions of experience in connection with questions about the possibility of knowledge seems misplaced. In contrast, Kant clearly thinks that experience *is* a form of knowledge.² To have what Kant calls 'experience' is to have perceptual knowledge of objects, so necessary conditions for experience are necessary conditions for perceptual knowledge of objects. But this still doesn't show that we can answer a how-possible question like (HP_{pk}) by arguing transcendently. An analogy might help: if someone asks how it is possible to travel from London to Paris in less than three hours, it would be perverse to think that what this question calls for is a specification of the *necessary* conditions for travelling from London to Paris in this time. It's possible to get from London to Paris in less than three hours by catching the Eurostar but this is a means rather than a necessary condition; it's also possible to get there in less than three hours by plane. In contrast, starting in London is a necessary condition for travelling from London to Paris but hardly a means of making this journey. A good answer to the question, 'How is it possible to get from London to Paris in less than

² As Kant puts it 'experience is an empirical knowledge, that is, a knowledge that determines an object through perceptions' (B218).

three hours?’ would be in terms of means (‘by catching the Eurostar’) rather than in terms of necessary conditions (‘first you’ve got to be in London’). Why, then, is it any more plausible to suppose that the way to explain how perceptual knowledge is possible is to identify its necessary conditions, a priori or empirical? Again, what we need are means rather than necessary conditions, so transcendental arguments are still beside the point.

I believe that these criticisms of the suggestion that transcendental arguments have a bearing on epistemological how-possible questions are sound, but they need spelling out. That is what I’m going to be doing in this chapter. It will be helpful to focus on a specific version of the proposal which I want to criticize, so let’s begin by examining the following representative passage from Arthur Collins:

Taking ordinary knowledge at face value, Kant asks how it is possible for us to have this knowledge. This is the order of Kantian transcendental arguments. It is the stance from which Kant formulates the famous ‘how-possible’ questions in the introduction to the first *Critique* and elsewhere... ‘How is such and such knowledge, knowledge that we do possess, possible?’ That is Kant’s question... The knowledge is given. The philosophical problem is to account for the possibility of this knowledge. (1999: 91–3)

One thing that is striking about Collins’s discussion is the suggestion that Kant’s how-possible questions are concerned with the possibility of ‘ordinary knowledge’ as well as the possibility of synthetic a priori knowledge. Presumably, ‘ordinary knowledge’ includes ordinary *perceptual* knowledge, that is, knowledge by means of perception of the existence and properties of objects. This implies that Kant’s how-possible questions include (HP_{pk}) as well as (HP_{sap}). By relating these questions to ‘the order of Kantian transcendental arguments’ Collins is suggesting that Kant saw his transcendental arguments as the key to answering both varieties of how-possible question.

Collins doesn’t explain how transcendental arguments help with how-possible questions. Let’s assume, however, that the proposal is that questions like (HP_{pk}) and (HP_{sap}) *can’t* be satisfactorily answered without using transcendental arguments (transcendental arguments are necessary) and that questions like (HP_{pk}) and (HP_{sap}) *can* be satisfactorily answered just by using transcendental arguments (transcendental arguments are sufficient). If transcendental arguments are necessary, then the multi-levels response must be no good unless it is itself a kind of transcendental argument. And if transcendental arguments are sufficient, then there is no need to look elsewhere for an answer to

(HP_{pk}) and (HP_{sap}); the multi-levels response would be superfluous even if it is independently viable.

Let's begin by examining the idea that transcendental arguments are sufficient in relation to (HP_{pk}). I have already suggested that Kant regards the 'experience' which is the focus of transcendental arguments as a cognitive achievement, a form of knowledge. Knowledge of what? The problem with stipulating that experience is, or involves, perceptual knowledge of external objects is that it would undermine what many regard as the anti-sceptical role of transcendental arguments.³ Suppose that we represent such arguments as having the following form: there is experience, necessarily if there is experience then p, therefore p. On an anti-sceptical reading, p is a proposition which is the target of sceptical attack, and the argument proceeds by showing that the truth of p is a necessary condition for something which the sceptic does not and cannot doubt. Thus, for anti-sceptical purposes, it must be, as Stern puts it, an 'indisputable fact about us and our mental life' (2000: 6) that we have experience but this won't be indisputable in the relevant sense if experience is defined as perceptual knowledge of external objects. The sceptic's question is whether there are any such objects or whether we can know anything about them, so in this context it won't do to assume at the outset that we have perceptual knowledge of external objects. If p is the proposition that such objects exist, a good transcendental argument for the truth of p must start with a 'thinner' notion of experience, one that doesn't beg any questions against scepticism about the external world.

In his *Refutation of Idealism*, Kant deals with this difficulty by introducing the notion of 'inner experience' and contrasting it with 'outer experience'.⁴ Inner experience is a form of self-knowledge; it is knowledge of the temporal order of one's experiences. Outer experience is perceptual knowledge of the existence of objects in space. Kant's claim is that outer experience is a necessary condition for inner experience. So if the sceptic grants that he has inner experience, then he must also grant that he has perceptual knowledge of external objects. This argument won't have any force against a sceptic who is prepared to question the existence of inner experience, so one issue is whether it

³ Stroud 2000c emphasizes the anti-sceptical role of transcendental arguments. See Stern 2000 for further discussion.

⁴ See B274–9.

is an indisputable fact about us and our mental lives that we have knowledge of the temporal order of our experiences.⁵ If not, then the notion of inner experience might need to be weakened even further to make the claim that we have such experience sceptic-proof. Even if we ignore this difficulty, however, there is still a problem with the suggestion that Kant's anti-sceptical transcendental argument in the Refutation of Idealism provides us with an answer (HP_{pk}). If this argument is successful, what it shows is that perceptual knowledge is *necessary* for inner experience but showing that perceptual knowledge is necessary for inner experience is not the same thing as explaining how perceptual knowledge is itself *possible*; we are none the wiser as to the best way of overcoming or dissipating apparent obstacles to its existence.

These considerations cast doubt on the suggestion that transcendental arguments are sufficient for the purposes of answering (HP_{pk}); we can't explain how perceptual knowledge is possible just by spelling out necessary conditions for experience if experience is understood in the way that it needs to be understood for anti-sceptical purposes. Might it nevertheless be the case that transcendental arguments are necessary in relation to (HP_{pk})? It seems not. We have already seen that (HP_{pk}) can be dealt with by a multi-levels response but this response doesn't proceed by identifying necessary conditions for experience. Although the multi-levels response talks about necessary conditions at Level 3, the level of enabling conditions, these are background necessary conditions for knowing about the external world by some specific means rather than necessary conditions for experience in general or for inner experience. This implies that a multi-levels response to (HP_{pk}) is not the same as, and does not incorporate, a transcendental response. So if, as I argued in Chapter 1, (HP_{pk}) can be satisfactorily answered by means of a multi-levels response, then it is false that transcendental arguments are necessary in relation to this how-possible question, or that the multi-levels response is no good unless it amounts to a transcendental argument.

If it's hard to see what anti-sceptical transcendental arguments have to do with (HP_{pk}), it's no easier to see what they have to do with (HP_{sap}). As we saw in the last chapter, an explanation of the possibility of synthetic a priori knowledge will need to identify means by which it is possible for us to acquire such knowledge, given the cognitive

⁵ Allison raises this question in his account of the Refutation. See Allison 1983: 304–9.

resources that are available to us. Kant's proposal in the case of geometrical knowledge is that the mathematician makes his way 'by means of intuitions' (A717/B745), but the idea that construction in pure intuition is a means of acquiring synthetic a priori geometrical knowledge explains how this kind of knowledge is possible without making any general claims about what is necessary for experience. Whereas in transcendental knowledge 'our guide is the possibility of experience', in mathematics 'all our conclusions are drawn immediately from pure intuition' (A782-3/B810-11). So it's not just that we *can* explain the possibility of synthetic a priori mathematical knowledge without employing transcendental arguments; the implication is that we *can't* explain how this kind of knowledge is possible by means of transcendental arguments.

What about other kinds of synthetic a priori knowledge? One thought is that if the outer experience is necessary for inner experience, then it is possible to know a priori that we have outer experience even though the proposition that we have outer experience is synthetic. This suggests that the conclusions of anti-sceptical transcendental arguments are synthetic a priori. Maybe this is the sense in which such arguments explain how synthetic a priori knowledge is possible. However, the problem with this suggestion is that it doesn't follow from the alleged fact that transcendental arguments make synthetic a priori knowledge available to those who grasp them that they explain how this kind of knowledge is possible; making something available is not equivalent to explaining its possibility. In any case, it's debatable whether the best thing to think about the conclusions of transcendental arguments is that they are synthetic a priori. This certainly won't be the best thing to think if one is already sceptical about the possibility of synthetic a priori knowledge; from this standpoint, the point to press is that transcendental arguments can't even provide us with synthetic a priori knowledge, let alone explain how they provide us with synthetic a priori knowledge and thereby explain how this kind of knowledge is possible.

Many of these problems for the proposal that transcendental arguments are a good way of answering how-possible questions stem from the assumption that transcendental arguments are anti-sceptical. Perhaps, in that case, we should consider the possibility that this is not the best or, at any rate, the only way of conceiving of transcendental arguments. On a different interpretation, such arguments spell out necessary conditions for experience but they define experience as perceptual or empirical knowledge of external objects. This will make transcendental

arguments ineffective against scepticism about the external world but, as Ameriks remarks, ‘not every interesting argument has to be a refutation of extreme skepticism’ (2003: 61). Furthermore, since Kant’s official definition of experience *simpliciter* identifies it with perceptual knowledge of objects, this implies that his transcendental arguments were not, in general, anti-sceptical. On this interpretation, the argument of the Refutation of Idealism is a special case.

What is the alternative to the anti-sceptical reading of Kant’s arguments? The obvious alternative is to read them as what Ameriks calls ‘regressive’ arguments. A regressive transcendental argument moves ‘from the assumption that there is empirical knowledge to a proof of the preconditions of that knowledge’ (Ameriks 2003: 51). Such arguments can still properly be described as spelling out necessary conditions for ‘experience’, on the assumption that ‘preconditions’ are necessary conditions and that experience is the same thing as empirical knowledge. What is the relationship between empirical knowledge and perceptual knowledge? Although Kant often treats them as equivalent this can’t be quite right. While perceptual knowledge is a form of empirical knowledge not all empirical knowledge is perceptual; empirical knowledge that has its source in the word of others isn’t perceptual even though the capacity to perceive is an enabling condition for the acquisition of knowledge by testimony. Nevertheless, as long as empirical knowledge is understood to include perceptual knowledge of external objects, the assumption that there is empirical knowledge is one which any self-respecting sceptic would want to question. That is why regressive transcendental arguments won’t be effective as anti-sceptical arguments.

What, then, is the point of a regressive transcendental argument if it takes empirical knowledge as ‘a premise to be regressively explained rather than as a conclusion to be established’ (Ameriks 2003: 55)? In what sense do regressive transcendental arguments promise to ‘explain’ our empirical knowledge? Let’s consider the suggestion that the point of such arguments is to provide an answer to this how-possible question:

(HP_{ek}) How is empirical knowledge possible?

If empirical and perceptual knowledge are the same thing, then an answer to (HP_{ek}) will also be an answer to (HP_{pk}). Even if the notion of empirical knowledge is broader than that of perceptual knowledge, we might still expect a good answer to (HP_{ek}) to provide the basis of a good answer to (HP_{pk}). So the proposal I’m now considering is that regressive transcendental arguments, understood in the way that

Ameriks understands them, are in the first instance trying to answer (HP_{ek}) and that they answer this question by identifying necessary conditions for empirical knowledge.

There isn't much to be said for this proposal as a reading of Kant's transcendental arguments. The problem is not that Kant doesn't use regressive transcendental arguments but that he doesn't use them to answer (HP_{ek}). This question doesn't appear in Kant's list of how-possible questions in the introduction to the first *Critique*, and this is a reflection of the fact that he doesn't regard the possibility of empirical knowledge as requiring any special explanation.⁶ In his view, there is no obstacle to the existence of empirical knowledge as such, so (HP_{ek}) is not a question which even arises. And if Kant thinks that (HP_{ek}) is not a question that arises, then the point of his regressive transcendental arguments can't be to provide an answer to this question. They must serve some other purpose in his system.

Is Kant right to be dismissive of (HP_{ek})? If we ignore the possibility of empirical knowledge that isn't perceptual and read (HP_{ek}) as equivalent to (HP_{pk}) then it might seem that (HP_{ek}) is a question which we should be taking seriously. After all, many philosophers have been persuaded that there are at least apparent obstacles that stand in the way of our possession or acquisition of perceptual knowledge of external objects, and this gives (HP_{ek}) an obvious point. But this vindication of (HP_{ek}) doesn't vindicate the proposal that regressive transcendental arguments provide a way of answering this question. I have already suggested that to explain how something is possible is to identify the means by which it is possible rather than necessary conditions for its possibility. This suggestion is enough to cast doubt on the idea that the possibility of empirical knowledge can be adequately explained just by identifying its necessary conditions. For example, there couldn't be any empirical knowledge without at least one knower. The existence of a knower is therefore necessary for there to be any empirical knowledge but this observation doesn't tell us how empirical knowledge is possible, any more than identifying the necessary conditions for travelling from London to Paris tells us how it is possible to travel from London to Paris. In both cases, the identification of necessary conditions is insufficient; what is missing is any reference to means.

One might think that this argument only seems compelling because it focuses on the wrong kind of necessary condition. While there might be

⁶ This is clear from Kant 1977: 275.

some necessary conditions for empirical knowledge the identification of which would not suffice to explain the possibility of empirical knowledge couldn't there be others the identification of which really would suffice to explain this possibility? As long as regressive transcendental arguments focus on this special sub-class of necessary conditions, perhaps we can continue to maintain that (HP_{ek}) can be answered by means of such arguments. The challenge for this line of thinking is, however, to explain what is distinctive of this sub-class of necessary conditions. They won't be causally necessary conditions, given that transcendental arguments in general are concerned with a priori conditions and that causally necessary conditions can't be known a priori. Could it be, then, that we can explain how empirical knowledge is possible by identifying its a priori necessary conditions?⁷ This doesn't seem right either; the existence of a knower is not just a necessary condition but an a priori necessary condition for empirical knowledge, yet more is needed to explain how empirical knowledge is possible than to point this out.

The remaining possibility is that we can answer (HP_{ek}) by identifying a special sub-class of the a priori conditions of empirical knowledge, and that these are the conditions that are the focus of regressive transcendental arguments. For example, Kant argues in the Transcendental Deduction that categorial thinking is an a priori necessary condition for empirical knowledge of objects. Categorial thinking is thinking by means of categorial concepts like *substance*, *causality*, and *unity*. Kant's claim in the Deduction is that we couldn't think about objects without using categorial concepts, and that we couldn't have empirical knowledge of objects without thinking about them. Furthermore, these are all things that we can know to be true by a priori philosophical reflection. That is why categorial thinking is not just a necessary condition but an a priori condition for empirical knowledge of objects.

Do we now have an answer to (HP_{ek})? Not if Kant is wrong to claim that we couldn't think about objects without using the categories. When I think that the cup I am holding is chipped, I am thinking about an 'object' (the cup) but I don't appear to have thought about it by means of a category; *cup* and *chipped* aren't categories. So it's certainly open to question whether objective thinking, thinking about objects,

⁷ A priori conditions can either be understood as ones which *can* be established without any empirical investigation or, additionally, as ones which *can't* be established empirically.

must be a form of categorial thinking.⁸ For the moment, however, let's just ignore this worry. The problem is that even if the categories really are necessary for empirical knowledge this doesn't suffice to explain how empirical knowledge is possible. Categorial thinking isn't sufficient for empirical knowledge, and the observation that there couldn't be empirical knowledge without the categories no more explains how empirical knowledge is possible than the observation that there couldn't be empirical knowledge without a knower. What is missing in both cases is any attempt to specify means of acquiring empirical knowledge or to address the apparent obstacles which led us to take (HP_{ek}) seriously in the first place. For example, suppose that one is persuaded by Stroud that what threatens to make empirical or perceptual knowledge problematic is the epistemic priority of sensory experiences over independently existing objects. In that case, a good answer to (HP_{ek}) will need to show one how to overcome or dissipate this apparent obstacle. It's hard to see how establishing the indispensability of the categories comes close to doing that.

It seems unlikely that Kant would have thought that the identification of a priori necessary conditions for empirical knowledge is *sufficient* to explain how empirical knowledge is possible, whether or not we are identifying empirical knowledge as perceptual. It's much more likely that he would have regarded the identification of necessary conditions as *necessary* for explaining the possibility of empirical or perceptual knowledge. If this is right then it might seem that there is still hope for the proposal that regressive transcendental arguments are necessary to answer questions like (HP_{ek}) even if they aren't sufficient. As we have seen, regressive transcendental arguments aim to establish necessary conditions for empirical knowledge. And if regressive transcendental arguments establish necessary conditions, *and* we can't explain how empirical knowledge is possible without identifying its necessary conditions, doesn't it follow that we can't explain how empirical knowledge is possible without relying on regressive transcendental arguments?

⁸ I don't take the claim that objective thinking is 'thinking about objects' to be equivalent to the claim that it is, or involves, thinking about objects *as* objects. To think about an object is simply to think about—to make a judgement about—what is in fact an object. So, for example, if I think that the cup in my hand is chipped then that is a piece of objective thinking; the cup is an object and I am thinking about it. Am I thinking about it 'as' an object? That obviously depends on what it would be to think about the cup in this way. In general, the more demanding one's conception of what it would be to think of objects as objects the less plausible it becomes to insist that this kind of objective thinking is necessary for knowledge of objects.

The first thing to notice is that this doesn't follow unless it is taken for granted that *only* regressive transcendental arguments can establish necessary conditions for empirical knowledge. What is true is that regressive transcendental arguments make claims about what is necessary for empirical knowledge but this doesn't mean that they themselves establish these claims or, even less plausibly, that they are the only way of establishing them. In any case, it's open to question whether explaining how empirical knowledge is possible requires the identification of its necessary conditions. Consider the claim that the cup in my hand is chipped. To know that the cup is chipped is to know something about the external world. To know that the cup is chipped by seeing that it is chipped is to be in possession of a piece of empirical knowledge. So we have explained how empirical knowledge is possible by explaining how it is possible to know by empirical means such things as that the cup is chipped, and we have explained how it is possible to know such things by pointing out that it is possible to know that the cup is chipped by seeing that it is. Yet seeing that the cup is chipped isn't a necessary condition for knowing that the cup is chipped. It isn't even a necessary condition for knowing by empirical means that the cup is chipped; one can also know that the cup is chipped by feeling that it is chipped or hearing from someone else that it is chipped. These are all ways of acquiring empirical knowledge of objects, and if we can explain how empirical knowledge is possible by reference to these ways of knowing there is no need to answer (HP_{ek}) by reference to necessary conditions. And if we don't need to answer (HP_{ek}) by identifying what is necessary for empirical knowledge, then we presumably don't need regressive transcendental arguments to answer (HP_{ek}).

To sum up, my claim is that even if we give up on the idea that transcendental arguments are anti-sceptical and read them as regressive instead, it's still not plausible that such arguments are either necessary or sufficient for the purposes of answering questions like (HP_{ek}) and (HP_{pk}). They aren't sufficient because they don't identify means of knowing or overcome obstacles to knowing by those means. They aren't necessary because we can explain how knowledge is possible by identifying ways or means of knowing that aren't necessary conditions. This is just another way of saying that once we have seen the possibility of a multi-levels response to (HP_{ek}) and (HP_{pk}), with its emphasis on means rather than on necessary conditions, we no longer need transcendental arguments. So the position is not that transcendental arguments make the multi-levels response superfluous but that the

multi-levels response makes transcendental arguments superfluous *if* the object of the exercise is to explain how knowledge, or knowledge of some specific kind, is possible.

Doesn't the multi-levels response talk about necessary conditions at Level 3, and doesn't this leave an opening for the idea that necessary conditions aren't superfluous from a how-possible perspective? I have already given a brief indication of what is wrong with this suggestion: in essence, the problem is that Level 3 conditions in the multi-levels response are necessary conditions for knowing about the external world by some specific means rather than necessary conditions for empirical or perceptual knowledge as such. Yet it is necessary conditions in the latter sense that are the focus of regressive transcendental arguments. From the standpoint of the multiple levels response, we should be sceptical about the idea of something as general as necessary conditions for empirical knowledge as such, as distinct from necessary conditions for knowing, or coming to know, by some specific means such as seeing or hearing or feeling. The suggestion, in other words, is that transcendental arguments are excessively general in their orientation, and that this is another reason for being sceptical about their use in connection with how-possible questions. I think that this point is important enough to merit a section to itself, so the next section is about this issue.

2.2 THE PROBLEM OF GENERALITY

Let's go back to the suggestion that it's possible to know that the cup is chipped by seeing that it is chipped. If I know that the cup is chipped by seeing that it is chipped then my knowledge that the cup is chipped is a piece of perceptual knowledge and therefore a piece of empirical knowledge. Specifically, it is empirical knowledge of the truth of a proposition about the external world. So we now have at least a partial explanation of the possibility of empirical knowledge of the truth of such propositions; we can know that they are true by seeing that they are true. This is a Level 1 explanation of the possibility of empirical knowledge, an explanation in terms of what Dretske calls 'epistemic seeing'.⁹ However, proponents of the multi-levels response insist that this is only a partial explanation. First, there is the worry

⁹ See Dretske 1969: 78–139 and Chapter 1 above for an account of this notion.

that it isn't possible to see that the cup is chipped because there are epistemological requirements on epistemic seeing that can't be met. This is the worry that is dealt with at Level 2 of the multi-levels response, the obstacle-removing level. Second, even if the alleged epistemological requirements can be met or be shown to be bogus, there is a further question which needs to be addressed: what makes it possible to see that the cup is chipped and thereby to know that it is chipped? This is a Level 3 question. The multi-levels response answers it by identifying background necessary conditions for epistemic seeing or for coming to know that the cup is chipped by seeing that it is chipped.

In Chapter 1, I briefly discussed the proposal that the perception of space is a background necessary condition for basic primary epistemic seeing, for seeing that *b* is *P* by seeing *b* itself in the case in which *b* itself is a material object. In order to see *b* itself I must be able to differentiate *b* from other things in its environment. Kant's proposal is that differentiating visual perception must be a form of spatial perception. Suppose that this is all true. Can we now reasonably claim to have identified a necessary condition for empirical knowledge of the external world? Obviously not. For a start, not all empirical knowledge is perceptual, and we shouldn't assume without further argument that the perception of space is a necessary condition for the acquisition of non-perceptual empirical knowledge. Indeed, from the fact that the perception of space is a necessary condition for knowing that the cup is chipped by *seeing* that it is chipped we can't even infer that the perception of space is a necessary condition for the acquisition of perceptual knowledge. There are many propositions about the external world which we can know to be true by smell or hearing, and we will see in the next chapter that there is a much looser connection between the perception of space and epistemic hearing or epistemic smelling than there is between the perception of space and epistemic seeing.

This is not a problem for the multi-levels response, given that the background necessary conditions that figure at Level 3 are necessary conditions for knowing that something is the case by some specific means. It is not committed to thinking that the background necessary conditions for seeing that *A*, where *A* is a proposition about the external world, will also be background necessary conditions for smelling that *A* or hearing that *A* or knowing that *A* by some other means. It certainly doesn't exclude this possibility but it doesn't bank on it either. On the other hand, regressive transcendental arguments do appear to bank on this possibility. To talk about necessary conditions

for empirical knowledge, and to suggest that such conditions are the focus of transcendental arguments, is to assume that there are such conditions, that is, necessary conditions for *any* empirical knowledge or for empirical knowledge *in general*.¹⁰ The question is whether this is a reasonable assumption. If not, this would call into question the viability of regressive transcendental arguments without calling into question the viability of the multi-levels response.

This issue can't be settled without examining specific proposals. Perhaps it's not plausible that the perception of space is a necessary condition for empirical knowledge in general but what about the categories? Couldn't it be that the categories are not just enabling conditions for knowing that A by seeing that A but necessary conditions for knowing that A by any empirical means? This is a question for Chapter 4, but at least we are now in a position to articulate an apparent problem for the project of constructing regressive transcendental arguments. The problem is what I'm going to call the problem of generality. The point is that regressive transcendental arguments can't move from the assumption that there is empirical knowledge to a proof of the preconditions of that knowledge unless there are preconditions or necessary conditions of empirical knowledge *per se*, but it is at least open to question whether there are any such conditions. If there aren't, then it is a weakness of regressive transcendental arguments that they are *excessively* general. This is the problem of generality; strictly speaking, it is the problem of excessive generality. This isn't a problem for the multi-levels response because it doesn't pursue generality at Level 3.

In one sense, there is a simple response to the problem of generality. Suppose that we take empirical knowledge to be a form of propositional knowledge. In that case, logically necessary conditions for propositional knowledge in general will also be necessary conditions for empirical knowledge in general. Here one might think of conditions such as truth, belief, and justification. Yet, however general these conditions might be, they are not the conditions that are at issue in regressive transcendental arguments. So appealing to what is necessary for propositional knowledge is not the way to solve the problem of generality. The necessary conditions which figure in regressive transcendental arguments are somehow less general than logically necessary conditions for propositional knowledge yet general enough to qualify as 'necessary

¹⁰ This assumption is explicit in Strawson 1997c and 1997d. See, for example, 1997c: 240.

conditions for empirical knowledge' rather than necessary conditions for knowing about the external world by some specific empirical means. Another way of formulating the problem of generality would therefore be to formulate it as the problem of identifying conditions of empirical knowledge which display just the right level of generality.

We are now in a position to understand why it's not right to view the multi-levels response as somehow incorporating regressive transcendental arguments. This can't be right because, as I have been arguing, Level 3 conditions in the multi-levels response needn't be anything as general as necessary conditions of empirical knowledge; they can be *means-specific* or *modality-specific*. That is why the multi-levels approach to how-possible questions is a genuine alternative to the transcendental strategy and isn't just a version of this strategy. The only way of bringing the two approaches closer together would be to suppose that regressive arguments can dispense with their quest for generality and concentrate on the identification of necessary conditions for the possession or acquisition of specific types of empirical knowledge, such as visual knowledge, auditory knowledge, and so on. At this point, however, we no longer have anything recognizable as a transcendental argument. Transcendental arguments are *essentially* general in their orientation. Their generality is not something that can be given up without giving them up.

This is as much as I want to say about the proposal that the best way to answer an epistemological how-possible question is by arguing transcendentially. It should now be clear that far from being the best way of answering questions like (HP_{ek}) and (HP_{pk}) , transcendental arguments aren't even *a* way of answering such questions. So what useful purpose do transcendental arguments serve in epistemology? If we are satisfied that the multi-levels response is the way to go with questions like (HP_{ek}) and (HP_{pk}) , why should we bother with transcendental arguments at all? As far as anti-sceptical transcendental arguments are concerned there is a simple answer to this question. We should bother with them if we want an answer to scepticism about the external world and think that anti-sceptical transcendental arguments such as the Refutation of Idealism can provide us with the kind of answer we are looking for. But what about regressive transcendental arguments? What useful purpose do they serve if they don't explain how knowledge is possible and don't aim to refute scepticism about the external world?

We can bring out the force of this question by taking a closer look at the central argument of Kant's Transcendental Deduction. The

sense in which this argument is regressive is that it tries to show that the categories are necessary conditions for empirical knowledge. So it clearly fits Ameriks's description of regressive arguments as moving from the assumption that there is empirical knowledge to a proof of the preconditions of that knowledge. If, as Kant argues, we couldn't think about objects without using categorical concepts then the categories are among the preconditions of empirical knowledge. But this still doesn't tell us what the *point* of establishing the indispensability of the categories is supposed to be; it doesn't identify the further question which a proof of the indispensability of the categories for objective thinking, and therefore for empirical knowledge, would help us to answer.

One possibility is that the proof tells us something about the concepts which we actually use in our objective thinking. If we *must* use the categories in order to think about objects, then we *do* use the categories when thinking about objects. On this reading, regressive transcendental arguments are *revelatory*; they reveal something about the way in which we actually think about the objects of our knowledge. On a different reading, regressive arguments such as Kant's argument in the Deduction are not revelatory but *validatory*. Their aim is to validate or legitimize the concepts we use in our objective thinking and which would otherwise lack what Kant calls 'objective validity'. Finally, one might think that regressive arguments are primarily *explanatory*. Their aim is to explain what empirical knowledge is by uncovering its a priori necessary conditions, the conditions under which it is possible.

On the face of it, none of these proposals is especially promising. For example, if we think about objects by means of the categories, won't that be evident to us on reflection? Why do we need a transcendental argument to tell us what we actually do? It's equally unclear why the categories need to be legitimized or how establishing their indispensability serves to legitimize them. Given the obscurity of Kant's distinction between legitimate and illegitimate ('usurpatory') concepts, we shouldn't attach too much weight to the idea that regressive transcendental arguments are needed to establish the objective validity of the categories. As for the idea that regressive transcendental arguments are explanatory, it's not obvious that the best way to explain what something *is* is to identify its necessary conditions, let alone its a priori necessary conditions. So it seems that we are no closer to understanding what the point of regressive transcendental arguments is supposed to be. But let's not be hasty. Those who think that such arguments serve a useful purpose in epistemology deserve a better run for their money.

In the following section, I'm going to consider whether, when the details of the various proposals I have just been describing are filled in, they provide us with a coherent and compelling account of the goal of regressive transcendental arguments.

2.3 REVELATION, VALIDATION, EXPLANATION

When we think about whether transcendental arguments are revelatory, validatory, or explanatory the first thing that comes to mind is that these ways of thinking about transcendental arguments aren't mutually exclusive. It's worth pointing out, however, that the second of the three readings comes closest to what Kant himself seems to have in mind. He explicitly represents the regressive argument of the Transcendental Deduction as validatory, and it would be natural to think that this tells us something about his conception of the goal of regressive arguments in general.¹¹ In any case, I'm going to spend more time on the proposal that transcendental arguments are validatory than on the other two proposals. But I want to start by saying something about the revelatory reading because understanding its limitations will make it easier to see why one might be drawn to the idea that legitimization is the real underlying aim of transcendental arguments.

Gary Hatfield gives a good account of the alleged revelatory function of transcendental arguments in the following passage:

Transcendental argument starts from some given body of knowledge, or some given cognitive achievement, and asks how it is possible. If it can be shown that the cognitive achievement in question is possible in only one way, then, given that the achievement is actual, the only possible means for its possibility must be actual, too. Kant's candidates for the starting point of transcendental arguments included: Euclid's geometry, Newton's physics, and the fact that we have 'experience', where experience is regarded as objective and demanding of intersubjective agreement. His transcendental arguments concluded by positing his famous 'categories' and 'forms of intuition'. (1990: 79–80)

The forms of intuition referred to in this passage are space and time. To 'intuit' something in Kant's sense is, roughly speaking, to perceive it,

¹¹ A 'deduction' in Kant's sense establishes a right or legal claim. Usurpatory concepts, such as *fortune* and *fate*, are problematic because 'no clear legal title, sufficient to justify their employment' (A85/B117) is obtainable either from experience or from reason. A deduction of the categories will show that they aren't usurpatory and that their employment is therefore justified. There is a helpful account of Kant's notion of a deduction in Henrich 1989.

and the things we intuit are individual 'objects'. When Kant describes space and time as forms of intuition what he means is at least that objects are given or presented in perception as spatially and temporally ordered. By making it plausible that the spatio-temporal forms of intuition and the categories are necessary for what it takes to be some actual cognitive achievement of ours, a transcendental argument makes it plausible that we actually employ the categories in the thinking that that achievement involves and that space and time are actually the forms of our intuition. This is presumably the sense in which Kant's arguments conclude by 'positing' his categories and forms of intuition; they reveal categorial thinking and spatio-temporal intuition to be actual by identifying them as necessary for some knowledge that we actually possess.

I have already raised one question about this reading of transcendental arguments: if their job is primarily revelatory, doesn't this threaten to make them superfluous? If space and time are the forms of our intuition, one might think that this should be directly evident to us; just by being introspectively aware of how we perceive objects, we should be able to tell that we always perceive them as spatio-temporally ordered. In that case, we don't need transcendental arguments to tell us that we perceive objects as spatio-temporally ordered. The same goes for the categories. We can tell that the categories figure in our objective thinking because they are necessary for objective thinking or, in Hatfield's terminology, the only possible means for the possibility of this kind of thinking. But can't we also tell that we use the categories in our objective thinking just by thinking about it? If this is right, then revelatory transcendental arguments are a very roundabout way of establishing something that can be established much more directly and simply.

One response to this line of argument would be to insist that our cognitive faculties aren't as transparent as they would need to be for revelatory transcendental arguments to really be superfluous. For a start, it isn't evident that we think about objects by means of categorial concepts. The concepts by means of which we usually think about objects are sortal concepts like *cup* or characterizing concepts like *chipped*, so the role of the categories in our thinking can't easily be established by direct inspection. That is why showing that objective thinking—thought about objects—*must* be a form of categorial thinking is a much more effective way of showing that our objective thinking *is* a form of categorial thinking. This is where transcendental arguments come into the picture. By revealing that even sortal thinking or thinking by means of characterizing concepts must be implicitly categorial they can reveal

something about the structure of our thinking that isn't otherwise apparent.

It's arguable that the forms of our intuition are easier to discern than the structure of our thinking. Even if we agree that it isn't evident that we think about object by means of the categories, we might be more reluctant to accept that it isn't introspectively evident that we perceive objects as ordered in space and time and that space and time are therefore the forms of our intuition. This assumes, however, that all it takes for space and time to be the forms of our intuition is for it to be the case that objects *are* given to us in perception as spatially and temporally ordered. In fact, Kant's thesis about our forms of intuition is that objects *must* be given to us as spatially and temporally ordered. This makes it easier to see why introspection can't tell us that space and time are the forms of our intuition; perhaps it can reveal how we do perceive objects but it can't tell us how we must perceive objects. A much better bet would be to argue transcendently that our intuition must be spatio-temporal because spatio-temporal intuition is a necessary condition for the perception of objects.

If we argue in this way, then we will be under pressure to explain why the perception of objects requires spatio-temporal intuition. A realist might say that the objects that are at issue here are fundamentally spatio-temporal objects, and that that is why they must be perceived as spatio-temporally ordered if they are to be perceived at all.¹² In contrast, Kant wants to say that the 'constitution of our faculty of intuition' (Bxvii) accounts for the connection between object perception and spatio-temporal perception. But if the constitution of our faculty of intuition consists in the fact that space and time are its forms, then his claim is not that our intuition must be spatio-temporal because spatio-temporal intuition is a necessary condition for the perception of objects. Instead, he would be claiming that spatio-temporal intuition is necessary for the perception of objects because our intuition must be spatio-temporal. As for the fact that 'space and time are the only forms of our possible intuition' (B146), this is something which Kant represents as inexplicable.

¹² This is Strawson's view in his later writings. The following passage is representative: 'the very notion of the generality of a concept implies the possibility of numerically distinguishable individuals falling under one and the same concept; and, once granted that objects are themselves spatio-temporal, then space and time provide the uniquely necessary media for the realization of this possibility in sensible intuition of objects' (1997c: 239–40).

I will have more to say about all of this in the next chapter. For the moment, the position we have arrived at is this: unlike anti-sceptical transcendental arguments, revelatory transcendental arguments are *self-directed* rather than *world-directed*.¹³ They uncover the structure of our cognitive faculties rather than the existence or structure of mind-independent reality. While some of what they reveal about our cognitive faculties might be discoverable in other ways, we shouldn't assume that this will always be the case. In principle, there could be facts about the structure of objective thinking or the structure of perception that aren't introspectively evident so it's not inevitable that revelatory transcendental arguments will end up revealing what we already know or can know in other ways. There might therefore be a role for such arguments in epistemology or the philosophy of mind. The fact is, however, that Kant doesn't see his regressive transcendental arguments as revelatory. The question is why not.

To answer this question, let's go back to the regressive argument of the Transcendental Deduction. In the opening paragraph of the Deduction, Kant distinguishes sharply between questions of fact and questions of right. A 'deduction' in Kant's technical sense addresses the latter rather than the former; it seeks to prove our right or entitlement to something:

Jurists, when speaking of rights and claims, distinguish in a legal action the question of right (*quid juris*) from the question of fact (*quid facti*); and they demand that both be proved. Proof of the former, which has to state the right or the legal claim, they entitle *deduction*. (A84/B116)

But showing that we do use categorial concepts in our thinking is a very different thing from showing that we are entitled to use them; it is to answer a question of fact rather than a question of right. That is why we shouldn't read the argument of the Deduction as revelatory; to read in this way is to miss its main point. Instead, we should read the argument as *validatory* if we want to do justice to the notion of a 'deduction'. Presumably, Kant wouldn't be trying to show that we are entitled to use the categories unless he thought that we do actually use them, but it is a presupposition of his argument that we actually use categorial concepts in our objective thinking; *that* this is what we do is not what he is trying to prove.

¹³ There is more on the distinction between world-directed and self-directed transcendental arguments in Cassam 1999.

Of course, it might be that the Deduction is a special case and that a Transcendental Deduction isn't the same thing as a transcendental argument. But if we are interested in regressive rather than anti-sceptical transcendental arguments, it's difficult to conceive of a better example of a regressive argument than Kant's argument in the Deduction. To the extent that this argument is validatory rather than revelatory, that tells us something about the nature of regressive arguments more generally. In the Deduction, Kant is assuming that empirical knowledge requires objective thinking and arguing that we couldn't think about objects without the categories. And the point of arguing in this way is that if we can prove that the categories are necessary for objective thinking then, as Kant puts it, 'this will be a sufficient deduction of them, and will justify their objective validity' (A97).

The two questions we now have to consider are these: first, why do the categories need to be validated? Second, how would proving their indispensability for objective thinking, and therefore their indispensability for empirical knowledge, serve to validate them? The first of these questions is a question about the question of right. Once we understand what this question means and why it arises we should be in a better position to figure out why Kant tries to legitimize the categories in the way that he does. At this stage, I'm not concerned with whether the argument of the Deduction actually works, that is, with whether it actually shows that the categories are preconditions of empirical knowledge. For the moment, my only question is whether we understand what Kant is trying to achieve in the Deduction and how he is trying to achieve it.

We can see why Kant thinks that categorial concepts like *substance*, *unity*, and *causality* need to be validated by comparing them with empirical concepts like *cup* and *chipped*. He says that 'many empirical concepts are employed without question from anyone' since 'experience is always available for proof of their objective reality' (A84/B116). For a concept to have 'objective reality' is for it to have 'application to objects which can be given us in intuition' (B150–1). So for experience to prove that a concept like *cup* has objective reality is for experience to prove that this concept has application to objects of intuition or, equivalently, to objects of experience. How does experience do that? By presenting us with cups. If our experience presents us with cups, then it presents us with objects to which the concept *cup* can correctly be applied. By presenting us with objects to which the concept *cup* can correctly be applied experience assures us that this concept has application to objects which *are* given us in intuition, and therefore to objects which *can be*

given us in intuition. And the thing about empirical concepts is that their objective reality can always be proved in this way; the fact that we can perceive cups means that we don't worry about the legitimacy of the concept *cup*.

This is not to say that empirical concepts can't be given a 'deduction', but the deduction of an empirical concept would be what Kant calls an 'empirical deduction'. An empirical deduction 'shows the manner in which a concept is acquired through experience and through reflection upon experience' (A85/B117). It might seem that showing that the manner in which a concept is acquired can't be a way of validating the concept since this only concerns the concept's '*de facto* mode of origination' rather than its legitimacy. This makes 'empirical deduction' sound like an oxymoron but Daniel Warren suggests a way round this difficulty. To acquire a concept from experience in Kant's sense is to abstract it from experience. For example, one acquires a simple concept like *cup* by abstracting it from one's experiences of cups. But if a simple concept has been abstracted from experience then it follows that it is instantiated in experience, and therefore that it can correctly or legitimately be applied to objects of experience. So, as Warren puts it, 'although, *in general* settling the question about origin will not settle the question about justification, *sometimes* it will'. For 'if an empirical concept is formed by abstracting it from experience, then it can be shown to be legitimate by appealing to the very experiences from which it is abstracted' (1998: 215). The very experiences from which the concept is abstracted serve to legitimize it, given that it is sufficient for a concept to be legitimate that it is instantiated in experience and therefore correctly applicable to objects of experience.

Now compare the categories or 'pure concepts of understanding'. Kant sometimes suggests that experience isn't available for proof of their objective reality. For example, he writes in the Schematism that:

[P]ure concepts of understanding being quite heterogeneous from empirical intuitions, and indeed from all sensible intuitions, can never be met with in any intuition. For no one will say that a category, such as that of causality, can be intuited through sense and is itself contained in appearance. (A137–8/B176–7)

This suggests that experience isn't available to prove the objective reality of the categories because they aren't and can't be instantiated in experience. So their legitimacy is always going to be in question in a way that the legitimacy of empirical concepts isn't open to question. In the case of the categories, we will find ourselves wondering whether

they can correctly be applied to objects of experience because experience can't present us with instances of them. And we shouldn't respond to this worry by trying to give the categories an empirical deduction. If they can't be instantiated in experience, they can't be abstracted from experience, and this means that they can't be legitimized by showing the manner in which they are abstracted from experience.

Does this explain why the question of right arises? The problem is this: it only makes sense to try to validate the categories if their legitimacy is in question in the first place, and it only makes sense to think of their legitimacy as being in question if we understand what it would be for a concept to be legitimate. But when we think about the notion of conceptual legitimacy in more detail, it's not at all clear what it amounts to or how it relates to the idea that a concept is, or can be, instantiated in experience. For a start, it's obviously not a necessary condition for conceptual legitimacy that a concept *is* instantiated in experience. If all the cups in the world were destroyed in some freak accident, the concept *cup* would no longer be instantiated in experience but there wouldn't be anything wrong in these circumstances with the *concept* of a cup; one could still legitimately employ this concept in thinking about cups. On the other hand, pejorative concepts like *Boche* are defective or illegitimate even though there is a sense in which they are instantiated in experience.¹⁴ *Boche* has the same reference as *German* so the sense in which *Boche* is instantiated in experience is that the concept *German* is instantiated in experience. Yet *Boche* is still defective since it licenses unacceptable inferences.¹⁵ Actual instantiation in experience is therefore neither necessary nor sufficient for conceptual legitimacy or even for the appearance of conceptual legitimacy. There is such a thing as a defective concept—Kant's own examples are *fate* and *fortune*—but we are still looking for a good account of what conceptual defectiveness consists in.

Suppose, then, that we try a different approach. Since it is clearly not Kant's view that what matters for conceptual legitimacy is actual instantiation in experience, *cup* and *Boche* aren't a problem for him. He thinks that there is only a question about the legitimacy of concepts which can't be instantiated in experience, not about the legitimacy

¹⁴ *Boche* is Dummett's example of a pejorative concept. See Dummett 1973: 454. For further discussion of pejoratives and other defective concepts see Boghossian 2003 and Williamson 2003.

¹⁵ For example, it licenses the inference from 'x is German' to 'x is cruel'. See Williamson 2003: 257 for more on this.

of concepts which aren't instantiated in experience. But if, as I have argued, the fact that a concept is actually instantiated in experience doesn't guarantee its legitimacy, it's hard to see how the fact that a concept is *capable* of being instantiated in experience guarantees its legitimacy. Consider the concept of a witch. This is presumably an example of a defective concept. Yet it isn't as if *witch* is incapable of being instantiated in experience; if there were witches, they would be objects of experience. The problem with the concept *witch* is not that it can't be instantiated in experience but that it 'carves reality up in an inappropriate way' (Kitcher 1982: 226) or doesn't pull its weight in a good explanatory theory of the world. Perhaps, in that case, we should read Kant as claiming that the fact that a concept can be instantiated in experience is only a necessary rather than a sufficient condition for it to be non-defective. Yet this is too strong. The concept of an object or event that is unperceivable in principle might figure in a good explanatory theory but it's built in to the concept that it can't be instantiated in experience.

As a matter of fact, Kant would have to agree that there is a problem with the suggestion that for a concept to be legitimate it must be capable of being instantiated in experience. He describes concepts like *God* and *soul* as 'transcendental ideas' because they transcend the possibility of experience.¹⁶ Yet he thinks that they still have a kind of subjective validity, which is presumably a kind of legitimacy.¹⁷ In contrast, categories like causality have objective validity, which is an even stronger form of legitimacy. So if it's true that the categories and transcendental ideas are legitimate concepts that can't be instantiated in experience, then it's false that a concept is legitimate only if it is capable of being instantiated in experience. Kant must think that this is false because in the end he himself is committed to the view that there are subjectively or objectively valid concepts which, in the terminology of the Schematism, 'can never be met with in any intuition'.

Of course, it might still be the case that the fact that the categories transcend the possibility of experience gives them the *appearance* of defectiveness and that this is enough to motivate the question of right. But is it true that the categories can't be instantiated in experience?

¹⁶ See the passage starting at A321/B377 for Kant's account of the notion of a transcendental idea.

¹⁷ Specifically, Kant claims to have provided a 'subjective derivation' of the transcendental ideas 'from the nature of our reason' (A336/B393).

Although I have quoted a passage from the Schematism which implies that this is precisely what Kant thinks, it's clear on reflection that this can't be his considered view. For his explanation of the difference between the categories and the transcendental ideas is that the latter 'have, in fact, no relation to any object that could be given as coinciding with them' (A336/B393) whereas this is not true of the categories. In other words, it turns out that the categories can be instantiated in experience after all. In that case, why do the categories need a transcendental deduction? If they can be instantiated in experience, then experience is available for proof of their objective reality, and if experience is available for proof of their objective reality it's hard to see why their legitimacy should be in question. They are now in the same position as empirical concepts, and this means that it ought to be possible to give them an empirical deduction if the question of right arises. What we are still missing, however, is a plausible story about how this question arises.

It seems, then, that there is a tension in Kant's position. When he wants to motivate the question of right he implies that the categories need a transcendental deduction because instances of them can't be encountered in experience. Yet when he wants to explain how the categories can be objectively rather than merely subjectively valid he insists that they can be instantiated or 'given' in experience. The only way of making sense of this is to argue that there are different senses in which a concept can be instantiated in experience, and that what Kant is denying when he says that the categories can't be instantiated in experience must be different from what he is asserting when he says that the categories can be instantiated in experience. So we now have the proposal that one kind of instantiation in experience is at issue when Kant is explaining why the legitimacy of the categories is in question and that another kind of instantiation in experience is at issue when he is explaining why the categories are different from the transcendental ideas.

What are these different kinds of instantiation in experience? According to Michael Friedman, Kant's view is that 'substances, causes and so on, are indeed given in experience—in virtue precisely of the fact that material substances interacting in accordance with the laws of motion through the fundamental forces of attraction and repulsion are given in experience' (1991: 83). Another way of putting this would be to say that the categories are *indirectly* instantiated in experience, via the instantiation in experience of other concepts such as *matter* and *force*.

If the transcendental ideas aren't even indirectly instantiated in experience, then this would explain why the categories are different from the transcendental ideas. And if empirical concepts are directly instantiated in experience, then this would explain why the categories are different from empirical concepts; it would be the fact that the categories can't be directly instantiated in experience that raises a special question about their legitimacy.

It's doubtful whether this proposal works. If the notion of indirect instantiation makes any sense, one would have to agree that lots of empirical concepts are only indirectly instantiated in experience. Yet their legitimacy isn't in question just because of that. For example, there are cases in which what one might call 'more specific' empirical concepts like *jockey* are only instantiated in experience via the instantiation of 'less specific' empirical concepts like *human being*. *Human being* can in turn only be instantiated via the instantiation of more specific concepts like man or woman or child. One perceives instances of the concept *jockey* by perceiving instances of the concept *human being*, and one perceives instances of *human being* by perceiving instances of *man* or *woman* or *child*. But none of this suggests that concepts like *jockey* or *human being* need a transcendental deduction or that they can't be derived from experience. By the same token, the alleged fact that the categories are only indirectly instantiated in experience shouldn't lead one to think that they need a transcendental deduction or that they can't be derived from experience.

It's also worth pointing out that there are differences between the individual categories that aren't addressed by the indirect instantiation proposal. Kant is exercised by the category of causality because he thinks that it embeds the concept of necessity and that necessity can't be experienced. In fact, this makes it difficult to see how *causality* can be instantiated in experience either directly or indirectly. Kant claims that the objective reality of this concept is demonstrated by 'the intuition of the movement of a point in space' (B292) but what has such an intuition got to do with necessity? On the other hand, there are other categories which don't appear to embed the concept of necessity. Consider the category of plurality. When one sees a collection of cups on the kitchen table, there is just as good a case for saying that the concept of plurality is instantiated in one's experience as there is for saying that the concept *cup* is instantiated in one's experience. It's true that one sees a plurality in this case by seeing the cups but it's also true that one sees the cups by seeing their spatial and other sensible properties. If one doesn't say that

the *cup* is only indirectly instantiated because of this, it's not clear why one should want to say that *plurality* is only indirectly instantiated.

The net result of all of this is that we still don't have a good account of why the categories as a whole need to be validated. The basic problem with Kant's account is that it is trying to do too many different things at once. It is trying to explain why each and every one of the categories need to be validated, and to do so in a way that respects the distinction between the categories and the transcendental ideas, as well as the distinction between the categories and empirical concepts. It must turn out that the categories have application to objects which can be given us in intuition, but it mustn't turn out that experience is available to prove the objective reality of the categories in the way that it is available to prove the objective reality of empirical concepts. The challenge is to find an account of the question of right which somehow meets all these constraints, and my discussion doesn't encourage the hope that this challenge can be met.

Even if we ignore all of this and just accept that the categories need to be validated, there remains a question about the proposed method of validation. What needs to be explained is how it can be a 'sufficient deduction' (A97) of the categories to show that they are necessary for objective thinking. If we can't think about objects of experience without the categories, why isn't this just a fact about us and our limitations?¹⁸ In what way is a concept shown to be legitimate just by being shown to be indispensable? Part of Kant's answer to these questions is that the indispensability he is after isn't just psychological indispensability; even Hume could grant that the concept of cause is indispensable in this sense but this wouldn't be enough by Kant's lights. If the categories are to be legitimized, they need to be shown to be indispensable in a stronger sense than that.

What is this stronger sense? Although Kant doesn't want to say that the categories are only psychologically indispensable, he does regard their indispensability as in some sense grounded in the structure of

¹⁸ Hannah Ginsborg has suggested to me that this is exactly the question which Kant himself is raising when he asks 'how *subjective conditions of thought* can have *objective validity*, that is, furnish conditions of the possibility of all knowledge of objects' (A89–90/B122). It looks as though his answer to this question is that the categories are objectively valid because they are conditions under which objects can be perceived as well as conditions under which they can be thought about. If this is right, then it is not Kant's considered view that establishing the indispensability of the categories for objective thinking 'will be a sufficient deduction of them' (A97). There is much more about the role of the categories in perception in Chapter 4 below.

the mind. The 'mind' in this context doesn't mean the 'human mind' but the mind in some more generic sense.¹⁹ His idea is that any being whose awareness of objects is perceptual must employ the categories in thinking about the objects of its perceptual awareness. So instead of saying that the indispensability of the categories is grounded in the structure of the mind, we might also say that it is grounded in the idea or concept of what it would be to think about an object of perceptual awareness, that is, what it would be to think about an object that is in some sense 'given'. On this account, the categories are something like conceptually or analytically necessary conditions for thinking about objects of perception, and we might think of regressive transcendental arguments more generally as spelling out conceptually or analytically necessary conditions. At any rate, if Kant were to deny that his conditions are necessary in this sense then it's hard to see how he could avoid having to view the indispensability of the categories as psychological.

If it's true that the categories are analytically necessary for objective thinking then this would obviously be of some interest. The categories would then be quite special when compared with other concepts. Even if concepts like substance and plurality are either directly or instantiated in experience, they wouldn't be like non-categorial concepts; the fact that there is no such thing as objective thinking without the categories would give them a very distinctive status. But the question is not whether establishing the analytic necessity of the categories would show that they are special but whether establishing their analytic necessity would suffice to validate them. Without an independent conception of what it would be for a concept to be objectively valid this question is impossible to answer. Kant doesn't explain why proving the indispensability of the categories in his sense amounts to a proof of their objective validity any more than he explains why it wouldn't be a proof of their objective validity just to show that they are psychologically indispensable or that they are instantiated in experience. He takes it as obvious that it will be a sufficient deduction of the categories 'if we can prove that by their means alone an object can be thought' (A97) but this looks like a stipulation rather than an argument.

We have now seen that there are two fundamental problems with the proposal that the regressive argument of the Transcendental Deduction

¹⁹ See Cassam 1998 for further discussion of the distinction between conditions that are grounded in the structure of the human mind and ones that are grounded in the structure of 'the mind' in a more generic sense.

is primarily *validatory*. The first, which is a reflection of the obscurity of the notion of a legitimate concept, is that it fails to explain why the categories need to be validated. The second, which is really an inevitable consequence of the first, is that it fails to explain the connection between legitimacy and indispensability. But the Deduction isn't Kant's only regressive transcendental argument. In addition to arguing that the categories are preconditions of empirical knowledge, he also argues that the forms of intuition—space and time—are necessary conditions for the perception of objects and therefore necessary conditions for any perceptual knowledge of objects. So before completely giving up on the idea that regressive transcendental arguments are *validatory* perhaps it would be worth taking a moment to consider whether a successful argument along these lines would amount to a 'validation' of the forms of intuition.

What would it be for a form of intuition, as distinct from a concept, to be objectively valid? Kant claims that the objective validity of the forms of intuition consists in their being conditions 'which limit our intuition and which for us are universally valid' (A27/B43). But why does the fact that objects must be presented to us spatio-temporally ordered 'validate' our awareness of them as spatio-temporal? Intuitively, the answer to this question depends on how we explain the 'must'. One possibility is that the objects we perceive as spatio-temporal really are spatio-temporal, and that we must perceive them as spatio-temporally ordered in order to perceive them as they really are. In contrast, if the objects we perceive as coloured aren't really coloured, then it is false that we must perceive them as coloured in order to perceive them as they really are. Instead, we might regard the perception of colour as embodying a kind of error. And one way of expressing the contrast between spatio-temporal perception and the perception of colour would be to say that the former has a kind of 'objective validity' that the latter lacks. On this account, the necessity of spatio-temporal perception is 'ontological' rather than psychological or conceptual. What this means is that it is ultimately grounded in the nature of the objects of perception rather than in human psychology or the concept of perception.

This somewhat 'realist' account of objective validity isn't available to Kant because his transcendental idealism commits him to denying that the objects which we perceive as spatio-temporal really are spatio-temporal. So he can't say that the sense in which we must perceive objects as spatio-temporal is that this is how we must perceive them in order to perceive them as they really are. But if we aren't transcendental

idealists then this is something which we might still want to say. What this shows is that the proposal that regressive transcendental arguments can be validatory isn't completely hopeless. When it comes to the forms of intuition, we can make some limited sense of the project of validating them, even if we are sceptical about the idea that the forms of intuition give rise to anything that is recognizable as a question of right. What isn't plausible, however, is that all regressive transcendental arguments can plausibly be thought of as validatory. Given the limitations of a validatory reading of the Transcendental Deduction, we are still missing a wholly convincing account of the point of regressive transcendental arguments in general. They might be revelatory and they might be validatory, but we are still looking for an account of the point of regressive arguments that are neither.

The remaining proposal is that regressive transcendental arguments are primarily explanatory, that is, that they aim to explain what empirical knowledge is by uncovering its a priori necessary conditions. When I first mentioned this proposal I said that it wasn't obvious that the best way to explain what something *is* is to identify its necessary conditions, let alone its a priori necessary conditions. We can bring out the force of this observation by means of an analogy. Suppose that we are trying to explain what cricket is to someone who doesn't know anything about it. A good start would be to say that cricket is a game played with bats and balls between two teams of eleven in which, to simplify a bit, the aim is to score more runs than the opposing team. One would then have to explain how runs are scored, the distinction between batting and bowling, and the various different ways in which wickets can be taken. In addition, one might point out that cricket is played on a pitch, and that the outcome of a cricket game often depends on the kind of pitch on which it is played. Finally, one might go on to say something about the history of cricket, its social and cultural significance in the countries in which it is popular, and the differences between it and other similar games such as baseball.

It looks as though we have now explained what cricket is. Clearly, this explanation won't help someone who doesn't know what a game is, but let's assume that we are speaking to someone who isn't in this position. In that case, the proposed explanation seems a perfectly good one, even though it isn't an explanation in terms of necessary conditions. To know what a game of cricket is, one needs to know something about how it is played and about what winning or losing at cricket consists in but to know these things is not, in any obvious sense, to know the

necessary conditions for cricket. Even if we can make sense of the idea that cricket has necessary conditions we can still explain what cricket is without talking in terms of necessary conditions. The stipulation that the necessary conditions must be a priori rather than ‘merely’ causally necessary conditions doesn’t help since, apart from anything else, it’s quite obscure how this distinction applies in the present case. For example, the existence of cricket bats, or of objects that can be used as bats, might be described as a necessary condition for the playing of cricket but one would be hard pushed to say whether this is a causally necessary or an a priori necessary condition.

The point of this analogy is, of course, to make it plausible that explaining what empirical knowledge is is no more a matter of identifying its necessary conditions than explaining what cricket is is a matter of identifying necessary conditions for the playing of cricket. The analogy suggests that we shouldn’t expect the distinction between a priori and causally necessary conditions of empirical knowledge to cut much ice, and that what is really problematic is the idea that the identification of necessary conditions is what really matters for explanatory purposes. So if the thought is that regressive transcendental arguments explain what empirical knowledge is by identifying its necessary conditions then there isn’t much to this thought. If we don’t need to know what is necessary for empirical knowledge in order to know what empirical knowledge is then we don’t need regressive transcendental arguments to explain what empirical knowledge is.

How convincing is this argument by analogy? So far in the book, I have been focusing on how-possible questions. What we now have is not a how-possible question but what might be called a *what-is-it* question. Specifically, we now have the question:

(W_{ek}) What is empirical knowledge?

The point of the analogy with cricket is to make it plausible that it is not true in general that a what-is-it question is to be answered by specifying necessary conditions for the existence or occurrence of whatever kind of thing the question is about. It might be objected, however, that the reason we don’t explain what cricket is by identifying its necessary conditions is that cricket is a game, and that we don’t usually explain games by reference to their necessary conditions. Instead, we explain them by outlining their rules and describing how they are played. But knowledge isn’t a game, and doesn’t have rules. Given this and other limitations of the analogy with cricket, it might appear that we are

none the wiser as to the best way of answering (W_{ek}). In particular, the fact that we are able to explain what cricket is without talking about its necessary conditions doesn't entitle us to conclude that (W_{ek}) can be answered without identifying necessary conditions for empirical knowledge.

How, then, should (W_{ek}) be tackled? We might begin by observing that empirical knowledge is knowledge that has its source in experience but this response to (W_{ek}) won't be of much use unless it is already clear what knowledge is and what it is for knowledge to have its source in experience. This implies that we can't answer (W_{ek}) without answering a much more general what-is-it question:

(W_k) What is knowledge?

It is in connection with this question that, for all its obvious limitations, the analogy with cricket begins to pay its way. The point is that when cricket is described as a game there is a further question that can always legitimately be asked. The further question, which can be asked about any game, is: how is it played? This is an example of a *how-question*. The relevance of this is that when someone is said to know that something is the case there is also a further how-question that can be asked, namely, *how* does he know?²⁰ My suggestion is that this is the key to tackling (W_k). To explain what knowledge is is to identify acceptable ways of answering such how-questions and to distinguish them from unacceptable ways of answering them.

A how-question is not the same as a how-possible question. To ask how cricket is played is not to ask how cricket is possible. One doesn't have to think that there is an obstacle that stands in the way of the playing of cricket in order to ask how cricket is played. The same goes for (W_k). One doesn't have to think that there is anything that stands in the way of S's knowing that A in order to ask how S knows that A. Although, as Austin points out, how-questions can be asked pointedly, with the implication that S doesn't really know what he is said to know, they can also be asked 'out of respectful curiosity, from a genuine desire to learn' (1979: 78). When they are asked in this spirit, a good answer will identify S's way of knowing, or means of coming to know what he knows, just as a good answer to the question 'how is cricket played?' will identify the methods or techniques of playing cricket. So, for example, saying that S can see that the cup is chipped is a perfectly good answer

²⁰ Cf. Austin 1979: 77.

to the question, 'how does S know that the cup is chipped?' And if this is how S knows that the cup is chipped, then S's knowledge will count as empirical. So we are now in a position to tackle (W_{ek}). To answer this question is to identify those particular ways or means of knowing which count as sources of empirical knowledge. The aim shouldn't be to come up with a complete list but to give examples of canonical sources of empirical knowledge with which more controversial putative sources can be compared. Thus, one might say that one's knowledge that A will count as empirical if one knows that A by seeing or hearing or smelling that A while leaving it open whether something like clairvoyance, if there were such a thing, would be a source of empirical knowledge.

Although (W_{ek}) is a different question from (HP_{ek}), 'how is empirical knowledge possible?', there is an obvious parallel between the response to (W_{ek}) which I am now recommending and the first level of a multiple levels response to (HP_{ek}). As I argued in the last chapter, the first thing that needs to be done in order to explain how empirical knowledge is possible is to identify means of acquiring this kind of knowledge. This Level 1 response to (HP_{ek}) is an example of what I called a Means Response to a how-possible question. What we now have is a parallel Means Response to (W_{ek}). This response trades on the idea that, as Williamson puts it, 'if one knows that A, then there is a specific way in which one knows' (2000: 34). In my terms, these specific ways of knowing that A are specific ways of coming to know that A. According to the Means Response to (W_{ek}), specifying appropriate means of coming to know is an appropriate means of explaining what empirical knowledge is.²¹

The Means Response to questions like (W_k) and (W_{ek}) is different from other popular responses. In particular, it is different from the analytic response to (W_k), according to which the way to explain what knowledge is is to analyse the concept of knowledge with a view to uncovering non-circular necessary and sufficient conditions for knowing. The Means Response doesn't imply that the concept is unanalysable in this sense, but it does suggest that analysing the concept of knowledge into more basic concepts is not the only or the best way of explaining

²¹ This fits in with Goldman's idea that knowledge is something that is attainable 'by a wide variety of sometimes independent and sometimes interconnected pathways' (2002: p. xi). If there are lots of different pathways to knowledge we can explain what knowledge is in part by identifying and ranking these pathways. In a fuller account mention would also need to be made of the various means by which knowledge can be retained and transmitted. Necessary conditions aren't to the point.

what knowledge is. The Means Response to (W_k) is also different from the idea that (W_k) is a question for empirical science. This approach is sometimes recommended on the basis that (W_k) is a question about 'knowledge itself, not our concept of knowledge' (Kornblith 2002: 1), and that knowledge itself is a natural kind. Apart from any scepticism one might feel about the idea that knowledge is a natural kind, there is no need to think of (W_k) as a question for science if it can be answered by an account of the different ways in which knowledge is acquired. It is philosophical reflection rather than empirical science which tells one that, say, seeing that A is a way of coming to know that A.

For present purposes, the most important difference is between the Means Response to (W_k) and (W_{ek}) and the transcendental response. Since means of knowing needn't be necessary conditions for knowing, the Means Response explains why we should be sceptical about the idea that regressive transcendental arguments are the best way of explaining what empirical knowledge is. If we read (W_{ek}) as calling for a specification of means rather than necessary conditions, then regressive transcendental arguments are no more to the point in relation to this question than they are in relation to (HP_{ek}) . Means of knowing have background necessary conditions, but I have already argued that these are unlikely to be the highly general conditions for empirical knowledge in general which figure in regressive transcendental arguments. If this is right, then the role of regressive transcendental arguments in epistemology can't be to explain what empirical knowledge is. They might have a very limited revelatory or validatory role but it's doubtful that they have much of an explanatory role.

3

Perceptual Knowledge (I): Space

3.1 WHAT MAKES IT POSSIBLE?

In the opening chapter of this book I gave an account of how-possible questions in epistemology. One such question is:

(HP_{ew}) How is knowledge of the external world possible?

I said that how-possible questions are obstacle-dependent questions; to ask how *x* is possible is to ask how something that looks impossible is nevertheless possible. But is there any reason to think that knowledge of the external world looks impossible? Yes, according to one line of thinking, because this kind of knowledge is possible only if perceptual knowledge is possible, yet there are apparently undeniable facts about sense perception that ‘make it difficult to understand how we could get any knowledge at all of the world around us on the basis of sense perception’ (Stroud 2000b: 5). So having started with (HP_{ew}) we now find ourselves facing another how-possible question:

(HP_{pk}) How is perceptual knowledge possible?

Perceptual knowledge is knowledge of the external world that has its source in perceptual experience. If perceptual knowledge is impossible it doesn’t immediately follow that knowledge of the external world is impossible. There are, after all, non-perceptual sources of knowledge of the world around us. In practice, however, perception is such a fundamental source of knowledge for us that it’s hard to see how one could reasonably think that perceptual knowledge is impossible without also calling into question the possibility of knowledge of the external world.

How, then, should we deal with (HP_{pk})? In Chapter 2 I argued that transcendental arguments aren’t a good way of tackling how-possible questions so let us stick with the multi-levels approach that I was recommending in Chapter 1. Multi-levels responses to how-possible

questions start by identifying means by which the allegedly problematic knowledge can be acquired. So, for example, one might think that seeing that the cup into which I'm pouring coffee is chipped is a way of coming to knowing that it is chipped, and that this kind of epistemic seeing is therefore one source of perceptual knowledge. This is an example of a Level 1 or Means Response to (HP_{pk}) but now we run into various considerations that have been thought to show that epistemic seeing is impossible. It might be argued, for example, that there are various conditions on epistemic seeing that we can't meet, and that that is why it isn't possible for us to know that the cup is chipped by seeing that it is chipped. Since these conditions are also conditions on knowing that the cup is chipped by feeling that it is chipped, or on knowing that any proposition about the external world is true by perceiving that it is true, we can now see why perceptual knowledge looks impossible.

One way of dealing with this worry would be to show that there are no genuine conditions on epistemic seeing or on epistemic perceiving that we are incapable of fulfilling. This is the Level 2 or obstacle-removing response to (HP_{pk}) that I was describing in Chapter 1. But even if we are satisfied that there is nothing that prevents us from seeing that the cup is chipped and thereby knowing that it is chipped there is a further question that we might ask. The further question is: what makes it possible to see that the cup is chipped? More generally, where A is a proposition about the external world, we might ask: what makes it possible to perceive that A is true? These *what-makes-it-possible* questions are what I have been calling Level 3 questions. They are questions about the enabling conditions, about the background necessary conditions, for the acquisition of perceptual knowledge by the various means identified at Level 1. Maybe we don't have to answer such questions in order to answer (HP_{pk}) but we might still wonder whether it is possible to say anything informative about what makes perceptual knowledge possible.

Here is one suggestion: it wouldn't be possible to see that the cup is chipped if one lacked eyes. So having eyes, or at least some functionally equivalent organ, is a background necessary condition for knowing that the cup is chipped by seeing that it is chipped.¹ The fact is that there are many physiological and biological enabling conditions for epistemic seeing or for epistemic perceiving, and this is not something that even those who are sympathetic to what I have called 'minimalism' in relation

¹ This is John Campbell's example. The argument of this paragraph owes quite a bit to his comments on an earlier version of this chapter.

to (HP_{pk}) should want to deny. What minimalists should deny is that there are substantive enabling conditions for the acquisition of perceptual knowledge that, unlike physiological and biological conditions, can be established without any empirical investigation. So-called ‘anti-minimalists’ think that there are such conditions. Enabling conditions that can be established without any empirical investigation are weakly a priori conditions.² It seems, therefore, that what is at issue between minimalism and anti-minimalism is whether there are any weakly a priori enabling conditions for knowing that the cup is chipped by seeing that it is chipped, or for knowing that a proposition about the external world is true by perceiving that it is true.

What would be examples of such conditions? It might seem that I have already given one: presumably, it doesn’t take any empirical investigation to work out that having eyes or some functionally equivalent organ is a background necessary condition for *seeing* that the cup is chipped. But this isn’t an especially interesting a priori enabling condition for seeing that the cup is chipped, and it isn’t an a priori condition for knowing that the cup is chipped by non-visual means. Ambitious anti-minimalists want more than this. They think that there are much more substantive a priori enabling conditions for epistemic seeing, and that some of these conditions are also background for acquiring knowledge of the external world by non-visual but nevertheless perceptual means. If they are right about this, then it looks as though armchair philosophy has something interesting to say about what makes perceptual knowledge possible.

Kant is the paradigm ambitious anti-minimalist. I have already mentioned his idea that the perception of space, the capacity to perceive space or spatial properties, is an a priori condition for basic primary epistemic seeing. In primary epistemic seeing, one sees that an object *b* is *P* by seeing *b* itself.³ In basic primary epistemic seeing, *b* itself is a material object. Seeing that the cup is chipped by seeing the cup itself is an example of this kind of epistemic seeing, and one reason for thinking that it must involve the perception of space is that it is the perception of space that enables one to differentiate the cup from its surroundings. If one couldn’t do that one wouldn’t be able to see the cup itself and therefore wouldn’t be able to see that the cup is chipped. According to Kant, however, the perception of space is not

² The contrast is with strongly a priori conditions, ones that *can’t* be established empirically.

³ The contrast is with secondary epistemic seeing. See Dretske 1969: 153–62 for more on this contrast.

just an enabling condition for basic primary epistemic seeing but also an enabling condition for epistemic perceiving generally and for the resulting perceptual knowledge. What he is committed to, in other words, is the following *Spatial Perception Requirement* (SPR): in order to perceive that something is the case and thereby to know that it is the case one must be capable of spatial perception.

A further consideration is this: it wouldn't be possible to see that the cup is chipped if one lacked the concepts *cup* and *chipped*. Epistemic seeing is, in this sense, conceptual. *Cup* and *chipped* are empirical concepts, and Kant argues that a capacity for categorial thinking, thinking by means of categorial concepts like *substance*, *unity*, *plurality*, and *causality*, is a background necessary condition for the possession and acquisition of empirical concepts. If he is right about this, and if all epistemic perceiving is conceptual in the way that epistemic seeing is conceptual, then he is committed to the following *Categorial Thinking Requirement* (CTR): in order to perceive that something is the case and thereby to know that it is the case one must be capable of categorial thinking.

Can SPR and CTR be established without any empirical investigation? If so, then this would represent a significant victory for anti-minimalism. Minimalists think that there are no background necessary conditions for knowing about the external world by perceptual means that are both substantive and can be established by armchair philosophy, yet there could scarcely be enabling conditions for epistemic perceiving that are more substantive than SPR and CTR. So the most important questions for proponents of SPR and CTR are whether they represent genuinely *necessary* conditions for epistemic perceiving and, if so, whether they are necessary conditions that can really be established non-empirically. The first of these questions might be called the *question of necessity*. The second might be called the *question of a priority*. Minimalism is committed to thinking that the answer to at least one of these questions is 'no'. To think that the answer to both questions is 'yes' is to be committed to a form of anti-minimalism.

Although SPR and CTR are especially interesting in the context of a debate between minimalism and anti-minimalism they are also interesting in their own right. Whether or not one cares about the prospects for anti-minimalism it's a good question whether spatial perception and categorial thinking really are a priori enabling conditions for knowing that propositions about the external world are true by perceiving that they are true. In relation to SPR, for example, one might be willing to grant that epistemic *seeing* is made possible by

the perception of space while being much more sceptical about the suggestion that the same is true of other forms of epistemic perceiving. As for CTR, the obvious question is whether it is remotely plausible that one needs to be able to think categorially in order to see that the cup is chipped. Maybe one needs some other concepts in order to have concepts like *cup* and *chipped* but why are categorial concepts necessary?

I will come back to CTR, and to questions about its plausibility, in the next chapter. In this chapter I'm going to focus on SPR. I'm going to be arguing for a suitably qualified version of SPR, that is, a version of SPR that acknowledges the differences between the role of spatial perception in epistemic seeing and its role in other forms of epistemic perceiving. Despite these differences, I believe that SPR is defensible. It is also plausible, I will argue, that SPR is at least weakly a priori. Not only can it be established non-empirically, it is doubtful that it can be established by empirical methods. This is not to say, however, that a philosophically satisfying answer to (HP_{pk}) *must* identify the a priori enabling conditions for the acquisition of perceptual knowledge. One can agree there *are* non-empirical enabling conditions for the acquisition of perceptual knowledge without also agreeing that in order to explain how perceptual knowledge is possible it is necessary to say what these conditions are. This means that, in the terminology of Chapter 1, the anti-minimalism for which I am going to be arguing in relation to (HP_{pk}) is 'moderate' rather than 'extreme'.

Before moving on, there is one more issue which I would like to introduce, even though it doesn't have much to do with the debate between minimalism and anti-minimalism. When something like spatial perception is described as a background necessary condition for perceiving that something is the case we can obviously ask *whether* this is true but we can also ask *why* it is true. In asking this question what we want to know is: what *makes it the case* (assuming that it is the case) that spatial perception is an enabling condition for knowing about the world around us by perceptual means? I'm going to call this the *question of foundations* since it is a question about the basis or foundation of requirements such as SPR and CTR. Idealism is the view that their sole basis is the structure of the human mind, and that they are therefore wholly subjective in origin.⁴ Realism is the view that they are not wholly

⁴ Allison would call SPR and CTR 'epistemic conditions', that is, necessary conditions for the 'representation of an object or an objective state of affairs' (1983: 10). An idealist in Allison's sense (and Kant's, if Allison reads Kant correctly) thinks that 'whatever is

subjective in origin. Realists think that what makes spatial perception an enabling condition for knowing things about the world around us by perceptual means is the nature of the world around us as well as our own cognitive constitution.⁵ Crudely, the realist's idea is that epistemic perception must be spatial because the world around us is a spatial world.

Clearly, minimalists who are sceptical about requirements like SPR and CTR won't want to get involved in an argument about their basis; if one thinks that SPR and CTR aren't correct, one doesn't have to worry about what makes them correct. On the other hand, the question of foundations should interest philosophers who endorse SPR and CTR; if there are spatial and categorial requirements on knowing about the world around us by means of epistemic perception, it would be natural to wonder whether this is so because of the way we are or because of the way the world is. Kant is an idealist anti-minimalist but it's not obvious that one has to be an idealist in order to be an anti-minimalist. My own sympathies are with realism, at least in relation to SPR, so one of the challenges I'm going to face is to show how the moderate version of anti-minimalism which I want to defend can be reconciled with realism. As we will see, meeting this challenge is much less straightforward than one might think.

3.2 SPATIAL PERCEPTION

Let's now consider the Spatial Perception Requirement in more detail. When I first introduced this requirement I said that spatial perception is the capacity to perceive space or spatial properties. It will be a little while before it is clear why SPR talks about the *capacity* to perceive space or spatial properties rather than the *actual* perception of space or spatial properties. In the meantime, perhaps it would be helpful to say

necessary for the representation or experience of something as an object ... must reflect the cognitive structure of the mind (its manner of representing) rather than the nature of the object as it is in itself' (1983: 27). Notice the slide from talk about what is necessary for the representation of objects to talk about what is necessary for the representation of something *as* an object. In general, it's a good idea not to assume that representing an object is the same as representing it as an object. An even better idea is to think of enabling conditions as necessary conditions for perceiving or thinking about objects rather than for 'representing' objects.

⁵ Strawson is a realist in this sense. See his 1997c. Cf. Guyer 1987.

something about the notion of a 'spatial property'. What I have in mind when I talk about spatial properties are specifically spatial properties such as extension, shape, or spatial location.⁶ In a derivative sense, colour and other properties might also be described as spatial. Maybe something can't be red unless it is extended, but this doesn't make redness a *specifically* spatial property, any more than the fact that one can't be an unmarried man without being extended makes the property of being a bachelor a specifically spatial property. Specifically spatial properties are non-derivatively spatial, and I'm going to take it that the capacity to perceive spatial properties is the capacity to perceive spatial properties in this sense.

When one perceives one or more of the spatial properties of an object, one might be said to be perceiving spatially. It's important to recognize, however, that spatial perception needn't be the perception of the spatial properties of an object. Suppose that one sees two parked cars and sees the space between them. Seeing the space between two cars is a form of spatial perception, but the gap between the cars is not a property of either car. Nevertheless, the perception of the gap can properly be described as the perception of space, or at least as the perception of *a* space—a parking space, as we say. That is why spatial perception can either be understood as the capacity to perceive spatial properties or as the capacity to perceive space. In much of what follows, the distinction between the capacity to perceive space or spatial properties and the actual perception of space or spatial properties won't matter very much, so I won't always be careful about respecting it.

Turning to the question of necessity, why should anybody think that SPR states a genuine requirement? I want to discuss two main arguments for SPR, a direct and an indirect argument. Both arguments start from the observation that in primary epistemic perceiving one perceives that an object *b* is *P* by perceiving *b* itself. So, for example, one sees that the cup is chipped by seeing the cup itself. If *b* itself is an object, then perceiving *b* itself is a form of object perception. Suppose that there are background necessary conditions for object perception. Assuming that one can neither perceive that *b* is *P* nor know on this basis that *b* is *P* without perceiving *b* itself, any background necessary conditions for object perception will come out as enabling conditions for the acquisition of perceptual knowledge by means of epistemic

⁶ In Warren's terminology, the role of SPR is to assign a special role in the perception of objects to the perception of 'specifically spatial features' (1998: 188).

perception. The direct and indirect arguments for SPR agree that object perception requires the perception of space, and this is the basis of their endorsement of SPR. The only difference between the two arguments is in how they attempt to establish a link between object perception and spatial perception.

Both arguments assume that 'objects' are material or spatial objects. The direct argument for SPR says that it's not possible to perceive such objects without perceiving some of their spatial properties or without perceiving space. But in order to perceive space or the spatial properties of an object one must obviously have the capacity to perceive space or spatial properties. Possession of this capacity can therefore be regarded as a background necessary condition for object perception. And if spatial perception is a background necessary condition for object perception, it's also a background necessary condition for epistemic perceiving and for the acquisition of the knowledge that epistemic perceiving makes available to one. So, for example, one couldn't see that the cup is chipped or thereby know that it is chipped unless one could at least see the cup, and it wouldn't be possible to see the cup in the absence of a capacity for spatial perception.

The indirect argument for SPR goes like this: it isn't possible to perceive an object without differentiating it perceptually from other objects in its environment. It's not possible to differentiate an object in this sense without perceiving space or spatial properties. The capacity to perceive space or spatial properties is therefore an enabling condition for object perception, and therefore also an enabling condition for epistemic perception. For example, suppose that the particular cup which I see to be chipped is one among a number of cups that I can also see. In that case, I couldn't see that a particular cup is chipped without differentiating it from the remaining cups.⁷ Doing this requires the perception of space or spatial properties, and that is why spatial perception is an enabling condition for seeing that the cup is chipped and thereby knowing that it is chipped.

Are these arguments for SPR any good? A natural reaction to the direct argument would be to accuse it of trivializing this requirement. Sure, if one is talking about what is required for the perception of *spatial* objects, it's not surprising that the necessary conditions include the perception

⁷ This is what Dretske is getting at when he claims that a subject S sees an object D if and only if D is 'visually differentiated from its immediate environment by S' (1969: 20). There is much more on Dretske below.

of *space* or *spatial* properties. Indeed, the connection between the two might seem so obvious that it's not clear why one should think that spatial perception is only a *background* necessary condition for the perception of spatial objects. Why not a necessary condition *simpliciter*? What work is being done by the insertion of the word 'background'? The interesting question, one might conclude, is whether SPR can be defended without stipulating that 'objects' are spatial objects. On this issue, the direct argument leaves us none the wiser.

I want to suggest that this reaction to the direct argument is mistaken. Far from being trivial, the problem with the proposal that the perception of spatial objects must be spatial is that it runs into some fairly obvious and potentially devastating counterexamples. One only has to think about the possibility of perceiving spatial objects by senses other than sight and touch to get a sense of how difficult it is to argue directly from the premiss that a particular object is spatial to the conclusion that one can't perceive it without perceiving one or more of its spatial properties. I will have much more to say about this issue when I look at the direct argument in more detail. First, I want to take a closer look at the indirect argument for SPR. The emphasis in this argument is on perceptual differentiation rather than on a direct link between the perception of spatial objects and the perception of space, so let's begin by examining its conception of what this form of differentiation involves.

3.3 PERCEPTUAL DIFFERENTIATION (I)

We rarely perceive objects without perceiving their surroundings. When I see the cup into which I'm pouring my morning coffee, I don't just see the cup. I also see my hand, the coffee, the coffee jug, the table that the cup is resting on, and so on. The first premiss of the indirect argument for SPR implies that I don't see the cup unless I visually differentiate it.⁸ To differentiate an object is to distinguish it from its surroundings, and the indirect argument says that object perception requires this form of differentiation.⁹ What is more, the differentiation must be perceptual.

⁸ I describe this as an *implication* of the indirect argument because this argument is explicitly concerned with what it takes to *perceive* an object rather than with what it takes to *see* one. The differentiation requirement applies to the visual case because seeing an object is a way of perceiving it.

⁹ The indirect argument's notion of 'differentiation' is similar to Wiggins's notion of 'singling out'. To single out an object x is 'to isolate x in experience; to determine or

In the present case, therefore, it isn't a matter of *judging* or *believing* that the cup is distinct from all the other things that I also see but of *seeing* its distinctness from these things. Unless the distinction between a particular object *b* and its surroundings registers with one perceptually, one simply isn't perceiving *b*.

We can bring out the force of this requirement by looking at an example of Dretske's.¹⁰ The thesis that Dretske wants to defend says that a subject *S* sees, in the relevant sense of 'sees', an object *D* if and only if *D* is 'visually differentiated from its immediate environment by *S*' (1969: 20).¹¹ Suppose, then, that you take nine cubes and arrange them in the form of a square. Next, you move far enough away so that the ensemble of blocks appears to you as a uniform mass, like a wall, without distinguishable parts. In these circumstances, can you be said to see one particular cube in the middle of the ensemble, say cube #5? Not if it isn't visually differentiated from the other cubes. As Dretske puts it, 'although cube #5 makes a positive contribution to the way the "square" looks, in the sense that without it the square might appear to have a hole in the center, and in the sense that the light from #5 is stimulating your visual receptors, I do not think we would go so far as to say that one could see cube #5' (1969: 23–4). Notice that what Dretske is saying here isn't just that you aren't aware of seeing cube #5. The claim is that you don't see cube #5, and that you don't see #5 because it isn't differentiated from the surrounding cubes.

How convincing is this example? Suppose that cube #5 is the central cube in the array and that you are looking directly at the centre of the square. Why couldn't one say that you do in fact see cube #5 without realizing it? Another of Dretske's examples shows why we shouldn't say this. Imagine an astronaut looking down from an orbiting satellite at a portion of the earth that looks uniformly green to him. Even if he knows that there is a hill beneath him somewhere, and is looking in the direction of the hill, it's just not plausible that he sees the hill. He doesn't see it because 'nothing marks it out as an isolable element in the

fix upon *x* in particular by drawing its spatio-temporal boundaries and distinguishing it in its environment from other things of like and unlike kinds' (Wiggins 2001: 6). One question about these formulations is whether differentiating or singling an object out is something that I literally *do*. If not, then the claim that I don't see the cup unless I visually differentiate it needs to be understood as the claim that I don't see the cup unless it is visually differentiated for me. I ignore this complication in what follows.

¹⁰ See Dretske 1969: 24.

¹¹ Dretske calls the kind of seeing that is at issue here 'non-epistemic'.

landscape' (1969: 26); he is too far away to differentiate it and therefore too far away to see it, even though there is actually a hill where he is looking. Similarly, even if you already happen to know that cube #5 is the central cube in the square, you still don't *see* this cube as long as it isn't a visually differentiated element in the square; you are too far away to differentiate it and therefore too far away to see it even though it is actually where you are looking.

It seems, then, that the perceptual differentiation requirement has quite a lot going for it as long as one concentrates on the visual case.¹² But what the first premiss of the indirect argument for SPR says is that it isn't possible to *perceive* an object without differentiating it from other objects in its immediate environment, and this is a far more ambitious claim than the one about seeing that Dretske wants to defend. The question we have to ask, therefore, is whether the perceptual differentiation requirement applies to the perception of objects by senses other than sight. To see what the problem is, imagine a blind person S standing right in front of Dretske's square and reaching out until he touches cube #5. Here we might want to say that S perceives cube #5 even though it wouldn't be right to say that he differentiates this cube from its immediate environment. He obviously doesn't differentiate cube #5 visually and he needn't differentiate it in any other way either; S can touch cube #5, and thereby perceive it, without tracing its boundaries. If this is right then what we have here is a counterexample to the first premiss of the indirect argument.

Is it plausible, however, that S perceives cube #5? We might argue that S can't be perceiving cube #5 precisely because he fails to differentiate it. In that case, what does S perceive? If we don't want to say that S perceives cube #5, and are also reluctant to say that he doesn't perceive anything, the obvious alternative would be to say that he perceives the whole square of which cube #5 is a part. But the perceptual differentiation requirement is still in trouble if this is what we say. It's obvious what the problem is: from S's perspective the square isn't differentiated from *its* immediate environment any more than cube #5 is differentiated from its immediate environment. To differentiate the square S would have to trace its spatial

¹² Which is not to say that it is immune to counterexample. Presumably, the differentiation requirement wouldn't apply if I see cube #5 through a keyhole. Thanks to John Campbell for the example. Dretske describes a similar example of his own as a '*limiting case*' since it is one in which the object seen 'has no environment' (1969: 26). The upshot is that the differentiation requirement 'becomes inoperative when nothing appears to S that is not part of D' (1969: 27).

boundaries but he can touch the square without doing any such thing; as long as he is touching some part of the square he is touching the square, regardless of whether he has any sense of where the square ends and the rest of the world begins. So we are still left with an example of object perception without perceptual differentiation.

That takes care of the first premiss of the indirect argument for SPR. Does the second premiss fare any better? This is the claim that it's not possible to differentiate an object without perceiving space or spatial properties. A familiar argument in support of this claim is that in order to perceive two objects as numerically distinct from each other at a given time it's necessary to perceive them as being at different spatial locations, and this is clearly a form of spatial perception.¹³ So, for example, to perceive the cup into which I am pouring coffee as distinct from the cup next to it is to perceive the two cups as being in different places; perceiving the two cups as being in different places is not just a means of distinguishing them but a necessary condition for distinguishing them. Even if the cups are indistinguishable in every other respect, I can still differentiate them spatially.

There is a helpful statement of this line of thinking in the following passage from Henry Allison's commentary on the first *Critique*:

Kant can be taken to be arguing that in order to be aware of things as numerically distinct from one another, it is necessary to be aware, not only of their qualitative differences, but also of the fact that they are located in different places. In other words, the representation of place, and therefore of space, functions within human experience as a necessary condition of the possibility of distinguishing objects from one another. (1983: 83)¹⁴

Putting all of this together we now have the following indirect argument for SPR: in order to know that b is P by perceiving that b is P one

¹³ Notice that this is only an argument for the second premiss of the indirect argument if it is assumed that perceptually differentiating two objects is a matter of perceiving them *as* distinct from each other. One might question this assumption but I'm not going to make anything of this here. There are other, more obvious problems with the indirect argument, and, in any case, it is not obviously false that someone who perceptually differentiates two things must, in *some* sense, perceive them as distinct. It all depends on how much is built into the notion of perceiving two things as distinct from each other.

¹⁴ The insistence on awareness of qualitative differences is clearly a slip on Allison's part. The implication of his argument is that one can be aware of two objects as numerically distinct from each other just by being aware of them as being in different places. The objects don't need to be qualitatively distinct, and even if they are one doesn't need to be aware of their qualitative differences in order to be aware of them as numerically distinct.

must (leaving aside cases of secondary epistemic perceiving) perceive *b* itself. Assuming that *b* itself is an object, one can't perceive *b* itself without perceptually differentiating *b* from its surroundings, and therefore without perceptually differentiating *b* from other objects in its vicinity. In order to differentiate *b* from a neighbouring object *a* one must perceive *b* and *a* as being in different places. Finally, in order to perceive two objects as being in different places one must have the capacity to perceive the locations of objects, that is, a capacity for spatial perception. That is why possession of this capacity is a necessary condition for perceiving that *b* is *P* or for knowing on this basis that *b* is *P*. But it's only a background necessary condition; it's only a necessary condition for the acquisition of perceptual knowledge by being a necessary condition for perceptual differentiation.

Everything now depends on whether it's true that it's not possible to differentiate objects other than on the basis of location. On the face of it, this is straightforwardly false. As Daniel Warren points out, 'if I know that (at a given time) *a* is pink and *b* is not pink, then I can infer that *a* and *b* are numerically distinct. I don't need to consider the spatial features of *a* and *b* in order to distinguish them' (1998: 187). As we have seen, colour isn't a specifically spatial property of objects so when one distinguishes two objects by reference to a difference in colour one isn't distinguishing them on a specifically spatial basis. In such cases, location is irrelevant, and it doesn't help to point out that if *a* and *b* were perfect replicas a difference in location would be the only remaining basis for differentiating them. We don't regularly come across perfect replicas in our experience, and we don't need to appeal to spatial considerations in order to distinguish objects that *aren't* perfect replicas.¹⁵

None of this really should come as much of a surprise. It's sometimes suggested that we must differentiate objects spatially because two objects must *be* in different places in order to *be* numerically distinct. The implication is that the basis on which we differentiate objects must track the conditions under which they are in fact distinct, and one question is whether this tracking requirement is plausible. But even if we think that something along these lines is defensible Warren's objection still goes through. After all, it's false that numerically distinct objects must be in different places; a statue and a lump of clay can be in exactly the same place at the same time and still be numerically distinct. In

¹⁵ These are all points that Warren also makes in his very helpful discussion. See Warren 1998: 188–9.

that case, why should one think that representing them as being in different places is the key to differentiating them? All we need to do in order to differentiate them is to regard them as having different persistence conditions, and we can do that without viewing them as being in different places. It seems, then, that the indirect argument for SPR is once again in trouble. It is mistaken when it claims that objects must be differentiated spatially.

In fact, this is slightly unfair to the indirect argument. As I have been emphasizing, what is supposed to be at issue here is *perceptual* differentiation. The question is whether it is possible to differentiate two objects perceptually without perceiving them as being in different places, yet Warren's example works at the level of judgement rather than perception. I can know that a is pink and that b is not pink without ever having perceived them, and I can infer that they must be numerically distinct without ever having perceived them. For example, if someone tells me that a is pink and that b is not pink I can infer that a and b are distinct. But this isn't an example of perceptual differentiation. For Warren, the capacity to distinguish objects is 'a capacity for making certain kinds of *judgements*' (1998: 192). This is not how the indirect argument for SPR conceives of differentiation, so the possibility of judging that a and b are numerically distinct without judging that they are in different places is neither here nor there as far as this argument is concerned. The same goes for the statue and the lump of clay. Although one can judge that the two are distinct without judging that they are in different places, this doesn't show that one can perceive two things as distinct at a given time without perceiving them as being in different places. Since the distinction between the statue and the clay is drawn at the level of thought rather than at the level of perception, this case doesn't have any direct bearing on SPR. The statue and the lump of clay aren't *perceived* as distinct and SPR can explain why not: they aren't perceived as distinct because they aren't perceived as being in different places.

In arguing in this way, I'm relying on an intuitive distinction between thought and perception.¹⁶ I'm assuming, that is, that one can judge that a and b are numerically distinct without perceiving them as distinct, and that one can perceive a and b as numerically distinct without judging

¹⁶ This distinction comes into play in many other contexts. In the Müller-Lyer illusion, for example, the natural thing to say is that one perceives the lines as being of different lengths even though one doesn't think or judge that they are of different lengths.

that they are numerically distinct. Both of these assumptions seem quite reasonable. The first of them is uncontroversial and the second ought to be uncontroversial on any substantive conception of judgement. Infants and animals can presumably differentiate objects perceptually without making judgements of numerical distinctness. The important question, therefore, is whether it's possible to come up with a case in which it's not in dispute that two objects *a* and *b* are being differentiated perceptually and in which it's also plausible that *a* and *b* are not perceived as being in different places. If such cases are possible then they really would be a problem for SPR in a way that Warren's example is not.

One way of generating such a case would be to modify Warren's example so that it's clearly perceptual. Suppose that I can *see* that (at a given time) *a* is pink and *b* is not pink. Am I seeing that *a* and *b* are numerically distinct without seeing them as being in different places? The idea would be that I see *a* and *b* as numerically distinct by seeing them as qualitatively different rather than by seeing them as being in different places. Or, to take another example, suppose that the cup into which I'm pouring coffee is standing on a saucer. I don't in any ordinary sense see the cup and the saucer as being in 'different places' but I still see them as distinct things. In general, we don't have any trouble perceiving the distinctness of adjacent objects yet it's hard to believe that this has much to do with perceiving a difference in location. In such cases, given the proximity of the objects, we rely on other factors to differentiate them, factors such as colour differences and figural goodness.¹⁷ Against this background, the privileging of spatial location as the basis for perceptual differentiation has no basis.

In fact, these considerations aren't decisive. For a start, one couldn't see *a* as pink and *b* as not pink without seeing them as being in different places; it isn't possible to perceive both pinkness and its absence in the same region of space at the same time. The most we can say, therefore, is that seeing *a* and *b* as being in different places needn't be the *means* by which one sees them as numerically distinct. It doesn't follow that seeing them as being in different places isn't a background necessary condition for seeing them as numerically distinct by seeing one of them as pink and the other as not pink. Seeing *a* and *b* as being in different places is what makes it possible to distinguish them on the basis of

¹⁷ Elizabeth Spelke and Gretchen Van de Walle describe 'colour similarity, smoothness of edges and figural goodness' as 'Gestalt relationships that specify the boundaries of stationary objects for adults' (1993: 137).

colour, and it wouldn't be possible to distinguish them on this basis without perceiving them as being in different places. The perception of a difference in location is, in this sense, an enabling condition for distinguishing them on the basis of colour. Similarly, while it seems unlikely that I distinguish the cup and the saucer by perceiving them as being in different places, it's still true that they don't occupy exactly the same region of space and that it wouldn't be possible to perceive them as distinct without perceiving them as being in different places. In particular, it wouldn't be possible to see their shapes and edges, and to differentiate them on this basis, without seeing where they are in relation to each other.¹⁸ If this is right, then the indirect argument's conception of what is necessary for perceptual differentiation is still in play.

For a better example of perceptual differentiation that isn't based on the perception of a difference in location we need to bring in senses other than sight. Suppose that I'm in a hotel room and overhear a man and a woman arguing in the next room. I can hear them as distinct without hearing them as being in different places in the next room. Since I don't merely judge that they are distinct but hear them as distinct this looks like a case of perceptual differentiation without the perception of a difference in location; in general, we don't need to hear two people as being in different places in order to tell them apart on the basis of audition. Another example: as I walk into a kitchen, I can smell an onion and a banana. I smell them as distinct things even though I have no idea where they are. What tells me that they are distinct is the way they smell rather than any perceptual awareness of where they are located in relation to each other. Again, this looks like a counterexample to the indirect argument for SPR. It's not just that perceiving the onion and the banana as being in different places isn't the means by which I smell them as distinct; it doesn't even appear to be a background necessary condition for smelling them as distinct.

Is it clear that these examples are a problem for the indirect argument for SPR? It's important to remember that this argument only says that it isn't possible to perceive two *material objects* as distinct at a given time without perceiving them as being in different places. Material objects are bounded, three-dimensional space-occupiers. Cups and people are objects in this sense, sounds or smells are not. If this is right, the indirect argument can happily accept that it's possible to distinguish

¹⁸ It's also true, of course, that the perception of shapes and edges is a form of spatial perception.

two sounds or two smells on the basis of perception without relying on any perceived difference in location. This is significant because one might think that when I overhear a man and a woman arguing in the next room it's the distinction between the *sounds* they are making rather than a distinction between *people* that is given to me perceptually. Although I take it two different people are involved in the argument, the distinction between them is drawn at the level of thought rather than at the level of perception. Similarly, in the other example, I am able to differentiate perceptually between the *smell* of the onion and the *smell* of the banana, which is not to say that the distinction between the onion and the banana is perceptual. In neither example, therefore, is it uncontroversial that two material objects are being perceptually differentiated without being perceived as being at different locations.

But why should one think that perception is only distinguishing between sounds in the first case and smells in the second? What would be wrong with saying that I literally hear the two people in the next room as distinct, or that I literally smell the onion as distinct from the banana? One thought might be that this wouldn't be the right thing to say because there is an element of inference in both cases. I infer that the argument is between two people, and there is no guarantee that I'm right about this; for all I know, I could be hearing one person mimicking an argument between a man and a woman. Equally, it isn't my sense of smell which tells me that what I can smell is *a* banana and *an* onion, as distinct from several onions and several bananas or a strange dish containing both onions and bananas. That's why it's more natural to say in such cases that what one can smell is *onion* rather than *an* onion. And one can smell onion in this sense even if there are in fact no longer any onions present; to smell onion is to smell a certain type of smell rather than to smell a material object.

On its own, the fact that there is an element of inference involved when one distinguishes between two things doesn't show that the distinction isn't perceptual. It's worth noting that in the first example I don't consciously infer that what I can hear is an argument between two people. If there is an inference here it's unconscious, and it's arguable that unconscious inference is involved in a lot of ordinary perception. That is why we are not yet in a position to conclude that when I hear two people arguing in the next room I don't *hear* them as distinct from each other. Even if there is in fact only one person there I am still under the impression that I can hear two people. To maintain that this isn't a case of genuinely *perceptual* differentiation one would have to

maintain that one doesn't hear people at all and that the proper objects of audition are sounds rather than material objects. Similarly, one might hold that the proper objects of olfaction are smells. If one doesn't hear or smell material objects *at all*, then one can't literally hear or smell one material object as distinct from another.

The problem with arguing in this way is that we don't actually think that material objects can't be the proper objects of audition or olfaction. Take the case of audition. We speak of hearing strange sounds in the middle of the night, but we also sometimes describe ourselves as hearing the things that are responsible for the sounds which we can hear. In such cases, hearing sounds isn't an alternative to hearing material objects, it's a way of hearing the objects that make them. We talk about hearing sounds when we are unfamiliar with their sources. When their sources are familiar, it's much more natural to describe ourselves as hearing their sources; we hear two *people* arguing or the *floorboards* creaking. The same goes for smell. Although we might talk about smelling a strange smell, or of smelling onion, there is also such a thing as smelling *an* onion.

If material objects can be the objects of audition or olfaction it's hard to see why it shouldn't be possible to hear or smell two material objects as distinct from each other. There are creatures which discriminate material objects largely on the basis of smell, and we often have no difficulty discriminating material objects on the basis of hearing. If, when I hear two people arguing in the next room, I hear the people themselves and hear *them* as distinct from each other without hearing them as being in different places, it can't be right that perceptual differentiation of material objects *must* be based on the perception of a difference in location; the counterexamples can't all be dealt with by denying that the objects being differentiated are material or by insisting that none of them is a bona fide example of perceptual differentiation. The obvious conclusion is that the perception of space is not the only possible basis for differentiation of material objects in perception.

Although it looks like this is the end of the road for the indirect argument for SPR, there are a couple of additional moves which proponents of this argument might want to make before giving up completely. To get a handle on the first of these moves, let's go back to the suggestion that it's possible to hear two objects as distinct from each other. Although, as we have seen, hearing two people as distinct is a form of perceptual differentiation, we also need to recognize that it is significantly different from some other forms of perceptual

differentiation. We can bring out one important difference by looking at an example of John Campbell's:

Suppose that outside in the street you hear a bulldog and a pekingese start to fight; you hear the whole thing from the initial growls and yaps to the final triumphant or defiant yowls, and can tell exactly when one dog leaves off and the other begins. Well, you might say, that bulldog put up a good fight, but it had no chance. Here the demonstrative 'that bulldog' is an auditory demonstrative—you may at no point have looked at the scene—but it refers to an ordinary physical object. (1997: 65)

In such cases, how are we able to put together just those bits of auditory information which come from a single source, say one particular bulldog? Again, location can't be the key since one might easily fail to hear the dogs as being in different places; auditory localization just isn't that precise. Instead, Campbell's explanation draws on the notion of a 'schema'.

[I]f your dog has been in many fights, you may know exactly how he sounds in one, whereas the lay person listening to the scene might have some trouble in sorting out where one dog leaves off and the other begins. You have built up a schema for your dog's performance, which you can use to filter out all but the auditory information coming from it, a schema which the lay person lacks. (1997: 66)

A schema in this sense is specific rather than general; it is, as it were, an auditory picture of what one particular dog sounds like in a dog fight. In Campbell's example, this is what is needed to differentiate the two dogs. Often, however, auditory differentiation is underpinned by a general schema. In my example, I am able to distinguish between the man and the woman in the next room because I know what men and women sound like. I don't need a schema for that particular woman's performance in an argument to be able to distinguish her from the man she is arguing with. I might not have the faintest idea who they are but I can still tell them apart.

Possession of a specific or general schema is usually the result of one's familiarity with whatever thing or kind of thing is in question. I have a schema for my dog's performance in a fight because I have heard it in one before, and I'm obviously also familiar with what men and women generally sound like when they speak. With the help of the appropriate schema, it's possible to hear two things as distinct without hearing them as being in different places. Let's call this *schema-based* perceptual differentiation. Loosely speaking, all auditory discrimination of material

objects is schema-based or at least schema-involving. Although it's possible to hear two objects as distinct by hearing them as being in different places, I have already suggested that we count as hearing objects rather than just sounds when we have a conception of the kinds of object involved. But to know what kind of object is the source of a particular sound is to have a kind of schema for that object. Without a schema, we wouldn't be hearing the *object* at all and so couldn't be differentiating *it* from other objects on the basis of location.

Vision doesn't work like this since not all visual discrimination is schema-based. Here is an example of schema-based visual discrimination: as I look at a city skyline I see a bunch of skyscrapers and other buildings clustered together. From a distance, it might not be immediately apparent where one building ends and the next one begins, but I can still see two adjacent buildings as distinct because I have a grasp of the appropriate grouping or demarcation principles; for example, I have a good enough grasp of the relevant principles to be able to tell by looking that the Gherkin and the Lloyds Building in London's financial district are distinct buildings even though, from a certain distance and angle, they appear to run into each other. On the other hand, I can also see two strange and unfamiliar objects as distinct from each other even though I have no idea what they are; they look different just by appearing in different regions of space with a gap in between. I don't need a schema in order to see them or in order to see them as distinct things, so here we have an example of visual discrimination that isn't schema-based.

Armed with the distinction between schema-based and non-schema-based perceptual differentiation we could try the following move on behalf of the indirect argument for SPR: instead of insisting that all perceptual differentiation requires the perception of space we could simply insist that non-schema-based or non-schema-involving perceptual differentiation wouldn't be possible without the perception of space, that is, without the perception of a difference in the locations of the differentiated objects. We can now ignore the auditory and olfactory counterexamples to the unmodified indirect argument on the grounds that they are examples of schema-based or schema-involving perceptual differentiation. Where we don't have a schema to go on, we have to fall back on spatial considerations. This is what happens when we discriminate unfamiliar objects by sight or touch. In these cases, we are *presented* with the distinction between one object and another, and it's only possible for objects to be presented as distinct in

this schema-independent sense by being presented as being in different places.

Unfortunately, this attempt to rescue the indirect argument only solves one problem by creating another. Although it's not implausible that genuinely schema-independent perceptual differentiation must be spatial this doesn't help the indirect argument unless object perception requires schema-independent perceptual differentiation. But why should we think that? In order to perceive that *b* is *P*, or to know that *b* is *P* by perceiving that *b* is *P*, I must be able to perceive *b*. As we have already seen, it's not obvious that I only perceive *b* if I differentiate it from other objects in its immediate environment. It's even less obvious that I only perceive *b* if, in addition, my differentiation of it is schema-independent. The perceptual differentiation requirement makes the most sense in the visual case but there is no reason to think that I only know that *b* is *P* by seeing that *b* is *P* if my visual differentiation of *b* is schema-independent. I can see that the Gherkin is glinting and thereby know that it is glinting even if I need a schema to differentiate it from adjacent buildings. So the net result of focusing on perceptual differentiation which isn't schema-based is to make it more plausible that perceptual differentiation wouldn't be possible without the perception of space but correspondingly less plausible that perceptual differentiation is necessary for the perception of objects.

Where does this leave SPR? To recap, the indirect argument's defence of this requirement relies on two controversial premisses. The first states that it's not possible to perceive an object without differentiating it perceptually from other objects in its environment. The second states that it isn't possible to differentiate two objects perceptually without perceiving them as being in different places. What has emerged is that both premisses are too strong because they are too general. One can touch an object without differentiating it from neighbouring objects and one can hear or smell one object as distinct from another without hearing or smelling the two objects as being in different places. What is true is that it isn't possible to *see* an object without visually differentiating it from other objects in its environment and that one couldn't see two objects as distinct without *seeing* them as being in different places.¹⁹ Even in the case of schema-based visual differentiation the perception of numerical difference wouldn't be possible without the perception of a difference in location. Although I don't see the Gherkin and the

¹⁹ But see n. 12 for a qualification.

Lloyds Building as distinct by seeing them as being in different places, this doesn't alter the fact that I couldn't even begin to apply the various grouping principles by means of which I differentiate them unless I see them as occupying different though adjacent regions of space; even adjacent objects must be seen as being in different places in order to be seen at all.

The obvious reaction to the discussion so far would be to regard it as an illustration of the dangers of looking for highly general enabling conditions for the acquisition of perceptual knowledge. We can now see why, instead of thinking in terms of enabling conditions of epistemic perception *per se*, we should be thinking in terms of enabling conditions for knowing that something is that case by seeing that it's the case, by hearing that it's the case, and so on. Instead of trying to defend an unrestricted Spatial Perception Requirement, we should be trying to defend a restricted version of this requirement, something along the lines of: the visual perception of space or spatial properties is an enabling condition for knowing that b is P by seeing that b is P. As I have been emphasizing, the multi-levels approach to the explanation of the possibility of perceptual knowledge can and should take this restriction in its stride. On the other hand, it seems almost inevitable that there will be those who think that I haven't given SPR in its unrestricted form a run for its money, and it's worth considering what someone who says something along these lines might have in mind.

One thing they might have in mind is that the emphasis on *differentiation* in the indirect argument is misplaced. While it might be possible to smell or hear two things as distinct from each other without smelling or hearing them as being in different places, it isn't possible to hear or smell something without hearing or smelling it as being *somewhere*. Even sounds and smells, which aren't material objects, must be perceived as located in space or as coming from directions in space. So the obvious move should be to abandon the idea that it isn't possible to perceive two objects as distinct from each other without perceiving them as being in different places and to concentrate instead on the idea that it isn't possible to perceive an object *at all* without perceiving it as located. In effect, this amounts to the recommendation that we should give up on the indirect argument for SPR and settle instead for the direct argument, in which perceptual differentiation isn't the key.

I will have much more to say about the direct argument below. First there is another proposal which I would like to discuss. According to this proposal, the problem isn't that the indirect argument concentrates

on perceptual differentiation but that it concentrates on the wrong kind of perceptual differentiation. One question is whether the perception of space is an enabling condition for the perception of objects as distinct from *each other*. Let's call this *object-object* perceptual differentiation. A different question is whether the perception of space is an enabling condition for the perception of objects as distinct from *the self and its states*. Let's call this *subject-object* perceptual differentiation. In this context, 'states' are 'mental states' or 'states of consciousness'. We might be reluctant to agree that object-object perceptual differentiation requires the perception of space without disagreeing that subject-object perceptual differentiation requires the perception of space. If the perception of space is what makes it possible to perceive something as distinct from oneself and one's states, and so to differentiate it in this sense, then there is still hope for the idea that the perception of space is an enabling condition for the acquisition of perceptual knowledge by means of epistemic perception. More cautiously, there is still hope for this idea *if* we are prepared to accept that in order to know that b is P by perceiving that b is P it is necessary not only to perceive b itself but also to differentiate b itself from oneself and one's states. Before moving on to the direct argument, therefore, I want to take a moment to look at this variation on the indirect argument.

3.4 PERCEPTUAL DIFFERENTIATION (II)

The idea that subject-object differentiation is tied to the perception or 'representation' of space is often attributed to Kant. According to Allison, for example, Kant thinks that 'the representation of space is the condition or presupposition of human awareness, but not any conceivable awareness of objects as distinct from the self and its states' (1983: 83). Combining this with the alleged connection between spatial representation and object-object differentiation we get the claim that 'it is only because I already have the capacity to represent objects as spatial or as in space that I also have the capacity to represent these same objects as distinct from myself... and as numerically distinct from each other' (1983: 344). Let's not quibble about whether this is a good interpretation of Kant and concentrate instead on whether there is a good argument for this conception of subject-object differentiation.

Here is one argument: to differentiate an object from the self and its states can only be to conceive of it as distinct from the self and its

states. To conceive of something perceivable in this way is to conceive of it as capable of existing unperceived. We can only make sense of perceivable but unperceived existence if we think that perception has enabling conditions. These are conditions which a perceiver must meet in order to perceive what is there to be perceived. Specifically, we must have the idea that in order to perceive what is there to be perceived the perceiver must be in the right place at the right time. We can then think that what is there to be perceived will exist unperceived if there are no perceivers in the right place at the right time. This assumes, however, that both perceivers and the objects we conceive of as capable of existing unperceived are located in space and time. To think of something as spatially located is to represent it as spatially located, so we can now see why subject–object differentiation requires the representation of space. If we failed to represent subjects and objects as spatial, we wouldn't have a grasp of the spatial enabling conditions of perception and so wouldn't be able to distinguish objects from subjects.²⁰

Despite its popularity, there are lots of problems with this argument. I will mention just two. To begin with, it's not clear why we have to think of enabling conditions of perception as spatial. An alternative would be to think in terms of *receptivity*. Someone who does this sees the course of his experience as 'simultaneously determined by the way the world is and his changing receptivity to it' (Evans 1980: 91). If he isn't receptive, he won't perceive what is there to be perceived. And if nobody else is receptive either then what is there to be perceived will exist unperceived. Since there isn't any mention of space or spatial location in this story it looks as if it provides us with a non-spatial way of making sense of existence unperceived. But once we agree that it's possible to make sense of existence unperceived in non-spatial terms, it's going to be hard to maintain that subject–object differentiation requires the representation of space.

One response to this objection would be to argue that we can't make sense of existence unperceived just in terms of failures of receptivity. If I am to suppose that I am not now perceiving the occurrence of ϕ -ing simply because I'm not receptive, and that I would perceive ϕ -ing if I were to become receptive, there needs to be a criterion for my becoming receptive. But in a non-spatial scheme there is no criterion for my becoming receptive to ϕ -ing other than my perceiving ϕ -ing. According

²⁰ This is basically Evans's argument in his 1980 paper 'Things without the Mind'. See Cassam 2005 for further discussion.

to Evans, what this means is that ‘the factor accounting for the presence or absence of perception of perceptible phenomena’ won’t be subject to ‘significant empirical control’ in a scheme using receptivity (1980: 94). If empirical control is what we are after, only a spatial scheme will do. In order to perceive the ϕ -ing that is there to be perceived I must be in the right place, but my present experience isn’t the only criterion of my being at a certain location; recent experiences of adjacent places are also relevant.

The implication of this argument is that we can hang on to the suggestion that subject–object differentiation requires the representation of space as long as we accept that this form of differentiation must be subject to significant empirical control. If we don’t accept this requirement, or can show that such control is possible in a scheme using receptivity, then we should continue to be sceptical about the suggestion that the representation of space is necessary for subject–object differentiation. I’m not going to pursue these issues here, partly because there doesn’t appear to be a conclusive argument either way but also because there is a much simpler objection to the argument from subject–object differentiation. The objection is that, even if it is successful on its own terms, this argument doesn’t help SPR. The most the argument shows is that subject–object differentiation requires the *conception* of subjects and objects as spatially located whereas what is at issue in SPR is whether the *perception* of space is necessary for subject–object differentiation. It’s easy to lose track of this point if one talks about the role of the ‘representation’ of space in subject–object differentiation. Spatial thinking and spatial perceiving can both be described as forms of ‘spatial representation’, so the thesis that subject–object differentiation requires the ‘representation’ of space is ambiguous. Once we recognize this ambiguity it’s easy to see that an emphasis on spatial thinking won’t be appropriate in an argument for the Spatial *Perception* Requirement.

Just as there are two ways of understanding what it is to ‘represent’ space, so there are two ways of understanding subject–object differentiation. To differentiate something from the self and its states can either be a matter of perceiving it or of conceiving of it as distinct from the self and its states. The differentiation which SPR is interested in is perceptual rather than conceptual. The question is whether the *perception* of something as distinct from the self and its states requires the *perception* of space, and it is not an answer to this question to show that the *conception* of something as distinct from the self and its states requires the *conception* of it as spatially located. At any rate, it is not

an answer to this question as long as we think that perception and conception can be separated in the way that I have been assuming.

In general, we don't have any difficulty with the idea that perceiving is different from conceiving. We understand, for example, that it's possible to perceive one line as longer than another without conceiving of it as longer than the other. This is what happens in the Müller–Lyer illusion. To say that one perceives one line as longer than the other is to make a point about the content of one's experience, about how things seem to one perceptually, rather than about what is going on at the level of thought or judgement. Similarly, to say that one perceives an object as distinct from the self and its states or as distinct from some other object which one also perceives is to make a claim about how things seem to one perceptually. Just as an experience can seem to be an experience of objects which are distinct from each other, so it can seem to be an experience of objects which are distinct from oneself and one's states. When this happens, we have a case of perceptual subject–object differentiation.

What does this form of differentiation involve? Even if we accept that there is such a thing as perceptual subject–object differentiation we might still think that it needs to have a conceptual dimension. The Müller–Lyer illusion provides a helpful contrast. As well as taking it to show that it's possible to perceive one line as longer than the other without conceiving of it as longer than the other, we might also take it to show that it's possible to perceive one line as longer than the other without having the concept of one line's being longer than another. If this is correct then perceiving one line as longer than another is a form of what might be called 'non-conceptual' perceiving. But perceiving an object as distinct from the self and its states looks like a more sophisticated cognitive achievement, and this makes it more difficult to think of perceptual subject–object differentiation as non-conceptual. Maybe it's possible to perceive an object as distinct from oneself and one's states without actually conceiving of it as distinct from oneself and one's states, but is it right to allow for the possibility of perceptual subject–object differentiation in a case in which the perceiver doesn't have the concept of an object's being distinct from himself and his states?

That depends on, among other things, one's underlying conception of perceptual content. In the Müller–Lyer illusion what it comes to say that one experiences one line *as* longer than the other is that one's experience won't be veridical unless one line *is* longer than the other.

The conditions under which an experience is veridical are what Peacocke calls its 'correctness conditions', and it's helpful to think of perceptual differentiation in terms of correctness conditions.²¹ To experience a and b *as* distinct from each other is to have an experience that is veridical only if a and b *are* distinct from each other. Equally, to experience a and b *as* distinct from oneself and one's states is to have an experience that is only veridical if a and b *are* distinct from oneself and one's states. If the correctness of one's experience didn't depend on whether this is true of a and b it couldn't properly be described as an experience of a and b *as* distinct from oneself and one's states.

In a way, this isn't much of an explanation of perceptual subject-object differentiation. What we have so far is the idea that to experience something as distinct from the self and its states, and therefore as mind-independent or objective, is for one's experience to have a certain correctness condition. Now we need an explanation of what it is for an experience to have this correctness condition. As it happens, this is precisely where Kant thinks that the perception of space comes in. His proposal is that I perceive a and b as mind-independent, and therefore have an experience that is veridical only if a and b are mind-independent, just if I perceive a and b as 'outside me' (A23/B38). To perceive something as outside me is to perceive it as being 'in another region of space from that in which I find myself' (*ibid.*). By perceiving a and b as outside me in this sense I perceive them as distinct from me, and I couldn't perceive them as distinct from me and my states without perceiving them as outside me. Since perceiving something as being in another region of space from that in which I find myself is a form of spatial perception, we can conclude that spatial perception is necessary for perceptual subject-object differentiation.

For the purposes of this argument, a and b needn't be physical objects. Suppose that they are sounds. I'm going to take it that to hear a sound as 'outside me' is to hear it as *coming from* another region of space. In these terms, Kant's idea is that hearing a sound as distinct from me and my states is closely related to hearing it as outside me in this sense. Specifically, the thought is that I hear a sound as mind-independent *by* hearing it as outside me. Hearing it as outside me is therefore the *means* by which I hear it as mind-independent. And if I didn't hear the sound as coming from another region of space my experience wouldn't qualify as one that is only veridical if what I hear is mind-independent.

²¹ See Peacocke 1992: 107–8 for more on correctness conditions.

The implication is that hearing a sound as outside me is not just a background necessary condition but a straightforwardly necessary condition for hearing it as mind-independent.²² More generally, an experience can only *seem* to be of something objective by being a spatial experience.

We now have an argument for SPR, as well as an explanation of the possibility of non-conceptual subject–object perceptual differentiation. Assuming that epistemic perceiving requires perceptual subject–object differentiation, spatial perception is necessary for epistemic perceiving because it is necessary for perceptual subject–object differentiation. Yet perceiving something as outside me doesn't require spatial concepts or the thought of something's being distinct from me and states. So if all it takes to perceive something as mind-independent is to perceive it as outside me, perceptual subject–object differentiation needn't have a conceptual dimension. Finally, there is now what looks like a solution to the problem of generality. The argument for SPR from perceptual subject–object differentiation doesn't only work for sight or touch or for the perception of physical objects. We have just seen that the perception of space is necessary for the perception of sounds as objective, and we can run the same line for smell: to perceive a smell as distinct from oneself and one's states is to smell it as being in another region of space from that in which one finds oneself.

How good is this argument for SPR? Here is one problem: when I hear the sound of my own voice as I speak, I don't hear it as coming from another region of space from that in which I find myself. I hear it as coming from within my own body, which isn't outside me. Yet this doesn't prevent me from hearing the sound as distinct from myself and my states. I'm conscious of the sounds I make as I speak as distinct from me and I'm also conscious of them as public in a way that I'm not aware of my own states of consciousness as public. In this sense, I'm conscious of the sounds I make as I speak as distinct from my states of consciousness. But if I'm conscious of the sounds I make as distinct from me and my states of consciousness then it can't be true that it's only possible to hear sounds as mind-independent if one hears them as coming from somewhere else. So, for example, knowing that my voice

²² This claim has a bearing on Nudds's remark that 'we can hear a sound without hearing it to have any spatial properties at all' (2001: 215). What it suggests is that hearing a sound without hearing it to have any spatial properties would amount to hearing it without hearing it as mind-independent.

is hoarse by hearing that my voice is hoarse is a form of epistemic perceiving for which the perception of space doesn't appear to be a necessary condition.

Since this problem for SPR arises because I'm taking 'outside me' to mean 'outside my body', one solution would be to deny that this is what 'outside me' means. Maybe this is what Kant is getting at when he claims that 'I distinguish my own existence as that of a thinking being from other things outside me—among them my own body' (B409). On this account, perceiving something as coming from within my own body is still perceiving it as coming from outside me because my body is outside me. On the other hand, this can't easily be reconciled with the idea that for something to be outside me is for it to be in another region of space from that in which I find myself. The region of space in which I find myself is the region of space occupied by my own body, so it's hard to see how my body can be outside me. But if my body isn't outside me, then Kant's conception of perceptual subject-object differentiation is in serious trouble. As well as the auditory counterexample which I have just been discussing, we might also imagine experiencing one's own body as distinct from oneself and one's own states of consciousness. In such a case, one would be aware of one's own body as mind-independent without being aware of it as being outside itself.

It might seem that there is an obvious solution to these problems. The obvious solution would just be to take 'outside me' to mean 'in space'. In this sense, my own body *is* outside me and I'm conscious of it as such. And I'm also conscious of the sounds I make as I speak as outside me to the extent that I'm conscious of them as coming from *somewhere*. I don't hear them as coming from nowhere so my awareness of them still looks like a form of spatial awareness; directional perception is still a form of spatial perception. A genuine counterexample to the Kantian conception of what is necessary for perceptual subject-object differentiation would be a case in which one is aware of something as distinct from oneself and one's states without being aware of it as having *any* spatial properties. As yet we haven't managed to come up with a case of this kind.

Even so, we should be disappointed if this is the best we can do for SPR. The original idea was that the perception of space or of specifically spatial properties is necessary for perceiving that *b* is *P* or for knowing that *b* is *P* by perceiving that *b* is *P*. The perception of sounds is a problem for this idea because we can hear sounds as mind-independent without hearing them as shaped, extended, or as coming from a region

of space which isn't occupied by one's own body. SPR is now being defended against this objection on the basis that when one hears the sound of one's own voice one still hears it as possessing at least one spatial property, the property of coming from somewhere. But even if hearing a sound as coming from within my own body is a form of spatial perception, the perception of what might be called 'inner' space, it's only minimally spatial compared with other forms of spatial perception. This means that the most we can say is that the perception of something as distinct from the self and its states must be minimally spatial. It needn't involve the perception of shape, extension, or other specifically spatial properties. Although this emasculated version of SPR isn't trivial, it looks a lot less interesting and substantial than the robust requirement we started out with.

In fact, the position we are now in is in some ways even worse than this discussion suggests. We need to bear in mind that the point of insisting that spatial perception is necessary for perceptual subject-object differentiation can't just be to insist that there can be no perceptual subject-object differentiation without spatial perception. The point is presumably also to suggest that the fact that an experience is an experience of something as distinct from the self and its states is to be *explained* by reference to the fact that it is a spatial experience. Yet in the case in which one hears the sound of one's own voice, it's not clear in what sense hearing it as coming from somewhere within one's own body is supposed to explain one's perception of the sound as mind-independent. It could just be an ineliminable feature of auditory perception that it is directional, from which it doesn't follow that its possession of this feature is doing any explanatory work when it comes to accounting for its capacity to represent its objects as distinct from the self and its states. Maybe the very limited sense in which auditory perception requires the perception of space has nothing to do with its capacity to represent objects in this way.

The line of thinking threatens to make the notion of perceptual subject-object differentiation inexplicable if not ineffable. Starting with the idea of a link between the content of an experience and its correctness conditions, we found ourselves having to give an account of the conditions under which experiences of objects as distinct from the self and its states have their distinctive correctness conditions. Although the suggestion that this is to be explained in spatial terms looked promising at the outset, much of its initial plausibility has evaporated. If we take what I have been saying about auditory perception to

its logical conclusion we will be forced to concede that perceptual subject–object differentiation doesn't have any genuinely explanatory connection with the perception of space. Yet other explanations of subject–object differentiation fare even worse, and this might lead one to endorse the thesis that no informative account can be given of what it takes to experience something as distinct from the self and its states.

I'm not suggesting that we should endorse this thesis. It's not an attractive position to end up in, and there is no getting away from the intuitive pull of the thought that the perception of something as distinct from oneself and one's states does have *something* to do with the perception of space. The difficulties begin when we try to explain exactly *what* these forms of perception have to do with each other. Since we haven't been able to explain the connection in a way that sustains a robust version of SPR, that is one reason for not making SPR depend on Kant's conception of perceptual subject–object differentiation. Another more obvious reason is this: we are supposed to be talking about the enabling conditions for perceiving that *b* is *P* and thereby knowing that *b* is *P* but why should we think that knowing that *b* is *P* by perceiving that *b* is *P* requires perceptual subject–object differentiation? In primary epistemic perceiving, one can't perceive that *b* is *P* without perceiving *b* itself, yet even if *b* itself *is* distinct from the self and its states it doesn't follow that in order to perceive it and thereby to know something about it one must perceive it *as* distinct from the self and its states. This is just another way of saying that the Kantian argument for SPR can't even get started unless epistemic perceiving requires perceptual subject–object differentiation.

I don't want to rule out the possibility that a plausible argument can be constructed in support of this requirement. My point is simply that this is another potentially problematic dimension of the Kantian argument. Again, the obvious lesson is that one should look elsewhere for a defence of SPR. And this brings us back to the direct argument for SPR. From the perspective of this argument, none of the difficulties the indirect argument runs into should come as much of a surprise. Whether it is object–object or subject–object perceptual differentiation that the indirect argument concentrates on, the basic problem is the same: either we find that epistemic perception can do without perceptual differentiation or that perceptual differentiation can do without spatial perception. That is why one might think that the emphasis on perceptual differentiation is misplaced. From the perspective of the direct argument, the discussion of perceptual differentiation is an unnecessary detour since

SPR can be argued for much more directly. So now would be a good time to take a closer look at the direct argument.

3.5 THE DIRECT ARGUMENT AND PRIMARY QUALITIES

The direct argument claims that it's not possible to perceive a spatial object without perceiving any of its spatial properties or without perceiving space. To put it another way, it claims that there is a direct connection between the perception of spatial objects and spatial perception. Unlike the indirect argument, the direct argument isn't concerned with perceptual differentiation. In the case in which I hear two people arguing, the fact that I hear them as arguing *in the next room* shows that my perception of them is still spatial. In the context of the direct argument, it doesn't matter that I don't differentiate them by hearing them as being in different places. It also doesn't matter in this context if it is possible to differentiate objects of sight or touch other than on the basis of a perceived difference in location. Given that object perception is necessary for primary epistemic perception, the existence of a direct connection between object perception and spatial perception *already* implies that spatial perception is an enabling condition for perceiving that something is the case and thereby knowing that it is the case.

When I first discussed the direct argument I mentioned the worry that it trivializes SPR. I now want to explain why this isn't so. The worry was that it's trivial that *spatial* perception is necessary for the perception of *spatial* objects. We can see that this isn't trivial by thinking about auditory perception. When I hear two people arguing in the next room without perceiving them in any other way I don't hear their shapes even though I might infer something about their shapes from how they sound. Extension is another spatial property of spatial objects that can't be heard. It's true that if I'm aware of two people arguing in the next room I am aware of where they are, and that location is a spatial property. But what if I hear them without having any sense of where they are? Given that people are spatial objects, and that there are no other spatial properties of spatial objects that we are capable of hearing, what we would then have is a case in which I perceive two spatial objects without perceiving *any* of their spatial properties. And if this is a genuine possibility then it's not even true, let alone trivially true, that

the perception of spatial properties is necessary for the perception of spatial objects.²³

It only seems obvious that the perception of spatial objects requires spatial perception because the perception of spatial objects by sight or touch must be spatial and because it's easy to forget that seeing or touching an object aren't the only ways of perceiving one. When I see an object, I see it as shaped, extended, and located in relation to myself or to other objects in its immediate environment. Even though I might misperceive some of the object's spatial properties there is no such thing as seeing an object without seeing any of its spatial properties. In addition, seeing an object usually involves seeing the empty space around it, as well as the region of space it occupies. In contrast, the objection to the direct argument is that there is such a thing as hearing an object without hearing any of its spatial properties and no such thing as hearing the empty space around an object or the region of space which it occupies.

This objection suggests that the direct argument faces the same basic problem as the indirect argument. Both arguments make excessively general claims about the role of spatial perception in object perception, and in both cases the excessive generality becomes apparent when one thinks about auditory perception. So is there any way of defending the direct argument against this criticism? Let's return to the suggestion that it's possible to hear two people arguing without being aware of where they are. If we accept that this is possible we are also going to have to accept that it's possible to perceive a spatial object without perceiving any of its spatial properties. One option for the direct argument would therefore be to insist that, contrary to what I have been assuming, it isn't possible to hear a spatial object without perceiving its location.

It's not difficult to see why one might think otherwise. Given that auditory localization isn't very precise, we often hear things without having any very clear sense of where they are. Imagine being woken up in a strange hotel room in the middle of the night by the sound of two people arguing and not being able to tell exactly where the noise is coming from. It's tempting to regard this as a case in which one hears two spatial objects without hearing any of their spatial properties,

²³ Nudds is someone who appears to think that this is a genuine possibility. He gives the example of hearing a sound 'which initially one hears as coming from a certain direction but which, as one listens, one ceases to hear as coming from anywhere' (2001: 214).

and therefore as a problem for the direct argument. But the obvious point to make in response is that not having a *clear* sense of where something is doesn't amount to not having *any* sense of where it is. However disorientated one might be, and however vague one's sense of where the argument is taking place, one must surely hear it as taking place *somewhere*. Even if one can't narrow its location down to a specific spatial region one must hear it as being more or less to the left or the right, in front or behind, above or below.

This amounts to the proposal that hearing must be at least roughly directional so everything now depends on whether this proposal is correct. The problem is this: localization is one aspect of hearing and it seems conceivable that someone should lose the capacity to localize without losing the capacity to hear. The direct argument implies that this isn't possible whereas the objection I have been considering insists that it is possible. How is this disagreement to be resolved? A good start would be to figure out what notion of possibility is at issue in this dispute. Suppose that it is logical possibility. In that case, it would be difficult to sympathize with the direct argument. It doesn't appear to be built into the concept of hearing that it must be directional, and there is no other reason for thinking that non-directional hearing is logically impossible. The alternative is to maintain that it is physically or physiologically impossible to hear something without hearing it as being anywhere but it's not clear that this is correct either. In any case, since it is an empirical question whether non-spatial hearing is physically or physiologically possible, the net result of relying on these notions of possibility will be to make it difficult to think of the connection between the perception of space and the perception of spatial objects as an *a priori* connection.

It seems, then, that auditory object perception is still a problem for the direct argument. Although there isn't much doubt that this form of perception usually involves some form of spatial awareness, the link with spatial perception isn't clear enough or strong enough for the purposes of the direct argument. Perhaps, in that case, this argument should deny that we literally hear material objects, as distinct from the sounds they make. If hearing isn't a way of perceiving objects at all, then the direct argument doesn't have to worry about the possibility of hearing objects without hearing them as being anywhere. But I have already argued that we don't and shouldn't deny that objects can be objects of audition. If I'm right about this, then the direct argument is between a rock and a hard place: it has to choose between

denying something that we don't usually have any trouble accepting or insisting on an implausibly strong connection between auditory object perception and spatial perception. Either way, the direct argument looks shaky.

In fact, the situation isn't quite as bleak for SPR as the discussion so far suggests. To see why not, we need to remember that SPR is the thesis that the *capacity* to perceive space or spatial properties is an enabling condition for knowing about the external world by means of epistemic perception. This means that, unlike the direct argument, it is not committed to denying that it is possible to perceive a spatial object without *actually* perceiving any of its spatial properties or without perceiving space. What it is committed to denying is that object perception doesn't require a capacity for spatial perception. So while the possibility of hearing a spatial object without hearing any of its spatial properties is a problem for the direct argument, it isn't necessarily a problem for SPR. This requirement would still be defensible if a background capacity for spatial perception is needed in order to hear a spatial object without hearing any of its spatial properties. On this account, which is the one I want to defend, spatial perception in the 'capacity' sense is a background enabling condition for object perception even if there are examples of object perception which do not involve the exercise of this capacity.

Why would one need a capacity for spatial perception in order to hear a spatial object without hearing any of its spatial properties? How can *non-spatial* auditory perception depend on a capacity for spatial perception? In order to make any progress with these questions we need to return to the proposal that material objects can't be objects of audition. Although I have already rejected this proposal I haven't identified any of the considerations which might be thought to support it. I'm now going to suggest that we can only understand why the most compelling argument in its favour fails if we accept that a background capacity for spatial perception is necessary for non-spatial object perception. What I'm suggesting, in other words, is that a proper understanding of what it takes for material objects to be objects of audition will lead us to an argument for SPR. As we will see, this argument fares better than either the direct or the indirect argument.

The best case for thinking that material objects can't literally be heard is this: material objects are fundamentally space-occupiers, and their most basic properties are ones that they have in virtue of being space-occupiers. Here one might think of shape and extension. The

problem with hearing is that it doesn't provide us with perceptual access to such properties. As Campbell points out:

Listening to an object can tell you a lot about it. The mechanic listening to the car engine ticking over, or the doctor listening to his patient cough, can thereby find out a lot about the behaviour of the thing. But it is arguable that these properties are in a sense less fundamental to the objects in question than the basic volumetric properties—size, shape, solidity and so on—about which vision and touch inform us; it might be said that it is only through vision and touch that we have information about the properties of the object in virtue of which it counts as a space-occupant. (1997: 68)

Although Campbell isn't claiming that material objects can't be heard, the considerations he outlines can be used as the basis of a simple argument for this view. The argument is that in order to hear material objects, it would be necessary to hear their basic volumetric properties; their basic volumetric properties can't be heard, so material objects can't be heard. Even if one has a schema for an object this doesn't alter the fact that its size, shape, and solidity aren't objects of audition. Material objects can only be perceived through vision and touch because only vision and touch provide us with perceptual access to those properties of material objects in virtue of which they count as material objects.

What is plausible about this argument is its insistence that hearing doesn't inform us about the basic volumetric properties of material objects. These properties of material objects are what Locke calls their 'primary qualities', so another way of putting this would be to say that hearing doesn't inform us about the primary qualities of material objects. However, it is an overreaction to conclude from this that material objects aren't objects of audition. Just because we don't hear the size, shape, and solidity of material objects it doesn't follow that we don't hear them at all. Still, there is something right about the principle that underlies the overreaction. The underlying principle is that the ability to perceive material objects must be somehow connected to the ability to perceive their primary qualities. Specifically, the thought is that there would be no justification for supposing that someone who is simply incapable of perceiving properties such as size, shape, and solidity is nevertheless capable of perceiving material objects. The reason that this doesn't force one to concede that material objects can't be heard is that audition provides us with perceptual access to objects whose basic or primary properties are perceptually accessible to us in other ways. Vision and touch provide us with access to these properties, and it is only because

this is so that we are comfortable with the idea that material objects are objects of audition.

A good way of putting this would be to say that our ability to hear material objects is parasitic upon our ability to see or touch them. When we see or touch material objects we have perceptual access to their primary qualities. If primary qualities were perceptually inaccessible material objects would be perceptually inaccessible. Yet properties such as size, shape, and solidity can be inaccessible to *audition* without being *perceptually* inaccessible. Material objects can only be heard because their size, shape, and solidity can be seen or felt, and this means that a being with hearing but no touch or vision would be incapable of perceiving material objects. If we can conceive of such a being at all we would have to conceive of it as one which can only hear sounds.

Putting all this together, we now have the basis of a straightforward argument for SPR: being able to perceive material objects by sight or touch is a necessary condition for being able to perceive them at all. It isn't possible to see or touch a material object without perceiving any of its spatial properties or the region of space it occupies. To perceive the spatial properties of an object or the region of space it occupies is to exercise a capacity for spatial perception. Possession of this capacity is therefore a necessary condition for being able to perceive material objects at all. In brief: object perception requires spatial perception, but *not* because whenever one perceives a physical object one must actually perceive one or more of its spatial properties. This is the requirement that gets the direct argument into trouble. The alternative I am recommending avoids this difficulty by insisting on a less direct connection between spatial perception and auditory object perception.

Since the recommended argument for SPR trades on the idea that one can only perceive material objects if one can perceive their primary qualities I might as well call it the *primary qualities argument* (PQA). This argument explains why the distinction between the actual perception of spatial properties and the capacity to perceive spatial properties is important for the purposes of defending SPR. If one can hear an object without actually perceiving any of its spatial properties, the 'actual perception' requirement is too strong. If, as PQA claims, one can only hear objects without actually perceiving any of their spatial properties because one has the ability to perceive their spatial properties in other ways then it is the capacity to perceive spatial properties which is functioning as a necessary condition for object perception. But it's only a background necessary condition because it doesn't actually need to

be exercised in every case of object perception. This is the work which talk of 'background' conditions is doing. Finally, we also now have a response to the problem of generality. Because the primary qualities argument allows that auditory object perception needn't be spatial it can't be accused of ignoring the differences between the senses. On the other hand, its conception of the way in which auditory object perception is parasitic upon sight and touch allows us to stand by SPR. So it turns out that we don't have to choose between abandoning SPR and making excessively general claims about what is necessary for epistemic perception.

We now have a positive answer to the question of necessity. PQA makes it plausible that spatial perception is a background necessary condition for object perception and therefore also a background necessary condition for perceiving that *b* is *P* and thereby knowing that *b* is *P*. Although the sense in which spatial perception is necessary for auditory object perception is different from the sense in which it is necessary for visual or tactile object perception, PQA shows that all genuine object perception depends in some way on the capacity to perceive the spatial properties of objects or the region of space they occupy. The difference between the different forms of object perception is in the way that they depend on spatial perception. In some cases the connection is relatively direct and straightforward whereas in other cases it is indirect and anything but straightforward.

The remaining questions are the question of a priority and the question of foundations. The first of these asks whether SPR can be established without any empirical investigation. Since I am now representing the primary qualities argument as the basis of SPR another way of asking the question of a priority would be to ask whether this is an empirical or an a priori argument for SPR. Assuming that SPR is true, we can also ask what makes it true. This is the question of foundations. Since the two main options in response to this question are idealism and realism we can now make the question of foundations a bit more precise by asking which of these two responses is supported by PQA. Finally, having tackled these questions we can return to the issue of minimalism and to the prospects for defending a fully-fledged multi-levels response to (HP_{pk}).

In the next and last section of this chapter I'm going to argue as follows: first, I will attempt to show that PQA sustains a realist rather than an idealist conception of the basis of SPR. Next, I will suggest that spatial perception should be regarded as at least a weakly a

priori enabling condition for epistemic perception and for the resulting perceptual knowledge. Finally, I will argue that this conception of the status of SPR is incompatible with minimalism in relation to (HP_{pk}). A more difficult question is whether PQA supports moderate anti-minimalism or whether it counts in favour of extreme anti-minimalism. I'm going to argue that PQA doesn't settle this question but that there are good independent grounds in favour of moderate anti-minimalism.

3.6 FOUNDATIONS, A PRIORITY, AND MINIMALISM

The idealist thinks that what makes spatial perception a necessary condition for knowing about the external world by perceptual means is the 'cognitive structure of the human mind' (Allison 1983: 29) and nothing else. This is Kant's view. He claims that space and time are the two 'forms' of human sensibility and that it is only because this is so that spatial and temporal perception are 'necessary conditions under which alone objects can be for us objects of the senses' (A29). It can't be anything about the intrinsic nature of the objects we perceive which accounts for SPR since, according to transcendental idealism, these objects are not intrinsically spatio-temporal. We only experience them as spatio-temporal because of the way we are, not because of the way they are.

What does it mean to describe space as a 'form of human sensibility'? One possibility is that the forms of sensibility are the conditions under which alone objects can be for us objects of the senses, that is, necessary conditions for object perception. On this reading, idealism can't be claiming that what makes spatial perception necessary for object perception is the fact that space is a form of sensibility. This would be equivalent to the claim that what makes spatial perception necessary for object perception is the fact that spatial perception is necessary for object perception, and this can't be what the idealist has in mind. But if idealism lacks a substantive account of what makes spatial perception a form of sensibility then it also lacks a substantive account of what makes spatial perception a necessary condition for knowing about the external world by perceptual means.

This suggests that Kant shouldn't be thinking of the forms of sensibility simply as necessary conditions for object perception. If he wants to explain these conditions by reference to the forms of sensibility

he needs some other account of the sense in which space is a form of sensibility. His alternative depends on a distinction between the ‘matter’ and ‘form’ of perception. According to Kant, the senses present us with arrays of sensations (the matter) and are so constituted that they can only receive spatio-temporally ordered sensations (the form).²⁴ This account raises a lot of questions which I don’t intend to discuss in any detail here because it doesn’t look at all promising. One question is whether we can make sense of the distinction between the matter and the form of perception. Even if we can, it’s not clear how the fact that *sensations* are received in a certain order is supposed to explain the fact that the perception of *objects* must be underpinned by capacity to perceive properties such as extension, shape, and location. Since objects aren’t sensations, and sensations aren’t shaped or extended, there appears to be a large gap between Kant’s *explanandum* (a claim about how objects must be perceived) and his *explanans* (a claim about how sensations must be received).

Transcendental idealism isn’t the only form of idealism and the defects in Kant’s account of SPR aren’t necessarily defects in idealism *per se*. Nevertheless, if we argue for SPR in the way that I have been recommending we shouldn’t, or at least needn’t, be idealists. In essence, the primary qualities argument for SPR turns on the idea that material or spatial objects are perceptually accessible only if their primary qualities are perceptually accessible, and that the perception of primary qualities is a form of spatial perception. But the capacity to perceive their primary qualities is only a necessary condition for the perception of material objects because it is in virtue of their possession of primary qualities that they count as material objects in the first place. If it were not in the nature of the *objects* to have spatial properties such as shape, extension, and location it would not be in the nature of *object perception* to be spatial or to be underpinned by a capacity for spatial perception.

There isn’t the slightest trace of idealism in this account. There is no suggestion in PQA that it is the cognitive constitution of the perceiving subject that explains why SPR holds. It is the nature of the objects, along with a principle linking the perception of objects to the capacity to perceive their basic properties, which is doing all of the explanatory work. This is realism not idealism. What the linking principle says is that the capacity to perceive material objects is tied to the capacity to

²⁴ See Falkenstein 1995 for a reading of Kant along these lines.

perceive those of their properties which make them material objects.²⁵ This would be compatible with idealism if one thought that the nature of material objects is itself a reflection of the cognitive structure of the human mind but it's not clear why one should think that that's true. As long as our cognitive constitution doesn't account for the spatiality of primary qualities it doesn't account for SPR; instead, realism says that we've got to be able to perceive spatial properties of material objects in order to perceive them at all because material objects *are* fundamentally spatial.

This is as much as I am going to say about the question of foundations. Where does my realist response to this question leave the question of a priority? On a weak reading, a priori conditions or requirements are ones that can be established non-empirically. On a strong reading, a priori conditions can only be established non-empirically; they can't be established empirically. Suppose, then, that we divide PQA into two main components: first, there is the linking principle connecting what it takes to perceive material objects with what it takes for an object to be material. Then there is the idea that the basic properties of material objects are volumetric properties such as size, shape, solidity, and location. With respect to each of these components of PQA we can ask whether it can be, or can only be, established non-empirically. This will enable us to determine whether SPR states an a priori requirement on the acquisition of perceptual knowledge and, if so, whether it states a strongly or a weakly a priori requirement.

Starting with the linking principle, this doesn't look like something that can be established by empirical investigation. The principle doesn't say that the perception of material objects is, as a matter of fact, tied to the perception of primary qualities. The claim is that *what it is* to be able to perceive physical objects is to be able to perceive their primary qualities, and this looks like just the sort of constitutive claim about perception that can only be established by armchair reflection if it can be established at all. So what we have here is something that is strongly rather than weakly a priori. A more difficult question is whether what I have identified as the second main component of PQA is weakly or strongly a priori. On the one hand, one might think that questions

²⁵ This is similar to Peacocke's principle that 'if an account of what is necessarily involved in something's having a certain property makes reference to some substantial condition which must be met with by things which have it, a thinker's mental representation of that property must be suitably sensitive to the existence of this substantial condition' (1993: 171). The crux is what 'suitably sensitive' means.

about the nature of material objects must be empirical since armchair reflection can't tell us what material objects are actually like. On the other hand, PQA is concerned with *what it is* for an object to be a material object and this doesn't look like a straightforwardly empirical question.

The way forward with this issue is to remember that size, shape, and solidity count as basic properties of material objects because they are properties of material objects which they must have in order to count as space-occupants. This assumes that material objects are, in essence, space-occupants and this comes out in the way that they are individuated: no two material objects of the same kind can occupy exactly the same region of space at the same time. Since Locke states and endorses this principle in the *Essay* let's call it Locke's Principle.²⁶ What it is for an object to be a material object is for it to be governed by Locke's Principle, yet it's hard to see how this principle could be established empirically. Like the linking principle, Locke's Principle looks like a constitutive principle that can only be established by armchair reflection. Crudely, this principle defines what material objects are and thereby implicitly defines what is going to count as a primary quality since the primary qualities of material objects are consequential upon their being space-occupants.

We now have an argument for the strong a priority of SPR: as long as we think of this requirement as grounded in reflection on what it is to be a material object and on what it is to be able to perceive material objects we are effectively thinking of SPR as grounded in non-empirical considerations. Unless we think that the primary qualities argument isn't the only argument for SPR or that the principles on which PQA relies can be established empirically, we won't be able to make sense of the idea that we can know by empirical investigation that spatial perception is a background necessary condition for object perception and therefore for epistemic perception. Strictly speaking, I haven't ruled out the possibility of basing SPR on empirical grounds and I haven't proved that Locke's Principle can't be established empirically but this doesn't matter. Even if SPR can be established empirically, it still comes out as weakly a priori as long as it can also be established non-empirically along the lines that I have just been describing.

Where does this leave the debate between minimalism and anti-minimalism? Minimalism claims that perceptual knowledge has no

²⁶ See book II, chapter 27, section 1 of Locke's *Essay*.

substantive a priori enabling conditions, and that interesting philosophical Level 3 explanations of perceptual knowledge are therefore impossible if we accept the assumption that the distinguishing feature of such explanations is that they are non-empirical. Since I have just identified SPR as a weakly and perhaps also a strongly a priori enabling condition of perceptual knowledge that's pretty much the end of minimalism; it turns out that armchair philosophy can tell us something about what makes perceptual knowledge possible. The remaining options are moderate anti-minimalism (the view that Level 3 explanations are possible but not necessary), and extreme anti-minimalism (the view that Level 3 explanations are both possible and necessary). I'd like to bring this chapter to a close by giving a short argument for moderate anti-minimalism.

Suppose that we are interested in (HP_{pk}) and that we get as far as Level 2 in dealing with it. By this point, we have identified basic primary epistemic perceiving as a means of acquiring perceptual knowledge of the external world, and have removed apparent obstacles that stand in the way of this kind of perceiving. Is it necessary to go any further? That depends on what going further is supposed to be necessary for. Extreme anti-minimalism can be read as saying that a response to (HP_{pk}) that doesn't go to Level 3 is *incomplete*, but we saw at the end of Chapter 1 that this can't be right. Consider the following analogy: I am asked how it is possible to get from London to Paris in less than three hours, and my answer is that it is possible to do this by catching the Eurostar. I have now explained how it is possible to get from London to Paris in less than three hours even though I haven't said anything about the enabling conditions for getting from London to Paris by train. I haven't said anything about this because there is no need. Similarly, once I have explained that it's possible to know that the cup is chipped by seeing that it is chipped I have already explained how it is possible to know that the cup is chipped; there is no obvious sense in which my explanation is incomplete without any mention of enabling conditions so it isn't necessary to say anything about such conditions in order to provide a satisfactory answer to (HP_{pk}) . The further questions about a priori enabling conditions are ones that we *can* go into but they are not ones which we are *required* to go into in order to answer (HP_{pk}) .

Perhaps, in that case, what is at issue isn't exactly completeness but a certain kind of philosophical or epistemological satisfaction. Perhaps the extreme anti-minimalist's thought is that a Means Response to (HP_{pk}) that identifies means of knowing but not enabling conditions

for knowing by the proposed means can't be philosophically satisfying. But why not? What counts as a philosophically satisfying answer to (HP_{pk}) is always going to be relative to one's philosophical interests, and while one might think that explaining how perceptual knowledge is possible is fundamentally a matter of explaining what makes it possible, one might also think that explaining how perceptual knowledge is possible is fundamentally a matter of overcoming apparent obstacles to its existence. From this perspective, the source of one's epistemological satisfaction will be what happens at Level 2 rather than at Level 3 of the multi-levels response, and nothing that happens at Level 3 seems essential.

When we think about (HP_{pk}) in this way, it becomes apparent that extreme anti-minimalism is unacceptably dogmatic. The problem with asserting that a Level 3 explanation is necessary is that one might have no interest in what makes perceptual knowledge possible. If one is satisfied with what one has at Level 2 any further discussion is bound to strike one as superfluous. The important point, therefore, is not that we must say something about a priori enabling conditions if we are serious about answering (HP_{pk}) but that there are a priori enabling conditions of perceptual knowledge and that philosophical reflection can tell us what they are if we are interested. This is a case for moderate rather than extreme anti-minimalism. At the same time, the identification of spatial perception as a genuinely a priori enabling condition for the acquisition of perceptual knowledge shows that minimalism is not a viable alternative to extreme anti-minimalism. The only viable alternative is moderate anti-minimalism.

4

Perceptual Knowledge (II): Concepts

4.1 THE ROLE OF CONCEPTS

As I pour my first cup of coffee of the day I see that the cup is chipped and thereby come to know that it is chipped. Seeing that the cup is chipped is, like feeling that the cup is chipped by running one's finger along its rim, a form of epistemic perceiving. Perceiving that a proposition about the external world ('the cup is chipped') is true is a means of coming to know that it is true. If I come to know that the cup is chipped by seeing that it is chipped my knowledge is a form of perceptual knowledge. Perceptual knowledge is empirical knowledge though not all empirical knowledge is perceptual; I can also discover that the cup is chipped by being told that it is chipped or by inferring that it is chipped from other things I came to know by empirical means.

What makes it possible to see that the cup is chipped? This is a question about the background necessary conditions for seeing that the cup is chipped. In the last chapter, I said that there are many such conditions, including ones that can only be established by empirical means. But there are also background necessary conditions for seeing that the cup is chipped, and for epistemic perceiving generally, that can be established non-empirically. The perception of space is one such condition. Are there any others? In this chapter I want to examine the suggestion that the a priori enabling conditions for epistemic perceiving include not just narrowly *perceptual* capacities but also the capacity to *think* in certain ways; specifically, they include the capacity to think categorially. This is the point of what I have referred to in previous chapters as the *Categorial Thinking Requirement* (CTR). Categorial thinking is thinking by means of categorial concepts, and what CTR implies is that coming to know that a proposition about the external world is true by perceiving that it is true depends on one's possession of certain specific categorial concepts.

On the face of it, CTR doesn't have much going for it. It's true that in order to see that the cup is chipped I need the concepts *cup* and *chipped* but these aren't categorial concepts. *Cup* is a sortal concept, the concept of a sort of object, while *chipped* is a characterizing concept, the concept of a property or characteristic of objects. In contrast, categorial concepts like *substance*, *unity*, *plurality*, and *causality* are neither sortal nor characterizing, and it's not obvious why one would need any categorial concepts in order to see that the cup is chipped. If this is right then the most that one can say is that epistemic seeing and other forms of epistemic perceiving depend on one's possession of a repertoire of sortal and characterizing concepts. It looks as though categorial concepts have nothing to do with it.

Suppose, however, that it turns out that it wouldn't be possible for one to have any sortal or characterizing concepts if one lacked the capacity to think categorially. In that case, CTR would be in much better shape. Its point would be to draw attention to the background necessary conditions for possession of concepts like *cup* and *chipped* and, by implication, for epistemic perceiving. The proposal would be that being able to think categorially makes it possible for one to see that the cup is chipped by making it possible for one to have sortal and characterizing concepts in the first place. Alternatively, it might be thought that one needs categorial concepts simply in order to perceive objects like cups. If a capacity for categorial thinking is necessary for object perception, and object perception is a component of primary epistemic perceiving, then that would go a long way towards vindicating CTR.¹ Either way, there is now the prospect of a range of arguments in support of an initially unpromising Level 3 requirement on the acquisition of perceptual knowledge.

Before spelling out and assessing these arguments for CTR we need a better account of what makes a concept categorial. It's not enough just to give examples of such concepts; we need something more general. On one reading, categorial concepts are *formal* concepts.² On this account, the relationship between a formal concept like *substance* and a sortal concept like *cup* is the relationship between a 'determinable' and one of its 'determinates': being a cup is a particular way of being

¹ In primary but not in secondary epistemic perceiving one perceives that b is P by perceiving b itself. This is how Dretske distinguishes between primary and secondary epistemic seeing. See Dretske 1969: 79–80 and Chapter 1.4 above.

² My account of formal concepts follows Wiggins's account. See Wiggins 1980: 63–4 and 1997: 417–18.

a substance.³ Formal concepts are, in this sense, abstractions from sortal concepts; they are what Wiggins calls ‘dummy sortal concepts’ (1997: 418) and cannot be used to say what something is in the way that genuine sortal concepts can be used to say what something is. So, for example, the question ‘what are you pouring coffee into?’ can be satisfactorily answered by saying ‘a cup’ but not by saying ‘a substance’.⁴

On a different reading, categorial concepts are ‘categories’ or ‘pure concepts of understanding’ in roughly Kant’s sense. Kant says that the categories are ‘concepts of an object in general, by means of which the intuition of an object is regarded as determined in respect of one of the logical functions of judgement’ (B128). It’s not clear at this point what this means but what is clear is that Kant thinks that there are twelve categories in all: *unity, plurality, totality, reality, negation, limitation, substance, causality, community, possibility, existence, and necessity*. Although these are all formal concepts in a loose sense of ‘formal concept’, they are not all dummy sortals; for example, being a cup is not a particular way of being a totality, and the relationship between the concept of a cup and the concept of negation is not the relationship between a determinate and a determinable. By the same token, there are dummy sortal concepts that are not Kantian categories; being a cup is a particular way of being an object but *object* is not officially one of Kant’s categories.

If categorial concepts can either be understood as dummy sortals or as Kantian categories then categorial thinking can either be understood as thinking by means of dummy sortals or by means of Kant’s categories. I am going to take it that CTR is the thesis that categorial thinking in one or other of these senses is an a priori enabling condition for epistemic perceiving. With this in mind, CTR raises three questions corresponding to the three questions about SPR that were the focus of the previous chapter. The first is whether categorial thinking is a genuinely necessary condition for epistemic perceiving. This is the

³ See Prior 1949 for more on the determinable/determinate distinction.

⁴ As Wiggins puts it, ‘a formal concept like *entity* or *substance* has no autonomous individuating force of its own, and must be variously supplemented, wherever it appears in contexts of identification, according to the kind of individual in question. If supplementation yields all sorts of different principles of individuation according to the compliant and context, and if there are no restrictions on how it is filled out except for context-relative or merely categorial ones, then a concept is too “high”, too unspecific, to count as an answer to the question *what is x?* It is not a sortal concept’ (1980: 63–4).

question of necessity. Then there is the *question of a priority*: can CTR be established non-empirically if it can be established at all? Finally, there is the *question of foundations*: what makes it the case, if it is the case, that the acquisition of perceptual knowledge by means of epistemic perception must be underpinned by a capacity for categorial thinking? Is it the structure of the knowing mind, the structure of the objects of perceptual knowledge, or something else?

The most important of these questions is that of necessity. Since there is no obvious connection between, say, seeing that the cup is chipped and being able to think categorially, why should one think that CTR states a genuine requirement? I have already sketched two answers to this question. The first says that categorial thinking is necessary for epistemic perceiving because it is necessary for object perception and because object perception is part and parcel of primary epistemic perceiving. I'm going to call this the *argument from object perception* for CTR. The second argument, which I'm going to call the *argument from epistemic perception*, says that categorial thinking is necessary for epistemic perceiving because epistemic perceiving is conceptual and because it wouldn't be possible for one to grasp any of the sortal or characterizing concepts that are needed for epistemic perceiving if one lacked the capacity to think categorially. Let us now take a closer look at these arguments, starting with the argument from object perception.

4.2 OBJECT PERCEPTION

Why should one think that one needs to be able to think categorially in order to perceive objects? After all, infants and animals can presumably perceive objects but do we really want to say that they have concepts like *substance* and *causality*? We could try to get round this difficulty by insisting that if infant and animals lack categorial concepts then they can *see* objects but they can't *perceive them*. On this account, *perceiving* a cup is concept-involving in a way that simply *seeing* a cup is not. But why should one think that one needs to be able to perceive objects in this technical sense if one is to acquire knowledge of the external world by perceiving epistemically? In order to know that the cup is chipped by visual means why must I 'perceive' the cup and not just see it? And even if object perception in the concept-involving sense is required, why should one think that one needs *categorial* concepts, as distinct from sortal or characterizing concepts, in order to 'perceive' objects?

The distinction between seeing and perceiving an object is one that Dretske introduces in his account of non-epistemic or 'simple' seeing.⁵ For Dretske, simple seeing is concept-free since it implies nothing about the conceptual resources of the perceiver, and belief neutral in the sense that simply seeing X is compatible with no beliefs about X. In both of these respects seeing X is like stepping on X; one can step on X without recognizing it as an X or having any beliefs about X. In contrast, *perception* is '(either by stipulation or common understanding) cognitively loaded' since 'some degree of recognition or categorization is essential to our perception of things' (2000: 100). If this is right then 'it is by no means obvious that one must perceive something in order to see it' (ibid.). Not knowing what a cup is prevents one from perceiving a cup but not from seeing a cup.

The obvious question about this distinction is whether non-epistemic seeing can be quite as simple as it implies. As we saw in the last chapter, Dretske claims in *Seeing and Knowing* that a subject S sees an object D if and only if D is 'visually differentiated from its immediate environment by S' (1969: 20). What does it take for S to differentiate D in this sense? Suppose that it turns out that sortal or categorial concepts are needed for visual differentiation. In that case one would have to conclude that even simple seeing is cognitively loaded.⁶ It would be a further question whether it is cognitively loaded in exactly the same sense as perceiving but there would be no getting away from the fact that simple seeing couldn't be concept-free. So even if one only needs to see rather than to perceive the cup in order to see that it is chipped, there is now the prospect of an argument to the effect that one needs sortal or categorial concepts simply in order to see the cup.

But why should one think that one needs sortal or categorial concepts in order to differentiate and therefore to see objects? Here is an argument for the *sortal-dependence* of visual differentiation: to differentiate an object is to isolate it in experience, to draw its spatio-temporal boundaries, and one can't do that without knowing what kind of this it is.⁷ For example, when I see the cup into which I am pouring coffee I

⁵ See Dretske 2000.

⁶ Or, to put it another way, one would have to conclude that there is no such thing as 'simple' seeing.

⁷ This argument for the sortal-dependence of visual differentiation is similar to Wiggins's argument for the sortal-dependence of individuation. To individuate an object x is to single it out, to 'isolate x in experience; to determine or fix upon x in particular by drawing its spatio-temporal boundaries and distinguishing it in its environment from

visually differentiate it from the saucer it is standing on but not from its handle. On what basis? The sortal-dependence thesis implies that I can only visually delineate the cup because I recognize it as a cup and because my concept of a cup is the concept of something that incorporates a handle rather than the concept of something that incorporates a saucer. But if visual differentiation is, in this sense, sortal-dependent, then so is simple seeing.

What about categorial concepts? Where do they fit in? Suppose that sortals like *cup* are taken to be determinations of the formal concept of an object, the concept of a 'bounded, coherent, three-dimensional object with some particular way of behaving, coming to be, being, being qualified and passing away' (Wiggins 1997: 417). Given that visual differentiation is sortal-dependent, and that sortal concepts are determinations of formal concepts, we might think that visual differentiation is at least implicitly dependent on formal concepts. Isolating and categorizing the cup into which I am pouring coffee is 'part and parcel with treating it as a thing with some specific way of behaving' (Wiggins 1997: 413–14), and to treat a thing in this way just is to be thinking of it as an object. So it can't be that visual differentiation is sortal-dependent without also requiring a basic form of categorial thinking.

Are these arguments any good? If they are, then the only way to hold on to the concept-freedom of simple seeing would be to give up the idea that this kind of seeing requires visual differentiation. In fact, however, there is no need to do this because the argument for sortal-dependence isn't any good. It can't be true that visual differentiation requires knowledge of the sortal concept that the object in question falls under because one can visually differentiate an object even if one has no idea what kind of object it is or has a false belief about which sortal concept it falls under. So, for example, I don't need to know what a cup is in order to see a cup as distinct from the saucer it is standing on. The cup and the saucer aren't attached to each other and don't move as one in the way that the cup and its handle are attached to each other and move as one. In distinguishing one object from another, and putting different properties together as properties of a single object,

other things of like and unlike kinds (at this, that and the other times in its life history); to articulate or segment reality in such a way as to discover *x* there' (1980: 5). The strongest version of the sortal-dependence thesis says that it isn't possible to single something out without knowing what kind of thing it is (which sortal concept it falls under). This is not Wiggins's considered view; he thinks that individuation is sortal-dependent in a slightly weaker sense. See the next note for more on this.

qualitative and spatial considerations are more important than sortal concepts. Sortal concepts are applied to objects that have already been singled out by our visual systems, and this implies that the singling out 'is done at a more primitive level than the one at which we have the application of sortal concepts' (Campbell 2002: 73). That is why we have no trouble understanding how infants and animals can visually differentiate objects.⁸

If visual differentiation isn't sortal-dependent then we can't argue that it requires categorial thinking because sortal concepts are determinations of categorial concepts. Indeed, even if visual differentiation were sortal-dependent the case for CTR would still be weak. Just because the concept *cup* is a determination of the concept *object* it's not obvious that someone who thinks of something as a cup must implicitly be thinking of it as an object. At the very least one would need some guarantee that someone who has the determinate has the corresponding determinable. So while there might be a sense in which someone who has managed to single out what is in fact a cup must be treating it as a thing with some specific way of behaving, we have not yet seen why *treating* something in this way must be a matter of *thinking* of it either as a cup or as an 'object'. The thesis that visual differentiation requires categorial thinking is therefore in no better shape than the thesis that it requires sortal thinking.

So far, then, we still don't have an argument for CTR. I can't see that the cup is chipped without seeing the cup itself but I can see the cup itself without knowing that it is a cup or being able to think of it as falling under a formal or categorial concept. Perhaps, in that case, we need to go back to the distinction between simple seeing and perceiving. Instead of trying to undermine this distinction by arguing that simple seeing is conceptual maybe it would be better to just accept it and concentrate instead on trying to show why one would need to be able to perceive the cup and not just see it in order to know that it is chipped

⁸ Campbell's point is similar to one made in Ayers 1974. In response to Ayers Wiggins concedes that 'it is perfectly possible for a thinker to qualify as singling something out, as being in the right *rapport* for that, without *knowing* what he is singling out or having any in the context informative answer to the question what he has singled it out as' (1980: 218). In what sense, then, is individuation sortal-dependent? In the sense that, even in cases of sortal-ignorance, the singling out of an object is only possible in so far as 'one is somewhere en route to grasp of a sortal concept' (Wiggins 1997: 414). But it's doubtful whether the fact that an infant is en route to a sortal concept explains its capacity to single out and differentiate objects. Many non-human animals can differentiate but do we want to say that they are en route to grasp of sortal concepts?

by seeing that it is chipped. One can argue that perceptual knowledge of objects wouldn't be possible without object perception and that object perception needs to be underpinned by a capacity for categorial thinking without needing to deny the existence of more primitive forms of seeing that are genuinely concept-free.

Kant is a good example of someone who thinks that we must be able to perceive objects in this sense in order to gain knowledge of them. To perceive an object in the cognitively loaded sense is, in Kant's terminology, to 'intuit' one. An intuition is an immediate sensible representation of an object, and one of the defining principles of Kant's epistemology is that intuitions and concepts are the 'elements of all of our knowledge' (A50/B74). So in order to *know* that the cup is chipped by seeing that it is chipped one would need to intuit or perceive the cup. But intuitions are conceptual. An intuition is a 'perception', and 'all possible perceptions' are 'subject to the categories' (B164–5). Without a capacity for categorial thinking it wouldn't be possible for me to intuit objects like cups and if I couldn't intuit the cup into which I am pouring coffee it wouldn't be possible for me to know that it is chipped by seeing that it is chipped or, for that matter, to know anything else about the cup by any other means.

However suggestive this line of thinking might be it doesn't ultimately get one very far. For a start, it's false that intuitions and concepts are the elements of *all* of our knowledge since it's false that all of our knowledge is perceptual.⁹ I can come to know that the chip is chipped by learning from someone else that it is chipped, and this doesn't require me to intuit the cup. Perhaps, in that case, what Kant is really saying is that intuitions and concepts are the elements of all of our *perceptual* knowledge. But what is missing from his discussion is any argument for this claim. By stipulating that perceptual knowledge of objects requires intuitions of objects he is assuming that a cognitively loaded form of object perception is required for knowing that the cup is chipped by seeing that it is chipped, but the validity of this assumption is precisely what is now in question.

As for the idea that all possible perceptions are subject to the categories, that depends on Kant's theory of synthesis. Specifically, the suggestion is

⁹ On the other hand, it might be true that perceptual knowledge is, for us, a basic form of knowledge. Saying that intuitions and concepts are the elements of all of our knowledge might just be a slightly misleading way of making this point, on the assumption that perceptual knowledge requires both intuitions and concepts.

that an intuition is a complex representation that contains a ‘manifold’ of disparate elements. These need to be combined or synthesized in order to form a sensible representation of an object. The combining of the manifold of intuition is the ‘*synthesis of apprehension*’ (B160) or ‘empirical synthesis’ (B164) but empirical synthesis presupposes a prior ‘transcendental synthesis’. And this is where the categories come into play. They are the rules or basic forms of transcendental synthesis and are derivable from the logical forms of judgement because the aim of transcendental synthesis is to produce perceptual representations on which judgements about how things stand in the world can be rationally based.¹⁰ That is why Kant describes the categories as means by which intuitions are ‘determined’ in respect of the functions of judgement.¹¹

Although it is difficult to swallow the theory that perceptions of objects are the products of acts of synthesis or the notion that categorial concepts are rules of synthesis McDowell argues in his Woodbridge Lectures that Kant is nevertheless on the right track in his account of perceptual experience and perceptual knowledge. On McDowell’s account it is plausible both that knowing that a particular cup is chipped by seeing that it is chipped requires a cognitively loaded form of object perception and that this form of object perception involves ‘something like the categories’ (McDowell 1998d: 465). In effect, therefore, what McDowell comes up with on the basis of his reading of Kant is a version of the argument from object perception for CTR. So what we now need to do is to figure out how this argument goes and whether it works.

McDowell focuses in his discussion on visual experience, say an experience in which it looks to one as if there is a red cube in front of one. Such experiences make or contain claims and so are ‘actualizations of conceptual capacities’ (1998d: 438).¹² When one *judges* that there is a red cube in front of one one is exercising two conceptual capacities together, the capacity to judge that something is red and the capacity to judge that something is a cube. In ostensibly *seeing* that there is a red cube in front of one the very same two conceptual capacities are

¹⁰ I’m assuming that Kant is talking about transcendental synthesis when he describes the categories as ‘original pure concepts of synthesis that the understanding contains within itself *a priori*’ (A80/B106).

¹¹ This is in line with Paul Guyer’s suggestion that ‘the categories are the concepts by means of which we organize our intuitions in order to make them accessible to judgements’ (1992: 130).

¹² McDowell takes the idea that visual experiences make claims from Sellars 1956: 271–2. The idea is that visual experiences ‘are to be understood on the model of linguistic performances in which claims are literally made’ (McDowell 1998d: 438).

involuntarily actualized with ‘the *same* mode of togetherness’ (1998d: 440). This is the sense in which visual experiences are actualizations of conceptual capacities. But an ostensible seeing that there is a red cube there will have ‘more specificity to its content’ (1998d: 459); since ‘the apparent red cube will be *placed* more determinately than just somewhere or other in front of one’ (ibid.) the ostensible seeing’s content will be that there is a red cube *there*.

Now consider an ostensible seeing, or an experience of a red cube there, that is actually a *seeing* of a red cube there. If it is a seeing then there is a red cube there. The red cube is ‘in the subject’s view as *that red cube*’ (1998d: 459). But *that red cube* is the content of an intuition. Intuitions bring objects into view, and ‘intuitional content is essentially a fragment of judgemental content’ (1998d: 463). So we have now arrived at the idea of a cognitively loaded form of object perception that is part and parcel of epistemic seeing. In order to see that there is a red cube there one needs to *intuit* the particular red cube that is in question, but intuitions are conceptual occurrences. Unlike a simple seeing, a visual intuition represents an individual as a ‘*this-such*’ (1998d: 452), as *this red cube*, for example. To visually intuit an object is to perceive and not just to see an object because intuition involves some degree of recognition or categorization.

Cube is a sortal concept and *red* is a characterizing concept. So even if one agrees that these concepts must be actualized in seeing that there is a red cube there we still don’t have an argument that this kind of seeing has anything to do with the categories. This is where McDowell’s holism comes into play. The idea is that to have concepts like *red* and *cube* one needs lots of other concepts. For example, one must have other colour and shape concepts. To have lots of concepts is to have a world view so we now have the idea that one needs a world view in order to have any concepts, even quite basic concepts like *red* and *cube*. The moral is that ‘we can make sense of objects coming into view in intuitions only because we can see how objects fit into a view of the world’ (1998d: 465). But seeing how objects fit into a view of the world depends on ‘something like the categories, and the principles Kant connects with them’ (1998d: 465–6). So, for example, to have shape concepts one must understand how the shape of an object affects its behaviour, and this means that having shape concepts goes hand in hand with a capacity for causal thinking. Since causal thinking is a form of categorial thinking we are now in a position to understand the point of CTR: the role of categorial thinking is to sustain the

world view that one needs to have in order to have objects in view in intuitions.

To sum up, we have finally arrived at a version of the argument from object perception for CTR. This is how the argument goes: to know that there is a red cube there by seeing that there is a red cube there I must visually intuit and not just see the cube. To do that I need the concepts *red* and *cube*, and to have these concepts I must be capable of categorial thinking. By the same token, to know that cup is chipped by seeing the cup itself I must perceive and not just see the cup. To do that I need the concepts *cup* and *chipped*, and to have these concepts I must be able to think categorially. Possession of a capacity for categorial thinking is therefore a necessary condition for seeing that there is a red cube there, that the cup is chipped, or for any other primary epistemic perceiving. But it's only a background necessary condition. One needn't actually be thinking categorially every time one sees that there is a red cube there or that a particular cup is chipped.

Is this argument any good? One of its advantages is that it doesn't require one to deny the existence of simple seeing. The point is not that there is no such thing as simple, concept-free seeing but that epistemic seeing requires more than simple seeing. One can simply see a red cube there even if one lacks the concepts *red* and *cube*, but if one fails to see the red cube as a red cube then one can't be said to see that there is a red cube there. It is because epistemic seeing is itself conceptual that one's awareness of the red cube, in so far as it is a component of an epistemic seeing, must itself be conceptual. Intuitional content is a fragment of judgemental content, and fragments of judgemental content can't be any less conceptual than judgemental contents themselves.

Where does this leave CTR? For McDowell, the categories make it possible for one to have a world view. But why are the categories and the principles Kant connects with them necessary for having a world view? Why wouldn't other concepts and principles do just as well? What one can extract from McDowell's discussion is a sketch of an answer rather than an answer to this question: intuitions aren't just presentations of object but presentations of object as objects, as objective phenomena. This means that intuitions must fit into a view of the world as capable of existing unperceived, and this is where the categories come in: without concepts like *substance* and *causality* one wouldn't be able to think of the world as objective. But even if this is true, why must visual intuitions of objects be conceived of as presentations of objects as objects? It is one thing to say that one can't see that there is a red cube there without the

cube being in one's view as *that red cube*; it is another matter whether intuitions of object 'as objects' are necessary for epistemic seeing.

For these reasons, the argument from object perception can't be regarded as a totally effective argument for CTR. In particular, McDowell's version of this argument fails to make it plausible that the sortal and other non-categorical concepts that one needs in order to have intuitions of objects are somehow dependent on categorical concepts and an associated capacity for categorical thinking. It's not that there is no case for tying non-categorical to categorical concepts but that the link needs to be established in some other way, ideally without relying on the assumption that one needs visual intuitions of objects as objects in order to see epistemically. A different approach is needed, and this is where it might be a good idea to turn from the argument from object perception to the argument from epistemic perception. The latter argument is explicitly concerned with the way that possession of sortal and characterizing concepts is sustained by a capacity for categorical thinking. So if the argument from epistemic perception can be made to work then it can also be used to bolster the argument from object perception.

4.3 EPISTEMIC PERCEPTION

Suppose that one isn't persuaded that it wouldn't be possible to see a chipped cup without perceiving it in a cognitively loaded sense. Still, there is no getting away from the fact that seeing that the cup is chipped is conceptual in a way that simple seeing is not. Maybe one can see the cup if one lacks the concept *cup* but one can't see *that* the cup is chipped unless one has the concept *cup* and the concept *chipped*. The argument from epistemic perception says that in order to have such concepts one must have a capacity for categorical thinking. The point is not that one must actively be employing a categorical concept every time one perceives epistemically but that being able to think categorially is a background necessary condition for epistemic perceiving. If one lacked this capacity one would also lack concepts like *cup* and *chipped*.

What have sortal and characterizing concepts got to do with categorical concepts or with categorical thinking? Why couldn't one have concepts like *cup*, *chipped*, *red*, and *cube*, and therefore be in a position to perceive epistemically, regardless of whether one has concepts like *object*, *substance*, or *cause*? Let's call the thesis that one must be able to

think categorially in order to have any sortal or characterizing concepts thesis C. What we now need is a good argument for C. If we can find one that would obviously be good news for the argument from epistemic perception, but it would also be good news for the argument from object perception; it would give substance to the suggestion that the non-categorial concepts that are actualized in visual intuitions of objects depend on something like the categories.

There are lots of different ways of arguing for C. One would be to argue for it on the basis that categorial thinking is necessary for the *acquisition* of non-categorial concepts. Another would be to argue for C on the basis of an account of the nature of concepts like *cup* and *chipped*. Let's call the first of these arguments the Acquisition Argument for C and the second the Conceptual Argument. The Acquisition Argument says that in order to have any non-categorial concepts one must have acquired them, and that categorial thinking is necessary for possession of non-categorial concepts because it is necessary for acquiring them. In contrast, the Conceptual Argument leaves it open that concepts can be possessed without having been acquired, at least if concept acquisition is understood in the way that the Acquisition Argument understands it.¹³ It argues that one must be able to think categorially in order to have any sortal or characterizing concepts regardless of whether or how one acquired them. Either way, if C is defensible then the ability to think categorially comes out as a necessary condition not just for seeing that the cup is chipped but for any epistemic perceiving.

The idea that categorial thinking is necessary for the acquisition of concepts like *cup* and *chipped* is one that Béatrice Longuenesse finds in Kant.¹⁴ Like other sortal and characterizing concepts, *cup* and *chipped* are empirical concepts. An empirical concept in Kant's sense is one 'which can only be derived from experience' (B3). So when the Acquisition Argument talks about what is necessary for the acquisition of the concepts *cup* and *chipped* it's really talking about what is necessary for the derivation of such concepts from experience. The question is whether one needs to be able to think categorially in order to get non-categorial concepts from experience. Longuenesse thinks that Kant's answer to this question is 'yes'. In effect, therefore, she reads Kant as

¹³ In Jonathan Bennett's terminology, concepts that are possessed without having been acquired would be innate. See Bennett 1966: 98. Unlike the Conceptual Argument, the Acquisition Argument shouldn't move anyone who thinks that non-categorial concepts are, or could be, innate in Bennett's sense.

¹⁴ See Longuenesse 1998.

putting forward a theory of concept acquisition which commits him to something along the lines of the Acquisition Argument for C.

What is it to 'derive' a concept from experience? Kant's theory of concept acquisition is organized around the notions of *comparison*, *reflection*, and *abstraction*. These notions are explained in the following passage from one of his logic lectures:

To make concepts out of representations, one must... be able *to compare*, *to reflect*, and *to abstract*, for these logical operations of the understanding are the essential and universal conditions for the generation of every concept whatsoever. I see, e.g., a spruce, a willow, and a linden. By first comparing these objects with one another I note that they are different from one another in regard to the branch, the leaves, etc.; but next I reflect on what they have in common among themselves, trunk, branches, and leaves themselves, and I abstract from the quantity, the figure, etc. of these; thus I acquire a concept of a tree. (1992: 592)

To see that something is a tree I would already need to have the concept *tree* so this can't be the kind of 'experience' from which this concept is derived. But I don't already need to have the concept *tree* simply in order to see a spruce, a willow, or a linden. It's no accident, therefore, that this is the kind of seeing from which Kant thinks that the concept *tree* is derived. To compare, reflect, and abstract is, in Kant's terminology, to 'analyse' one's representations, so we now have the proposal that analysis is the means by which 'the understanding elevates given representations to a discursive form' (Longuenesse 1998: 11). If this proposal is correct, then one way of showing that categorial thinking is necessary for the acquisition of sortal and characterizing concepts would be to show that this kind of thinking is necessary for analysis, for the derivation of concepts like *cup*, *tree*, and *chipped* from experience by comparison, reflection, and abstraction.

Longuenesse maintains that what links analysis to the categories is synthesis. In outline, her idea is that synthesis is necessary for analysis and that the categories are necessary for synthesis. The point is this: for analysis to be possible, for it to be possible for the understanding to analyse given representations to form concepts, representations must be susceptible to analysis. However, it can't just be a brute, inexplicable fact that representations are 'susceptible to being reflected under concepts' (1998: 12). There must be something that accounts for this susceptibility, some activity of the understanding or imagination that structures what is given to the senses so that concepts can be extracted from it by comparison, reflection, and abstraction. This something,

which ‘*must occur prior to analysis*’ (1998: 11) and whose role is ‘*to make analysis possible*’ (1998: 12), is ‘pure synthesis’. But the categories are conceptual representations of the different forms of pure synthesis that make it possible to acquire concepts by comparison, reflection, and abstraction.¹⁵ So if synthesis is necessary for analysis then so are the categories. Without the categories there would be no synthesis, without synthesis there would be no analysis, without analysis there would be no concepts like *tree* and *cup*, and without concepts like *tree* and *cup* there would be no epistemic perceiving.

There are lots of problems with this line of thinking. One is that it threatens a regress. For if pure synthesis is needed to prepare the ground for analysis, the obvious question is: what prepares the ground for pure synthesis? What ensures that the sensible given is susceptible to synthesis in the way that synthesis ensures that the sensible given is susceptible to analysis? If we answer this question by positing some form of proto-synthesis that makes synthesis possible then we will have to account for the susceptibility of the sensible given to proto-synthesis. If we don’t want to go down this path we will have to say that it’s just a brute fact that representations are capable of being synthesized so there is no need to talk about proto-synthesis. But if it can be a brute fact that representations are capable of being synthesized, why can’t it be just as brute fact that they are susceptible to analysis? Why think that we need to *do* anything to what is given to the senses to make it possible to extract concepts from it?

These questions suggest that the theory of synthesis isn’t a response to a genuine problem; the susceptibility of the sensible given to analysis is not something that needs to be or can be explained. That’s a good thing because if there were a genuine problem here the theory of synthesis wouldn’t be a solution to it. It’s not just that this theory generates a regress but also that it is very hard to make sense of the notion of a type of mental activity that is capable of transforming something that isn’t already analysis-friendly into something that is analysis-friendly. ‘Pure synthesis’ looks like a label for something of which we have no coherent conception because we have no coherent conception of what

¹⁵ On Longuenesse’s reading the categories are concepts that universally represent the different forms of the synthesis that must occur prior to analysis. This is what she thinks that Kant is getting at when he says, in Kemp Smith’s translation, that ‘pure synthesis, represented in its most general aspect, gives us the pure concept of the understanding’ (A78/B104). Longuenesse’s translation is: ‘Pure synthesis, *universally represented*, gives us the pure concept of the understanding.’ See her 1998: 11 for further discussion.

it would be to *make* comparison, reflection, and abstraction possible. If this is right, then we should refrain from characterizing the categories as representations of the different forms of pure synthesis or from relying on Kant's theory of synthesis in arguing for C; if categorial thinking is necessary for concept acquisition it had better not be because pure synthesis is necessary for concept acquisition.

Does this mean that the Acquisition Argument for C is no good? Not necessarily. Instead of arguing that synthesis is what links analysis to the categories we could jettison the notion of synthesis and try to establish a more direct connection between analysis and categorial thinking. The thought would be that categorial thinking is necessary for analysis because to compare, reflect, and abstract is, in effect, to be thinking categorially. So the position is not that thinking by means of categorial concepts prepares the ground for analysis but that analysis itself involves thinking by means of categorial concepts. This makes it hard to see how comparison, reflection, and abstraction could be the source of the categories but that's as it should be. Although Kant has a broadly abstractionist conception of how empirical concepts are derived from experience he doesn't think that all concepts are empirical.¹⁶ Some concepts can't be abstracted from experience and so are a priori rather than empirical. On this account, the categories qualify as a priori concepts because one must already have them in order to abstract a concept from experience.¹⁷

To get a feel for the role of categorial thinking in comparison, reflection, and abstraction consider Kant's story about the acquisition of the concept *tree*. The first thing that happens is that I notice that several individual trees have certain features in common. This means that I must be able to think of each tree as a distinct individual, that I must be able to think of the several trees as a plurality of distinct individuals, and that I must be able to think of certain features as ones which all the trees have in common. These are all forms of categorial thinking. We can see this by noting that *unity*, *plurality*, and *totality* are all categories in Kant's sense. Specifically, they are categories of quantity. It's natural to suppose that to have a categorial concept in one's repertoire is to have a certain cognitive ability. So, for example, to have the concept of unity is to

¹⁶ McDowell seems to want to deny that Kant has an 'abstractionist picture of the formation of basic empirical concepts' (1998d: 462). The passage from the Jäsche Logic, quoted above, suggests otherwise.

¹⁷ See Bennett 1966: 98 for a closely related proposal.

have the ability to think about individuals as such, to have the concept of plurality is to have the ability to think about distinct individuals as constituting a plurality, and to have the concept of totality is to have the ability to think about totalities as such or about what is true of a totality of distinct individuals. But these are precisely the cognitive abilities one must be exercising in order to acquire the concept *tree* by comparison, reflection, and abstraction. So if I can think of each tree as a distinct individual, of several trees as distinct individuals, and of what all of a group of trees have in common, it follows that I have the categories of unity, plurality, and totality.

As for the remaining categories, *causality* is one of the relational categories, along with substance and community. Kant describes the concept of causality as the concept of the relation of ‘ground to consequence’ (B112).¹⁸ To have this concept is to be able to think hypothetically, and it’s plausible that the formation of the concept *tree* involves this kind of thinking. As Longuenesse remarks, we generate the concept *tree* by ‘learning to attribute to it various characters dependent on *added* conditions: “If the weather gets cold, trees lose their leaves”, “If a tree gets no water, it perishes,” etc.’ (1998: 145). Here, getting no water is the ground and perishing the consequence, and thinking that if a tree gets no water it perishes is a form of causal thinking. Without the ability to engage in this kind of thinking one wouldn’t have a proper understanding of what trees are and so wouldn’t count as having got hold of the concept *tree*. Since causal thinking is a form of categorial thinking, this is an illustration of the point that categorial thinking is necessary for the acquisition of concepts like *tree* by comparison, reflection, and abstraction.

Further illustrations of the same point are not difficult to find when one considers the three categories of modality, possibility, existence, and necessity. To get hold of the concept *tree*, one has got to grasp the distinction between changes that it’s possible for a tree to undergo without ceasing to exist and changes that would amount to the tree’s ceasing to exist. To grasp this distinction is to be able to think modally, to be able to think in terms of existence, possibility, and necessity, and the indispensability of modal thinking for concept acquisition amounts to indispensability of the three modal categories for concept acquisition. The relational category of substance, the concept of ‘something which can exist as subject and never as mere predicate’ (B149), is also relevant here. One’s conception of the changes a tree can undergo without

¹⁸ This is the ‘pure’ rather than the ‘schematized’ concept of cause.

ceasing to exist must be the conception of changes in the state of a persisting subject of change, that is, the conception of changes in the state of a persisting *substance*.

We now have a synthesis-free argument for the thesis that the acquisition of sortal concepts requires categorial thinking. What about characterizing concepts? Assuming that concepts like *red* or *chipped* are acquired by comparison, reflection, and abstraction, categorial thinking must also be involved in their acquisition. Selective attention to features that different individuals have in common is as much a part of the abstractionist account of the acquisition of characterizing concepts as it is of the abstractionist account of the acquisition of sortal concepts.¹⁹ So unity, plurality, and totality are involved in both cases. The same goes for causality and possibility. To get hold of the concept *blue* is, among other things, to grasp a range of conditional and modal propositions about the colour blue, for example propositions about what happens if blue is mixed with brown or about the relationship between being blue and looking blue. These are examples of causal and modal thinking, and therefore also examples of categorial thinking.

To sum up, the challenge was to find an argument for C, for the thesis that one must be able to think categorially in order to have any sortal or characterizing concepts, and we have now outlined one response to this challenge. The Acquisition Argument says that one must be able to think categorially in order to have concepts like *cup* and *chipped* because this kind of thinking is required for the acquisition of non-categorial concepts from experience. We have seen that this argument can be, but needn't be, fleshed out by drawing on Kant's theory of synthesis. If the synthesis-free version of the Acquisition Argument is in good shape then so is C, and if C is in good shape then so is the argument from epistemic perception for CTR. To perceive epistemically one must have sortal and characterizing concepts and the Acquisition Argument says that one couldn't acquire, and so couldn't have, sortal or characterizing concepts if one lacked the capacity to think categorially.

¹⁹ Mackie talks about selective attention in his account of Locke's theory of abstraction. Lockean abstraction, as Mackie understands it, 'consists in paying selective attention to one feature in a complex particular object of experience and ignoring other features which are in fact occurring along with it, and in associating verbal expressions (or other signs) with the selected feature in such a way that one is ready to apply them to other objects that are like this one with respect to this one feature' (1976: 112). I think that it is helpful to read what Kant calls comparison, reflection, and abstraction as involving selective attention in something like this sense.

In fact, this isn't quite right. Strictly speaking, the most that the Acquisition Argument shows is that categorial thinking is necessary for the acquisition of the sortal and characterizing concepts *on an abstractionist conception of concept acquisition*. What it doesn't show is that sortal and characterizing concepts must be acquired or that they can *only* be acquired by comparison, reflection, and abstraction. It's true that Kant describes these three operations as the essential and universal conditions for the generation of every concept whatsoever but this is not a claim that the Acquisition Argument establishes or even argues for. If there are other ways of acquiring non-categorial concepts then there is always the possibility that they don't require categorial thinking in the way that the acquisition of concepts by comparison, reflection, and abstraction requires categorial thinking. Unless we can exclude this possibility we are not yet entitled to conclude that categorial thinking is a necessary condition for concept acquisition let alone for concept possession.

To make this worry about the Acquisition Argument more concrete consider the following example: yesterday I didn't know the meaning of the word 'cup' and was unable to recognize cups as cups. This morning my cortex was surgically rewired. Now I can use the word 'cup' correctly and am able to distinguish cups from non-cups. It looks as though I have acquired the concept *cup* as a result of having my cortex surgically rewired.²⁰ Yet acquiring a concept in this way is not acquiring it by comparison, reflection, or abstraction, or in a way that is anything to do with categorial thinking. So the Acquisition Argument fails to show that comparison, reflection, and abstraction are the essential and universal conditions for the acquisition of non-categorial concepts. By the same token, it fails to show that categorial thinking is a necessary condition for the acquisition of non-categorial concepts. The most that it shows is that it is a necessary condition for the acquisition of non-categorial concepts by one particular method or set of operations.

²⁰ As Fodor remarks, 'there are, in principle, many other ways than learning in which the repertoire of mental representations available to an organism might be affected by its experiences. Think of being hit on your head or having your cortex surgically rewired' (1981: 275). To abstract a concept from experience is, in Fodor's terms, to 'learn' it. Concept learning is a rational process but concepts can also be acquired by non-rational mechanisms. Cortical rewiring would be one such mechanism. Such examples make the point that 'learning how to do something is not the only way of coming to be able to do it' (Bennett 1966: 97).

We don't usually acquire concepts by cortical rewiring but that's not the point. The point is that if it is *possible* to acquire a concept in this way then categorial thinking can't be *necessary* for concept acquisition; having brain surgery isn't a way of acquiring a concept that requires categorial thinking. Kant thinks of concept acquisition as a rational process, and he regards categorial thinking as necessary for concept acquisition because he thinks that the rational acquisition of concepts from experience requires this kind of thinking. But concept acquisition needn't be a rational process, and that's the problem with the Acquisition Argument; even if the relation between a concept and the experiences which occasion its acquisition is normally rational it can be 'brute-causal' (Fodor 1981: 280).

Although this seems a compelling objection to the Acquisition Argument, C is still in the running because the Conceptual Argument is still in play. This argument shifts the focus from what is required to acquire a sortal or characterizing concept to the nature of such concepts themselves. The thought is that concepts like *cup*, *tree*, or *chipped* are such that one must be able to think categorially in order to possess them, and that this is so regardless of how one acquired them. Even if concepts can be possessed without having been acquired, there are still substantive constraints on what it is to have a sortal or characterizing concept. The Conceptual Argument identifies C as one such constraint, and it does so on the basis of considerations that are similar in some respects to ones that the Acquisition Argument exploits.

The first consideration is this: concepts are, first and foremost, constituents of thoughts and 'thoughts are *essentially* structured' (Evans 1982: 104). For example:

There simply could not be a person who could entertain the thought that John is happy and the thought that Harry is friendly, but who could not entertain—was conceptually debarred from entertaining—the thought that John is friendly or the thought that Harry is happy. Someone who thinks that John is happy must, we might say, have the idea of a *happy man*—a situation instantiated in the case of John (he thinks), but in no way tied to John for its instantiation. (Evans 1982: 103)

To say that the concept of a happy man is not tied to a particular man for its instantiation is to say that this concept can, in principle, be instantiated by a range of distinguishable individuals.²¹ What is

²¹ As Strawson remarks, 'the main point here is a purely logical one: the idea of a predicate is correlative with that of a *range* of distinguishable individuals of which the predicate can be significantly, though not necessarily truly, affirmed' (1959: 99 n. 1).

more, this is something that someone who has the concept *happy man* must grasp, at least implicitly. But someone who understands that a concept like *happy man* can be instantiated by a range of distinguishable individuals is someone who can think in categorial terms: to have the concept of an individual is to have the concept of unity and to have the idea of a range of individuals is to have the concept of plurality.

This argument for C trades on what Evans calls the 'Generality Constraint'. This says that the thought that *a* is *F* must be seen as 'lying at the intersection of two series of thoughts: on the one hand, the series of thoughts that *a* is *F*, that *b* is *F*, that *c* is *F*, ... , and, on the other hand, the series of thoughts that *a* is *G*, that it is *H*, ...' (Evans 1982: 104). Kant is making essentially the same point when he describes generality as the form of all concepts. The concept of a happy man is not tied to a particular man for its instantiation because concepts, as distinct from intuitions, are representations of what is, or can be, 'common to several objects' (Kant 1992: 589). It is because of this feature of concepts that one must be able to think categorially in order to have any sortal or characterizing concepts.

Other arguments can be used to establish a connection between non-categorial concepts and the category of causality. For example, sortal concepts are concepts of objects and concepts of objects are concepts of things that are capable of changing. Yet the changes that objects undergo do not, on the whole, make it impossible for us to re-identify them. We are able to re-identify objects because, for the most part, they change in predictable and regular ways and because our sortal concepts are 'linked with sets of conditional expectations about the things we perceive as falling under them' (Strawson 1966: 145). Since the regularities that sustain these expectations are causal regularities concepts of objects are 'always and necessarily compendia of causal law or law-likeness, carry implications of causal power or dependence' (Strawson 1966: 145–6). But if sortal concepts must be compendia of causal law then it looks as though one must be capable of causal thinking if one is to have any sortal concepts.

To describe sortal concepts as compendia of causal law is to make a point about the nature of concepts like *tree* and *cup*. If it is in the nature of sortal concepts to be compendia of causal law then it is only possible to acquire sortal concepts by having one's cortex surgically rewired if one already has a capacity for causal thinking or if this capacity can itself be acquired by cortical rewiring. Either way, what is doing the work in the argument for C is reflection on what it is for

a concept to be a sortal concept rather than on what it is to acquire one. By focusing on the nature of sortal and characterizing concepts the Conceptual Argument is able to sidestep the worry that such concepts might be acquired without having been abstracted from experience. Even innate concepts are concepts and there is no such thing as a non-categorical concept that one can have without being able to think categorially.

The obvious worry about the Conceptual Argument is that it only justifies CTR on an excessively deflationary reading of this requirement. For example, we can all agree that one hasn't got hold of a concept like *cup* if one lacks the notion of *a* cup but is this really all that the indispensability of the 'category' of unity comes to? If so, then it begins to look as though there is much less to CTR than meets the eye, and that the Conceptual Argument for this requirement is nothing more than a laborious exercise in stating the obvious. The contrast with Kant's own conception of CTR couldn't be sharper. The one thing that one couldn't possibly accuse Kant of when he says that the categories underpin the synthesis that makes sortal concepts available in the first place is stating the obvious. The Conceptual Argument responds to the problems with Kant's account by sanitizing CTR but it runs the risk of depriving this requirement of much, if not all, of its bite; if we rely on this argument then CTR will end up saying what hardly needs saying.

One thing that proponents of the Conceptual Argument can do to deal with this criticism is to point out that it's far from trivial that grasp of sortal concepts requires a capacity for causal thinking. So it's certainly not the case that the Conceptual Argument only justifies a version of CTR that is so deflationary in relation to the category of causality as to be uninteresting. As for the remaining categories, it may well be true that the suggestion that non-categorical concepts rely on the concepts of unity and plurality doesn't add much to the suggestion that the thoughts in which concepts like *cup* and *chipped* figure must conform to the Generality Constraint. But the latter suggestion is non-trivial and it has the additional merit of being plausible. It's no bad thing, therefore, if the Generality Constraint rather than, say, the theory of synthesis turns out to be the basis of CTR, at least when it comes to explaining why one would need categories of quantity in order to see that the cup is chipped.

Another worry about the Conceptual Argument is that it doesn't establish the indispensability of all twelve Kantian categories. In fact,

there was never much prospect of any argument being able to do that.²² What matters for the purposes of the Conceptual Argument is that one must have at least *some* categorial concepts—unity, plurality, and causality being the three most obvious candidates—in order to have *any* sortal or characterizing concepts, and that this is a reflection of what it is for a concept to be a sortal or characterizing concept. The categorial thinking that is required for possession of non-categorial concepts is only a background condition for epistemic perceiving because one needn't be thinking in explicitly categorial terms every time one perceives epistemically. But someone who doesn't have any categorial concepts is someone who doesn't have any non-categorial concepts, and someone who doesn't have any non-categorial concepts is someone who can't perceive epistemically. That is the point of the argument from epistemic perception for CTR.

The thesis that one needs some categorial concepts in order to have any non-categorial concepts also helps the argument from object perception. This argument turns on the idea that intuitions or perceptions of objects are cognitively loaded. What this means is that one must have a repertoire of sortal and characterizing concepts in order to intuit or perceive objects, but how is it supposed to follow that this form of awareness of objects depends on the categories? We now have an answer to this question: categorial concepts are needed not because intuitions are presentations of objects 'as objects' but because possession of non-categorial concepts can't be separated from possession of categorial concepts. If non-categorial concepts are necessary for object perception then so are categorial concepts.

This concludes the defence of CTR. If the arguments from object perception and from epistemic perception are successful they provide a positive answer to the question of necessity. They show that categorial thinking is a genuinely necessary background condition for epistemic perceiving. In fact, the two arguments for CTR aren't all that different from each other. To describe object perception as cognitively loaded, as involving an element of recognition or categorization, is effectively to represent object perception as a form of epistemic perception.²³ It's not surprising, therefore, that the categorial constraints on epistemic

²² Which is not to say that the Conceptual Argument couldn't easily be extended to cover the modal categories.

²³ This is clear from McDowell's discussion. At one point he describes visual intuitions of objects as 'seeings that ..., looked at as it were from a different angle' (1998d: 462).

perception are also constraints on object perception. The best arguments for the existence of such constraints are the arguments for C and the best argument for C is the Conceptual Argument.

4.4 FOUNDATIONS, A PRIORITY, AND MINIMALISM

The remaining questions about CTR can be dealt with more briefly. The first, the question of foundations, concerns the basis of CTR: what, if anything, accounts for the fact that categorial thinking is a background necessary condition for the acquisition of knowledge of the external world by means of epistemic perception? Is it the cognitive structure of the mind, as idealists claim, or the nature of the objects of knowledge, as realists claim, that accounts for the obtaining of CTR? The initial answer to this question appears to be ‘neither’. Part of what explains the truth of CTR is the fact that epistemic perceiving is conceptual. Yet neither the structure of the mind nor the nature of the things we perceive can plausibly be regarded as making it the case that epistemic perceiving is conceptual. What makes this the case is the nature of epistemic perception. But the involvement of sortal and characterizing concepts in epistemic perception only supports CTR on the assumption that such concepts are underpinned by a capacity for categorial thinking. According to the Conceptual Argument for C, what justifies this assumption is the nature of sortal and characterizing concepts. So unless we are prepared to attribute the nature of sortal and characterizing concepts to the workings of the mind or nature of the things to which they are applied neither idealism nor realism constitutes a viable response to the question of foundations.

What would be wrong with attributing the nature of concepts like *cup* and *chipped* to the way the mind is? After all, if concepts are constituents of thoughts and thoughts are ‘in the mind’ then concepts are mental. And if concepts are mental isn’t their nature bound to be fixed by the nature of mind? This looks like a non sequitur; one can think that concepts are mental without thinking that what makes a particular concept a sortal or characterizing concept is fixed by the structure of the mind. For example, it’s not clear what it would even mean to say that it is the cognitive structure of the mind that accounts for the fact that sortal concepts are compendia of causal law. If anything, it’s easier to attribute this feature of sortal concepts to the nature of the things that

fall under them than to the nature of thinkers. In any case, one might want to resist the idea that thoughts are mental. If, like Frege, one thinks that thoughts are ‘neither things of the outer world nor ideas’ (1967: 29) then one should be reluctant to attribute the nature of concepts either to the nature of mind or to the nature of the outer world.

Where does this leave the question of a priority? Can CTR be established without any empirical investigation? The easiest way of showing that something can be established without any empirical investigation is to establish it without any empirical investigation. That’s exactly what the argument from epistemic perception does in relation to CTR. It establishes this requirement by reflecting on the role of sortal and characterizing concepts in epistemic perception and on the links between sortal and characterizing concepts and categorial thinking. The method used by CTR to establish these links is philosophical reflection rather than empirical science. Indeed, it’s hard to see how something like the claim that a capacity for categorial thinking is required for the possession of non-categorial concepts could be established empirically. This suggests that CTR is not just weakly but also strongly a priori; it can only be established non-empirically.

Where does all of this leave the dispute between minimalism and anti-minimalism? CTR is a Level 3 response to (HP_{pk}). It attempts to explain how perceptual knowledge is possible by identifying one of its enabling conditions. Minimalism insists that the only enabling conditions for knowing that something is the case by seeing that it is the case are empirical conditions, and that there are no strongly or weakly a priori enabling conditions for the acquisition of perceptual knowledge. Since we have just seen that categorial thinking is a strongly a priori condition for coming to know things about the external world by means of epistemic perception minimalism must be wrong about this. The minimalist wants to represent all Level 3 questions as scientific questions and thereby to ‘naturalize’ the project of explaining what makes perceptual knowledge possible. The argument from epistemic perception for CTR shows that there are plausible Level 3 questions about perceptual knowledge that can only be answered by a priori philosophical reflection and that therefore can’t be naturalized.

At the same time, the extreme anti-minimalist is wrong to insist that a complete and philosophically satisfying response to (HP_{pk}) *must* identify a priori enabling conditions for the acquisition of perceptual knowledge. Seeing that the cup is chipped is a way of coming to know that it is chipped. Having pointed this out, and shown that there is

nothing that stands in the way of our seeing that the cup is chipped, we can go on to ask what makes it possible for one to have concepts like *cup* and *chipped*. Although this is a legitimate and interesting question, it's hardly one without an answer to which we can't properly claim to have explained how it is possible to know that the cup is chipped. In the context of (HP_{pk}) we can but needn't ask what makes it possible to see that the cup is chipped, just as we can but needn't ask what makes it possible to use a step ladder to catch a fly ball twenty feet off the ground. The difference between the two questions isn't that one of them is optional in a way that the other isn't; they are both optional. The difference is that the question about the enabling conditions of epistemic seeing is a philosophical question rather than a question for physics.

5

Other Minds

5.1 THE PERCEPTUAL MODEL

So far in this book my focus has been on perceptual knowledge. I've interpreted (HP_{pk}) , 'how is perceptual knowledge possible?', as an obstacle-dependent how-possible question and recommended a multi-levels response which operates at three levels: the level of means (Level 1), of obstacle-removal (Level 2), and of enabling conditions (Level 3). In this chapter I want to apply the multi-levels approach to another epistemological how-possible question, namely:

(HP_{om}) How is knowledge of other minds possible?

I'm going to argue that this is another obstacle-dependent question, and that a multi-levels response is as helpful in this context as it is in the context of (HP_{pk}) . To make this plausible I will start by identifying some of the obstacles to knowledge of other minds which might prompt one to ask (HP_{om}) . Then I will outline a Means Response to this question and suggest that the obstacles to knowing about other minds by the proposed means can be overcome or dissipated. Lastly, I will identify some enabling conditions for knowing about other minds by the proposed means.

Let's start by taking a closer look at (HP_{om}) . Why does this question arise? The first thing to note is that there is more than one way of understanding the kind of knowledge that is at issue here. One problem is that it's sometimes difficult to tell *what* other people are thinking or feeling; for example, I might believe that a colleague has views about some administrative question but his opaque manner of expression makes it difficult if not impossible for me to figure out what his views are. Or I might suspect that someone is angry without being able to tell whether he is angry. Again, the problem is that other people can be quite opaque. Sometimes we think that with a bit more effort we could find out what someone else is thinking or feeling but on other occasions

or with some types of people we are much less confident of that. As Austin remarks:

[W]e do not for a moment suppose that we *always* know, of *all* men, whether they are angry or not, or that we could discover it. There are many occasions when I realize that I can't possibly tell what he is feeling; and there are many *types* of people, and many individuals too, with whom I (they being what they are, and I being what I am) never tell. The feelings of royalty, for example, or fakirs or bushmen or Wykehamists or simple eccentrics—these may be very hard to divine: unless you have had a prolonged acquaintance with such persons, and some intimacy with them, you are not in any position to know what their feelings are, especially if, for one reason or another, they can't or don't tell you. (1979: 103–4)

It obviously doesn't follow from the fact that it's sometimes hard to divine what others are thinking or feeling that this is always hard to divine. But even if we suppose that it's sometimes possible to tell what others are thinking or feeling we might still wonder *how* this is possible. Even when it comes to people with whom we are well acquainted we can be mistaken or misled about their thoughts or feelings, and that is why we might find ourselves asking how knowledge of other minds is *ever* possible. So one way of reading (HP_{om}) would be to read it as asking:

(HP_{omw}) How is it possible to know *what* another person is thinking or feeling?

I will call this the 'what' version of the epistemological problem of other minds, hence the 'w'.

This version of (HP_{om}) might be thought to presuppose an answer to a much more basic question about the very existence of other minds. For if the possibility of knowing *what* other people are thinking or feeling needs explaining then so does the possibility of knowing *that* they are thinking or feeling anything at all. The fundamental question, it seems, is not what others are thinking or feeling but whether they think or feel *anything at all*. So now we have a 'that' version of the problem of other minds:

(HP_{omt}) How is it possible to know *that* others think or feel anything?¹

¹ Anita Avramides also distinguishes between the 'what' and the 'that' versions of what she calls the 'epistemological' problem of other minds. See Avramides 2001: 219.

To get a feel for the problem, let's begin by listing some basic sources of empirical knowledge. Such a list would include introspection, perception, testimony, and inference. Which, if any, of these sources can be a source of knowledge of the existence of other minds? Not introspection since it can only tell me that I think and feel; it can't tell me that others do too.² Perception is no use either, at least on the assumption that the thoughts and feelings of others are unobservable. As for testimony, I'm only going to take someone else's word for it that I am not alone in the world if I take it that his words are expressive of genuine thoughts, that is, if I already take myself to know of the existence of at least one other thinker. Since the possibility of knowing this is precisely what is at issue we still lack an answer to (HP_{omt}). That leaves inference. Yet it seems that neither inductive inference nor inference to the best explanation can put me in a position to *know* that others think or feel anything; the most I can conclude on the basis of either form of inference is that it's *probable* that there are other subjects of thought and experience. So we can now see why (HP_{omt}) is a genuine question if it's knowledge that we are after. The problem is that none of the four presupposed sources of empirical knowledge can provide us with a satisfactory answer to the 'that' version of the problem of other minds.

To read (HP_{omt}) in this way is to read it as an expression of what I described in Chapter 1 as the *problem of sources*. In relation to (HP_{omt}) the problem of sources is the problem of explaining how one could come to know of the existence of other minds, given that knowledge of their existence can't easily be accounted for on the basis of introspection, perception, testimony, or inference. A *presupposed sources* solution to the problem of sources will therefore be one which shows that, contrary to what I have so far been assuming, one or more of these presupposed sources of knowledge *can* provide us with knowledge of other minds. In contrast, an *additional sources* solution will agree that our knowledge of other minds can't come from any of the presupposed sources but argue that we have other ways of coming to know of their existence. Finally, scepticism about other minds can be understood as the view that the problem of sources has no solution. The sceptic thinks that none of the proposed sources is viable and that if our acquisition of a particular kind of knowledge can't be accounted for then that is a good reason for thinking that we don't actually have any such knowledge.

² See McGinn 2004 and Cassam 2004 on the question of whether one could have introspective access to another person's thoughts or feelings.

What would an additional sources solution look like? One possibility would be to agree that none of the presupposed sources of *empirical* knowledge can provide us with knowledge of the existence of other minds but to argue that this doesn't matter if our knowledge of their existence is *a priori* rather than empirical. On this account, what we should be looking for are additional sources of knowledge of other minds that are sources of non-empirical knowledge. If, on the other hand, we insist on viewing our knowledge of the existence of other minds as empirical then an additional sources solution is going to have to find additional sources of empirical knowledge. It's important, however, that the additional sources one comes up with are additional sources that are actually available to us. So, for example, it's no good claiming that it would be possible in principle to acquire empirical knowledge of the existence of other minds on the basis of telepathy. Since we aren't telepathic the appeal to telepathy gets us nowhere; telepathy isn't a source of knowledge for us.

I'm going to defend a presupposed rather than an additional sources solution to the problem of other minds because I think it's false that none of what I have been describing as our presupposed basic sources of empirical knowledge can provide us with knowledge of other minds. If I'm right about this there is no need to look for additional sources of knowledge. So which of the presupposed sources of knowledge can provide us with knowledge of the existence of other subjects and knowledge of what they are thinking and feeling? Defenders of what I'm going to call the *inferential model* think that inference can be a source of knowledge of other minds and that our knowledge of other minds must be a form of inferential knowledge.³ In contrast, defenders of what I'm going to call the *perceptual model* look to perception as a source of our knowledge of other minds. They think that it's possible to know *that* others think and feel by *perceiving* that others think and feel, and that it's sometimes possible to know *what* others think and feel on the same basis.

The model that I want to defend here is the perceptual model. This model emphasizes that there is such a thing as, say, seeing that someone else is angry and thereby knowing that he is angry.⁴ Seeing that someone

³ This is Mill's view. See Avramides 2001: 164–71 for an account and discussion. A more recent defender of the inferential model is Christopher S. Hill. See chapter 9 of Hill 1991.

⁴ This model has also been defended by, among others, Strawson, Dretske, McDowell, and Stroud. See Strawson 1959: 108–9, Dretske 1969: 189–90, McDowell 1998c, and Stroud 2004: 166–7.

is angry is a form of what I have been calling 'epistemic' perceiving, so the suggestion is that epistemic perception can sometimes provide us with knowledge of the thoughts and feelings of others. It can provide us with knowledge of *what* someone else is feeling or thinking, and therefore also puts us in a position to know *that* others think or feel. On this account, we shouldn't think of the 'that' version of (HP_{om}) as more basic or fundamental from the standpoint of epistemology than the 'what' version. The position is not that we first have to come up with a justification for thinking that there are other minds and then figure out what others are thinking and feeling. Rather, we can come to know *that* there are other minds by seeing *what* others are thinking and feeling.

The perceptual model has a lot going for it. It provides a simple answer to the two versions of (HP_{om}) and it takes what we ordinarily say about our knowledge of other minds at face value; we do, after all, often describe ourselves as seeing what someone else is thinking or feeling. Another attraction is that it avoids some of the problems with the inferential model. One question that has often been raised about this model is whether it's actually possible to establish the existence of other minds by relying on induction or inference to the best explanation. If not, and if we don't want to end up as sceptics, we had better accept that we can know that someone else is angry by seeing that he is angry, or that someone else is suffering by seeing that he is suffering, and so on. What is more, even if one thinks that knowledge of the existence of other minds *could be* inferential it doesn't follow that our knowledge of their existence *is* primarily inferential. Epistemic perception could still be a basic source of our knowledge of the existence of other minds even if it's not the only possible source.

This point is worth making because I don't propose to argue for the perceptual model on the basis that inference couldn't be a source of knowledge of the existence of other minds. To justify such a claim one would need to look in detail at the ins and outs of the inferential model, and I'm not proposing to do that here. Instead, I'm going to be defending the perceptual model on the basis that perception is *a* source of our knowledge of other minds, and that the perceptual model gives a plausible account of how *some* of our knowledge of other minds comes to be. The perceptual model clearly shouldn't say that none of our knowledge of other minds is inferential since it's obvious that we sometimes have to rely on inference to figure out what others are thinking or feeling. But it's equally obvious from the standpoint of the

perceptual model that we also sometimes see what others are thinking and feeling. This is what the perceptual model emphasizes, partly in the hope of demystifying our knowledge of other minds.

All of this assumes that there is a genuine contrast between inferential and perceptual knowledge but that seems a reasonable assumption. Although it's sometimes suggested that all perception involves inference it's a mistake to think that perceptual knowledge is a disguised form of inferential knowledge. To infer is to reason one's way from premisses to a conclusion, and it's certainly not true that every case in which one sees that *b* is *P* is a case in which one has explicitly or implicitly inferred that *b* is *P*. For example, when I see that the cup into which I'm pouring coffee is chipped I don't infer that it is chipped. Maybe it's true that in order to see that the cup is chipped I must have various background beliefs about the cup and other things but these background beliefs needn't figure as premisses in an inference to the conclusion that the cup is chipped; I just see that it is chipped.⁵ According to the perceptual model seeing that someone else is angry is in this respect like seeing that the cup is chipped; in neither case is the resulting knowledge inferential.

But can this really be right? Despite the obvious attractions of the perceptual model, it faces some equally obvious objections. Here is one line of attack: as Mill points out, 'we may fancy that we see and feel what we in reality infer' (1891: 4), and this might lead to the worry that the perceptual model mistakes what can only be inferential knowledge for perceptual knowledge. Specifically, the worry is that we never literally see that someone else is angry because there are insuperable obstacles that stand in the way of our literally seeing any such thing. All we can ever literally see in such cases is that some other person is behaving in a certain way or undergoing certain bodily changes, and that is why our knowledge of another's state of mind must be inferential. Because such inferences are very natural we express ourselves by saying that we see that someone else is angry but the objection is that this is no more than a figure of speech; in every case in which we describe ourselves as seeing that someone else is angry what we are really doing is inferring that he is angry from propositions about his behaviour or bodily changes. So if the inferential model is no good then we are really in trouble.

Faced with this line of attack what the perceptual model needs to do is to show that there aren't any insuperable obstacles to seeing and thereby knowing that someone else is angry, and that it therefore isn't

⁵ Cf. Dretske 1969: 118–20.

guilty of the mistake that Mill describes. This is what I now want to do. I want to show that the obstacles that allegedly stand in the way of seeing that someone else is angry can be overcome or dissipated, and that epistemic perception is therefore a bona fide means of acquiring non-inferential knowledge of other minds. What this amounts to is a multiple levels response to the two versions of (HP_{om}). At one level we have the idea that seeing that someone else is angry is a means of knowing that he is angry and therefore also a means of knowing that there are other minds. At the next level we have the attempt to remove the obstacles to literally seeing that someone else is angry. Finally, we can go on to ask what makes it possible to see that someone else is angry, given that it is possible. If all of this can be made to work then we will not only have made some progress with the problem of other minds but also provided another illustration of the value of the multiple levels approach to how-possible questions in epistemology. So let's start by focusing on the project of obstacle-removal since there is no hope for the perceptual model if it's really true that it isn't possible to see what another person is thinking and feeling.

5.2 OBSTACLE-REMOVAL

Imagine this: I'm at a college meeting watching the Bursar closely as one of our less agreeable colleagues makes a speech denouncing his handling of the college's finances. Having been to plenty of college meetings and knowing the Bursar quite well I can see that he is taking it badly. His face is flushed, his fists are clenched, and he appears to be twitching. Then, at the end of the speech, comes the anticipated explosion. Barely able to contain himself, and with his face now crimson, he launches into a violent tirade against the speaker; his vituperation seems to know no bounds and he shouts down all attempts to interrupt or question his words. What, then, is the Bursar's state of mind? It seems safe to say that he is angry. How do I know that he is angry? I can see that he is; I can see the anger in his face. By seeing that he is angry I acquire knowledge of his state of mind by visual means but vision isn't all I have to go on; I can also hear the anger in his voice and in his choice of words.

Although I have focused on seeing that someone else is angry there is nothing special or unique about anger; we can also see that someone else is nervous or depressed or is suffering. As Dretske observes, 'we can see these things, and hence know that they are so, in the same fashion

as we see (epistemically) a great many non-psychological features of our environment' (1969: 189–90). Moreover, seeing that someone is angry is not just a case of epistemic seeing but of primary epistemic seeing. In primary epistemic seeing one sees that *b* is *P*, and thereby knows that *b* is *P*, by seeing *b* itself. When I see that the Bursar is angry, I do so by seeing the Bursar himself. That is how I come to know that the Bursar is angry. And if, for some reason, I start wondering whether other minds exist I can reassure myself that if I can see that the Bursar is angry then I know of the existence of at least one other mind.

With this example in mind let's now address the worry that it isn't possible to see that someone else is angry, and that seeing that the Bursar is angry can't therefore be a means of knowing that he is angry. Why not? One thought is that I can't see that the Bursar is angry because I can't see his anger. So this is quite unlike the case in which I see that the cup into which I'm pouring coffee is chipped. When I see that the cup is chipped I see the cup and I see the chip, that is, the chipped portion of the cup. In contrast, when I allegedly see that the Bursar is angry I see the Bursar but not his anger. All I can see are signs or symptoms of his anger, from which I infer that he is angry. Maybe the inference is so rapid that I don't notice it, and that is why I say that I see that the Bursar is angry. But the fact that this is what I say doesn't make my knowledge of the Bursar's state of mind any less inferential; I can't literally see that the Bursar is angry if in order to see this I would need to see his anger.

Having identified this obstacle to seeing that the Bursar is angry we now have two choices. We can try to *overcome* the alleged obstacle by showing that it is possible for me to see the Bursar's anger or we could try to *dissipate* the obstacle by showing that I don't need to be able to see the Bursar's anger in order to see that he is angry. Let's consider these responses in turn. If emotions like anger are bodily changes then they are straightforwardly observable since bodily changes are straightforwardly observable.⁶ Thus one way of overcoming the obstacle to seeing that the Bursar is angry would just be to identify his anger with his bodily changes. But this seems a high price to pay for the perceptual model; it's hard to believe that the Bursar's twitching, the increasing redness of his face, and so on *are* his anger. The natural thing to think is that emotions like anger have a subjective dimension and that I can't see the

⁶ William James is someone who thinks of emotions as bodily changes. See James 1884 and Pickard 2003 for further discussion of this idea.

Bursar's anger because the way it feels is perceptually inaccessible to me. The behavioural and bodily changes which I see are manifestations of his anger but seeing manifestations of his anger is not the same thing as seeing his anger.

Perhaps, in that case, we should try a different tack. The discussion so far suggests that the supposed obstacle to seeing that the Bursar is angry can't be overcome so maybe we should try to dissipate it instead. Dretske has a nice example that has a bearing on this suggestion. Suppose that I'm looking at a piece of hot metal, call it *b*. If I see that *b* is hot 'by seeing it glow in the way characteristic of hot metal, then I see that it is hot in a primary epistemic manner' (Dretske 1969: 154). This would be a case of primary epistemic seeing as long as I see that *b* is hot on the basis of how *b* looks or behaves rather than on the basis of how some other object—say a temperature gauge—looks or behaves. To be more precise, I see and thereby know that *b* is hot because the following four conditions on primary epistemic seeing are fulfilled: (i) I see *b*, (ii) *b* is hot, (iii) the conditions under which I see *b* are such that *b* wouldn't look the way it looks unless it was hot, and (iv) believing that the conditions are as just described I take *b* to be hot.

Let's agree that these conditions are both necessary and sufficient for seeing *b* is hot. Does it follow from the fact that these conditions are fulfilled that I see *b*'s heat? Maybe not. For a start, it's not clear what it would be to *see* *b*'s heat. One might think that heat can be felt but not seen, and that it can't therefore be literally true that I see *b*'s heat. Nevertheless, I see that *b* is hot. So if I see that *b* is hot without seeing its heat then it's false that one has to see *b*'s heat in order to see that it is hot. In that case, the chipped cup example isn't representative of all primary epistemic seeing and the alleged obstacle to seeing that the Bursar is angry begins to look much less real. For even if one is inclined to say that the Bursar's anger can only be felt (by the Bursar) and not seen, that doesn't mean that I can't see that he is angry. Instead, we could think of seeing that the Bursar is angry as more like seeing that the metal is hot than seeing that the cup is chipped. The sense in which I see that the Bursar is angry is that (i) I see him, (ii) he is angry, (iii) the conditions are such that he wouldn't look the way he looks now unless he was angry, and (iv) believing that the conditions are like this I take him to be angry.

Does this make my knowledge of the Bursar's state of mind inferential? This doesn't follow. Propositions about how the Bursar looks don't figure as premisses in an inference to the conclusion that he is angry.

Although I know that the Bursar is angry by how he looks I do not use the way he looks as a premiss for concluding that he is angry.⁷ As long as the conditions *are* such that the Bursar wouldn't look the way he now looks unless he was angry I *have* in the way the Bursar looks a reason for believing that he is angry. In general, there is a difference between having a reason for believing that b is P and concluding that b is P by some process of reasoning. And there is also a difference between having a reason for believing that b is P and knowing that one has such a reason. Unless there is evidence to the contrary I am entitled to believe that the conditions in which I see b are such that b wouldn't look the way it looks now unless it was P. If this belief doesn't constitute knowledge that only goes to show that I don't know that I see that b is P or don't know that I know that b is P; it doesn't follow that I don't see that b is P or that I don't know that b is P.

So far so good. Starting with the suggestion that I can't see that the Bursar is angry without seeing his anger I quickly gave up on the idea that I can see his anger and instead used the hot metal analogy to argue that the proposed condition on seeing that the Bursar is angry isn't correct. This is an obstacle-dissipating rather than an obstacle-overcoming response. But once it's agreed that it's possible to see that the Bursar is angry the case for denying that I can see his anger begins to look less compelling. In the sense in which I can see that the Bursar is angry I can see the anger in his face, and seeing the anger in his face is a way of seeing his anger. While I don't feel his anger in the way that he feels his anger it doesn't follow that his anger is invisible, any more than it really follows from the fact that I don't feel the heat of a piece of red hot metal that its heat is invisible. We can stipulate that heat or anger can only be felt rather than seen but it's no longer clear what the point of this stipulation is supposed to be once it is conceded that it is possible to see that a piece of metal is hot or that the Bursar is angry.

One of the effects of arguing in this way is to blur the dividing line between the obstacle-overcoming and obstacle-dissipating responses, and that's not necessarily a bad thing. What has happened is that an argument for the view that in order to see that the Bursar is angry it's not necessary to see his anger has evolved into an argument for the view that it *is* possible, in a sense, to see his anger. If this argument is any

⁷ As Dretske emphasizes, epistemic seeing 'does not ... involve a reasoning or inferring that b is P on the basis of what one has seen to be the case or on the basis of how something looks' (1969: 159).

good, we can still describe the Bursar's behaviour and bodily changes as 'manifestations' of his anger but we should refrain from describing them as 'signs' or 'symptoms' of anger. That's because, as Austin remarks:

'Symptoms' or 'signs' of anger tend to mean signs of *rising* or of *suppressed* anger. Once the man has exploded, we talk of something different—of an expression or manifestation or display of anger, of an exhibition of temper and so forth. A twitch of the eyebrow, pallor, a tremor in the voice, all these may be symptoms of anger: but a violent tirade or a blow in the face are not, they are acts in which the anger is vented. (1979: 107)

One can agree with what Austin says in this passage without going back on the idea that emotions like anger aren't identical with patterns of behaviour or bodily changes. To describe an act as one in which the Bursar's anger is 'vented' is not to be committed to thinking that the act *is* his anger; that's the beauty of the notion of a venting. Similarly, one can regard the Bursar's bodily changes as manifestations or displays of his anger without thinking that they *are* his anger. The point of distinguishing between signs and manifestations of the Bursar's anger is that the former are indications of something that is not yet out in the open whereas the latter bring his anger out into the open. What displays or manifestations of the Bursar's anger do is to *reveal* his state of mind without *being* his state of mind; one perceives his anger by perceiving displays of it.

Isn't this all just too good to be true? So far I've been assuming that the obstacle to my seeing that the Bursar is angry is that I can't see his anger, and I have suggested that this obstacle can be dissipated or overcome. It might seem, however, that there is a far more obvious obstacle to seeing that the Bursar is angry that has yet to be addressed. What I have in mind is this: in order to see that the Bursar is angry I would need to be able to eliminate the possibility that he is only pretending to be angry. I can't eliminate this possibility so I can't really see that he is angry any more than I can see his anger. Of course it's true that if the Bursar is pretending to be angry then he is still minded, but this is not something that I can possibly know either. To know that the Bursar is minded I would need to be able to eliminate the possibility that he is in fact a zombie, a being that has no mental life at all but which nevertheless gives the impression of being minded. However, as some of his colleagues are unhelpfully keen to point out, that's not a possibility I can eliminate either; I can't see that the Bursar isn't a zombie because he wouldn't look any different if he was a zombie. And

if I can't eliminate the possibility that the Bursar is a zombie how can I possibly see his anger or see that he is angry?

We now have an obstacle to seeing that the Bursar is angry that takes the form of an epistemological requirement. The obstacle-generating requirement is that in order to see and thereby know that the Bursar is angry I must be able to eliminate the possibility that he is a zombie. The suggestion is that I can't see the Bursar's anger or see that he is angry because I can't eliminate this possibility so we have now arrived at what might be thought of as the ultimate obstacle to knowing that the Bursar is angry by seeing that he is. Faced with this obstacle we have the usual two options. The obstacle-overcoming option would be to show that the requirement can be met; if I can eliminate the possibility that the Bursar is a zombie then there is no problem. Alternatively, if that doesn't work, there is the obstacle-dissipating option of showing that the alleged epistemological requirement is bogus. Again, there is no problem if it's just false that I can't see that the Bursar is angry unless I can eliminate the possibility that he is a zombie.

Let's consider these options in turn, starting with the obstacle-overcoming option. How can I eliminate the possibility that the Bursar is a zombie? One suggestion is that I can eliminate this possibility on purely biological or physiological grounds. Although I haven't investigated this personally I take it that the Bursar has pretty much the same physiology as I do. This suggests that his inner life is unlikely to be fundamentally different from mine. In general, it's reasonable to think that mindedness and biology are correlated, so that possession of a standard human brain goes hand in hand with the ability to enjoy conscious experiences like my own. Since it's possible to tell that the Bursar is biologically normal it's possible to tell, to all intents and purposes, that he isn't a zombie. There is nothing surprising or mysterious about this, it's just a straightforward piece of inference to the best explanation.⁸

Unfortunately, this argument doesn't do what it is supposed to do. For a start, it can't be right to think of inference to the best explanation as the ultimate source of one's knowledge of the Bursar's state of mind if the object of the exercise is to show how my knowledge of his state

⁸ This is Christopher Hill's view in the following passage: 'when I compare my basic sensory and behavioural capacities with those of other people, I find that they are biologically normal, in the sense that they are shared by almost all other members of my species. Hence, probably all other biologically normal members of my species enjoy conscious experiences like my own' (Hill 1991: 212). Note the 'probably'.

of mind could be perceptual. If I *infer* that the Bursar is minded from reasonable assumptions about his biology or physiology then I don't *see* that he is minded. In any case, the issue was whether I can *eliminate* the possibility that the Bursar is a zombie and this possibility has yet to be eliminated. What inference to the best explanation shows is that it's highly unlikely that the Bursar is a zombie and that it's therefore reasonable for me to think that he isn't a zombie, but highly unlikely isn't good enough. What is highly unlikely is still possible so inference to the best explanation can't be a way of meeting the obstacle-generating epistemological requirement on its own terms.

Let's try another tack: if I can see that the Bursar is angry then he can't be a zombie. So I can eliminate the possibility that he is a zombie by seeing that he is angry. This is like McDowell's response to the suggestion that in order to see that some proposition about the external world is true I must be able to eliminate the possibility that I am dreaming. McDowell thinks that I can eliminate this possibility. His idea, which I discussed in Chapter 1, is that my knowledge that I am not dreaming owes its credentials as knowledge to the fact that my senses are yielding me knowledge of the environment, and that this wouldn't be the case if I were dreaming. By parallel reasoning, the present suggestion is that my knowledge that the Bursar isn't a zombie owes its credential as knowledge to the fact that I can see that he is angry, something that wouldn't be possible if he were a zombie. So I eliminate the possibility that he is a zombie by seeing that he is angry; that's how I know that he isn't a zombie.

Clearly, someone who thinks that in order to know anything about the external world by means of the senses one must know that one is not dreaming isn't going to be satisfied by McDowell's proposal. McDowell doesn't show, and doesn't claim to show, that one can know that one is not dreaming *independently* of what one takes oneself to know about the external world by means of the senses. Yet it is conformity to this more demanding requirement that his opponent thinks is necessary for knowledge of the external world. Similarly, those who believe that in order to see that the Bursar is angry I must be able to eliminate the possibility that he is a zombie typically believe that I must be able to eliminate this possibility independently of any appeal to what I think I can see. In other words, I can't eliminate the possibility that the Bursar is a zombie *by* seeing that he is angry because in order to count as seeing that he is angry I must *already* have eliminated this possibility. Since I can't already have eliminated this possibility by perceptual means I

must have eliminated it in some other way. The problem is that there is no other way of eliminating it. And that is why there is no way of meeting the obstacle-generating epistemological requirement when it is understood in the way that its proponents intend it.

Let's agree, then, that the more demanding version of what we might call the *zombie-elimination requirement* can't be met on its own terms. The question we should be asking, therefore, is whether these terms are the right ones, given how demanding they are. This brings us to the possibility of dealing with the obstacle by dissipating it rather than by trying to find a way of overcoming it. An obvious thought is that in order to see that the Bursar is angry it's not necessary that one has already eliminated the possibility that he is a zombie. In order to see that he is angry it must be *true* that he isn't a zombie but that's a very different thing from saying that I must know this independently of knowing that he is angry. The basic confusion here is between the conditions for seeing that something is the case and the conditions for knowing that one sees that something is the case. If I can't eliminate the possibility that the Bursar is a zombie and therefore don't know that the Bursar isn't a zombie it follows that I don't know that I see that he is angry. It doesn't follow that I can't see that he is angry or know that he is angry.

We can bring out the force of this response by relating it to the third and fourth conditions on seeing that the Bursar is angry. The third condition says that I see that the Bursar is a zombie only if the conditions in which I see him are such that he wouldn't look the way he looks now unless he was angry. The fourth condition says that I see that the Bursar is angry only if, believing that the conditions in which I see him are such he wouldn't look the way he looks now unless he was angry, I take him to be angry. The possibility that the Bursar is a zombie is the possibility that the third condition isn't fulfilled so if I can't eliminate the possibility that the Bursar is a zombie then I don't know that the third condition is fulfilled. If I don't know that the third condition is fulfilled then I don't know that I see that the Bursar is angry but I nevertheless see that he is angry. In order to see that the Bursar is angry what is necessary is that the third condition *is* fulfilled; it isn't necessary that I *know* that it is fulfilled. I must, of course, believe that the third condition is fulfilled, and that is what the fourth condition says. If I don't believe that the Bursar must be angry given the way he looks I can't be said to see that he is angry. Yet I can see that he is angry regardless of whether I know that the conditions in which I see him are as I take them to be.

This is obviously not a knockdown argument against the zombie-elimination requirement. After all, those who think that the more demanding version of this requirement constitutes a genuine obstacle to seeing that the Bursar is angry will presumably just deny that they are confusing what is necessary for seeing that he is angry with what is necessary for knowing that I see that he is angry. They will insist that in order to see that the Bursar is angry I must know and not merely believe that the third condition is fulfilled, and that I must therefore already have eliminated the possibility that the Bursar is a zombie. So what are we to say to someone who takes this line? Beyond pointing to a possible confusion between conditions for knowing and conditions for knowing that one knows what else can we do to talk the demanding zombie-elimination requirement out of existence?

The best bet at this point would be to return to the Moorean considerations which came up in Chapter 1. In its most schematic form Moore's idea is this: suppose that *p* is some proposition which I take myself to know and that *R* is an obstacle-generating epistemological requirement that looks like standing in the way of my knowing that *p* because I don't satisfy *R* with respect to *p*. In a situation like this the question I should be asking myself before concluding that I don't really know that *p* is whether I am more certain of the correctness of *R* than I am of the truth of *p*. If so, then *p* is in trouble. If not, or if I am more certain of the truth of *p* than of the correctness of *R*, the rational response will not be to give up on *p*. Instead of questioning *p* on the basis of one's commitment to *R*, the rational response would be to question *R* on the basis of one's commitment to *p*.

In the present context *p* is the proposition that the Bursar is angry and *R* is the zombie-elimination requirement. As I watch the Bursar explode I am more certain that he is angry than I am of the truth of the principle that in order to see that the Bursar is angry I must already have eliminated the possibility that he is a zombie. Since I haven't already eliminated this possibility but can plainly see that the Bursar is angry the more the demanding version of the zombie-elimination requirement can't be right. Acceptable epistemological requirements mustn't have unacceptable epistemological consequences and it's an unacceptable epistemological consequence of the zombie-elimination requirement that I can't see that the Bursar is angry. That's why we should refrain from endorsing this principle, especially when an alternative explanation of its epistemological significance is readily available; as we have just seen, the alternative explanation is that zombie-elimination is required

for knowing that one sees that the Bursar is angry rather than for seeing that he is angry.

It's worth emphasizing that this is not a refutation of scepticism about other minds. At any rate, it's not a refutation of scepticism about other minds on the sceptic's own terms. For no sceptic worth his salt is going to agree that it's an unacceptable consequence of the zombie-elimination requirement that I can't see that the Bursar is angry. His point is that there is no guarantee that our epistemological commitments aren't irrational or unfounded, and that an example of an unfounded epistemological commitment is my commitment to the proposition that the Bursar is angry. Since the sceptic thinks that what calls this commitment into question is precisely the uneliminated possibility that the Bursar is a zombie he will reject as question-begging any attempt to dissipate the zombie-elimination requirement on the basis of my alleged ability to see that the Bursar is angry.

It's fortunate, therefore, that the multiple levels response to (HP_{om}) isn't in the business of trying to prove to the sceptic's satisfaction that I can see that the Bursar is angry. As far as this response is concerned, all we should be trying to do is to establish to our own satisfaction that it's possible to know that the Bursar is angry by seeing that he is angry. The key to doing this is to make it plausible that the zombie-elimination requirement isn't mandatory. I have tried to make this plausible by suggesting that prior zombie-elimination is only necessary for knowing that one knows and by pointing out that conflicts between our epistemological verdicts in particular cases and abstract epistemological requirements that purport to bring our particular verdicts into disrepute are not always to be settled in favour of the abstract requirements. Such requirements can in principle be trumped by our knowledge of particular propositions, and the suggestion I have just been considering is that the possibility of knowing that the Bursar is angry by seeing that he is angry trumps and thereby dissipates the obstacle-generating zombie-elimination requirement. Clearly, there would be nothing to this suggestion if there are independent grounds for thinking that it's impossible to see that the Bursar is angry but there are no such grounds; we certainly shouldn't think that it's impossible to see that the Bursar is angry because it's impossible to see his anger.

To sum up, the perceptual model offers us a presupposed sources solution to the problem of sources. Perception is one of our presupposed sources of knowledge and the perceptual model claims that one can sometimes know what others are thinking or feeling by visual means.

If this is right then one can also know by visual means that there are minds other than one's own. Having got as far as the idea that epistemic perception is a means of acquiring knowledge of other minds we then considered a range of objections to this idea. The fundamental objection was that so-called 'perceptual' knowledge of other minds must be a disguised form of inferential knowledge because there are insuperable obstacles that stand in the way of one's ever seeing what someone else is thinking or feeling. It has turned out, however, that these alleged obstacles are far from insuperable; they can either be overcome or dissipated, and what looks like non-inferential knowledge of other minds really is non-inferential. So that's Level 2 of a multi-levels response to (HP_{omw}) and (HP_{omt}) taken care of.

That leaves Level 3. If seeing that the Bursar is angry is a means of coming to know that he is angry what makes it possible to see that he is angry? What are the enabling conditions for this kind of epistemic seeing? One might wonder why philosophy even has to think about these questions. Minimalists think that it doesn't. Their view is that by Level 2 armchair reflection has already done everything it can do to explain how knowledge of other minds is possible. Although we can ask what makes it possible to see that the Bursar is angry this has to be a question about causal enabling conditions and therefore not one for philosophy. In contrast, moderate anti-minimalists insist that there are a priori enabling conditions for seeing that the Bursar is angry and that philosophy can tell us what they are. Nevertheless, they agree with minimalism that a philosophical response to (HP_{om}) that stops at Level 2 isn't incomplete. The only people who think that a Level 3 response to (HP_{om}) is both possible and necessary are extreme anti-minimalists. They think that we haven't properly explained how knowledge of other minds is possible just by identifying perception as a means of acquiring this kind of knowledge; they maintain that we need enabling conditions as well as means.

Is extreme anti-minimalism any more defensible in relation to (HP_{om}) than in relation to (HP_{pk})? If so why? In order to answer these questions we need a clearer sense of what would count as an a priori enabling condition for the acquisition of knowledge of other minds by visual means. In the next section I will examine two such conditions, the *Spatiality Condition* and the *Identity Condition*. The idea behind the spatiality condition is that it wouldn't be possible to see and thereby know that the Bursar is angry if he lacked spatial properties; his possession of spatial properties isn't a means of knowing that he is angry

but a background necessary condition for knowing that he is angry by seeing that he is angry. What the identity condition says is that it wouldn't be possible for me to see that the Bursar is angry unless I can think of the state he is in when he is angry as a state of the very same type as I am in when I am angry. The suggestion is that grasp of this 'sameness relation' (Peacocke 1999: 110) is a background necessary condition for seeing that the Bursar is angry rather than a means for knowing that he is angry.

To say that it wouldn't be possible to see that someone else is angry if he lacked spatial properties is not to claim, as Strawson does in *Individuals*, that it wouldn't be possible to ascribe states of consciousness to others if they lacked spatial properties. The Spatiality Condition is more modest, and I will attempt to bring out its modesty by comparing it with Strawson's superficially similar claim about the necessary conditions for ascribing states of consciousness to someone else. Then I will move on to the Identity Condition. The problem with this condition is that it's hard to see how it is possible to satisfy it given that the basis for my ascribing anger to the Bursar is so different from the basis for my ascribing anger to myself. I know that the Bursar is angry because I can see that he is angry but I don't know that I am angry by seeing that I am angry. How, then, can 'angry' mean that same thing in 'I am angry' and 'he is angry'?⁹ This looks like another obstacle-dependent how-possible question yet it arises at what I am calling Level 3. Doesn't this put pressure on the alleged distinction between Level 2 and Level 3? I will address this question in the concluding section of this chapter, which is also the point at which I will return to the dispute between minimalism and anti-minimalism. My conclusion will be that extreme anti-minimalism has no more going for it in relation to (HP_{om}) than in relation to (HP_{pk}).

5.3 ENABLING CONDITIONS

I know that the Bursar is angry by seeing that he is angry. I see that the Bursar is angry by seeing the Bursar himself. The Bursar is an object. So

⁹ This is Strawson's worry about the idea that 'X's depression *is* something, one and the same thing, which is felt, but not observed, by X, and observed, but not felt, by others than X' (1959: 109). The question is: 'how can one ascribe to oneself, not on the basis of observation, the very same thing that others may have, on the basis of observation, reasons of a logically adequate kind for ascribing to one?' (1959: 110). Notice that this is a how-possible question. Cf. Peacocke 1999: 110–11.

the enabling conditions for seeing that the Bursar is angry will include the background necessary conditions for object perception. I argued in Chapter 3 that these conditions include the capacity to perceive space or spatial properties. This *Spatial Perception Requirement* (SPR) is as much a requirement on seeing that the Bursar is angry as for seeing that the cup is chipped. This is a reflection of the fact that both the Bursar and the cup are spatial objects. In assuming that the Bursar is a spatial object I am assuming that he has some spatial properties such as shape and extension. In effect, the Spatiality Condition says that it wouldn't be possible for me to see that he is angry if he lacked spatial properties because it wouldn't be possible for me to see *him* if he lacked spatial properties.

What is the relationship between the Spatial Perception Requirement and the Spatiality Condition? In defending SPR in Chapter 3 I stipulated that 'objects' are spatial objects, and one of the questions I was addressing was whether it's possible to perceive a spatial object without perceiving any of its spatial properties. I argued that this is possible but that it's not possible to *see* a spatial object without seeing any of its spatial properties. In contrast, the question raised by the Spatiality Condition is whether it would be possible to see an object which doesn't *have* any spatial properties. Unless it is assumed that objects must be spatial one could think that it's possible to see objects that don't have any spatial properties even if, in the case of objects that *do* have spatial properties, it isn't possible to see *them* without seeing any of their spatial properties. This is what the Spatiality Condition rules out. Although it doesn't stipulate that objects must be spatial, it does imply that objects that have no spatial properties would be invisible. If an object can be seen at all it must have spatial properties, and if it has spatial properties then one can't see it without seeing some of its spatial properties. So there is no such thing as seeing an object without seeing its spatial properties even if there are objects that don't have spatial properties.

Before we can assess the claim that it wouldn't be possible to see the Bursar if he lacked spatial properties we need to be clear about what would it be for the Bursar to have spatial properties. On a Cartesian view what might be called the Bursar's 'thinking self', the subject of his thoughts, experiences, and emotions, is immaterial. This means that the only sense in which he has spatial properties is that he has a body that has spatial properties. On this account, the Bursar is only *weakly* spatial since he is distinct from his body and has no spatial properties in his own right. A *strongly* spatial Bursar would be one

whose thinking self has spatial properties and who therefore doesn't only have spatial properties in virtue of 'having' a body. This is the kind of spatiality that the Spatiality Condition is interested in. When this condition claims that it wouldn't be possible to see that the Bursar is angry if he lacked spatial properties what it means is that it wouldn't be possible to see that the Bursar is angry if he lacked spatial properties 'qua subject'.

For this claim to have any plausibility 'seeing that the Bursar is angry' must be understood as a form of primary epistemic seeing. To bring this out suppose that the Bursar is only weakly spatial. In that case, I could still see that he is angry by seeing what his body looks like. But if the Bursar is distinct from his body then seeing that the Bursar is angry by seeing what his body looks like wouldn't qualify as primary epistemic seeing since strictly speaking it wouldn't be a case in which I see that the Bursar is angry by seeing *him*; schematically, it wouldn't be a case in which I see that *b* is *P* by seeing *b* itself. For primary epistemic seeing to be the source of my knowledge of the Bursar's state of mind the Bursar himself must be visible to me. And what the Spatiality Condition says is that for the Bursar himself to be visible to me, or to anyone else, he must have some spatial properties in his own right and so must be strongly spatial.

Notice that this argument doesn't claim that it would be incoherent to think of the Bursar as only weakly spatial. Perhaps it's true that subjects of emotions like anger must be strongly spatial but that's not the point of the argument for the Spatiality Condition. This argument is only concerned with what is necessary for knowing that the Bursar is angry by one particular means, namely, primary epistemic seeing. Just because the Bursar must be strongly spatial for me to know that he is angry by seeing that he is angry it doesn't follow directly that he couldn't exist without being strongly spatial or that one couldn't conceive of him at all without conceiving of him as strongly spatial. It also doesn't follow that if he weren't strongly spatial it wouldn't be possible for me to know that he is angry by other means. Primary epistemic seeing is one among a range of different ways of coming to know that the Bursar is angry and strong spatiality could be a background necessary condition for knowing that he is angry by seeing the Bursar himself without being a background necessary condition for knowing that he is angry by some other means.

Now for the comparison between the Spatiality Condition and Strawson's superficially similar claim in *Individuals*. Strawson takes it

that 'one can ascribe states of consciousness to oneself only if one can ascribe them to others' (1959: 100). He then argues that one can ascribe states of consciousness to others 'only if one can identify other subjects of experience', and that one can identify other subjects of experience only if they have 'certain corporeal characteristics, a physical situation &c.' (1959: 102). What is supposed to follow from this is that 'the concept of the pure individual consciousness—the pure ego—is a concept that cannot exist; or at least, cannot exist as a primary concept in terms of which the concept of a person can be explained or analysed' (1959: 102). So the basic idea is that it doesn't make sense to conceive of other subjects as lacking corporeal characteristics because this would make it impossible for one to ascribe states of consciousness to them and, by implication, to oneself.

Corporeal characteristics are bodily characteristics. Although they include what I have been calling spatial properties not everything that has spatial properties is a body. Shadows and rainbows are shaped but they aren't bodies. Yet shape is a specifically spatial property. So when Strawson claims that other subjects can only be identified if they have corporeal characteristics he doesn't just mean that they can only be identified if they have spatial properties. What he means is that they can only be identified if they are bodily entities, corporeal objects among corporeal objects. To be bodily entities in the intended sense it isn't enough that other subjects *have* bodies. What is necessary is that they *are* bodies or bodily things. The things to which we ascribe states of consciousness when we ascribe states of consciousness to others must therefore also be things to which we are prepared to ascribe corporeal characteristics. Such things are 'persons' in the Strawsonian sense.

This is an explicitly anti-Cartesian argument. It tries to show that dualism isn't viable because it can't account for the possibility of self-consciousness. Yet it's difficult to believe that Strawson's argument is successful. For example, one question that arises straight away is whether it's true that one must be able to identify other subjects of experience in order to ascribe states of consciousness to them. If 'ascribing' states of consciousness to others is the same thing as *believing* that there are other conscious subjects then this requirement is too strong; it doesn't look as though I have to be able to identify other conscious subjects simply in order to believe that they exist. Strawson only denies this because he thinks that one can't coherently conceive of something that one is unable to identify, but this is an expression of his explicit commitment

to a form of verificationism.¹⁰ Without the verificationism the argument doesn't go through.

Is there any other way of plugging the gap in Strawson's argument? We could try stipulating that to ascribe states of consciousness to others is to *know* of the existence of other subjects of experience and not just to believe that they exist. So now the claim is that one couldn't know of the existence of other subjects without being able to identify them, and this is less obviously objectionable than the claim that one couldn't believe that other subjects exist without being able to identify them. Even so, the argument is still no good. The obvious problem is that beefing up one's conception of what it is to 'ascribe' states of consciousness to others makes it less plausible that one can ascribe states of consciousness to oneself only if one can ascribe them to others. It is one thing for the ascription of states of consciousness to oneself to require the *conception* of other subjects of experience but it's much harder to see why it should require *knowledge* of other subjects.

In any case, it's not obvious that the dualist has any trouble identifying other subjects. As long as other subjects have bodies they can be identified by reference to their bodies; they don't need to be bodily things in any stronger sense in order to be identified in this way. So even if it's true that one can ascribe states of consciousness to others only if one can identify them this doesn't justify Strawson's claim that one must be prepared to ascribe corporeal characteristics to the very things to which one ascribes states of consciousness. The only things to which one must be prepared to ascribe corporeal characteristics in this context are the *bodies* in which other immaterial subjects are embodied. While it's true that on a dualist view there is no guarantee that there is exactly one immaterial subject per body no such guarantee is needed for the purposes of identifying subjects by their bodies; it is enough simply that the one-subject-per-body hypothesis is better than any alternative hypothesis, as it is according to the dualist.

The failure of Strawson's anti-Cartesian argument isn't all that surprising. The argument is a transcendental argument and faces the same basic problem as many other arguments of this form. It makes a claim about what is necessary for self-consciousness but the necessary

¹⁰ Here is a representative passage: 'You do not know what souls are unless you know how to tell one from another and to say when you have the same one again. And if someone should say that this is just old verificationism writ small, or loose, then I am quite content with that' (1997a: 51). Although these words were published a long time after *Individuals* they only make explicit something that was implicit in *Individuals*.

condition that it identifies is too demanding. The Spatiality Condition doesn't face this problem because it isn't concerned with what is necessary for self-consciousness. It isn't even concerned with what is necessary for knowledge of other minds. It is only interested in what is necessary for acquiring knowledge of other minds by primary epistemic seeing and is careful not to assume that the enabling conditions for acquiring knowledge of other minds by primary epistemic seeing are also enabling conditions for the acquisition of knowledge of other minds by other means. Unlike Strawson's argument the approach I have been recommending doesn't make excessively general claims about what is necessary for knowledge of other minds or for conceiving of other minds. For example, the discussion so far leaves it open whether the Bursar must be strongly spatial for me to come to know that he is angry by hearing that he is angry or by reading a report of his anger.

Yet despite its modesty the Spatiality Condition certainly isn't bullet-proof. The obvious question is whether it is actually true that for something to be visible it must have spatial properties. Earlier I argued that it is possible, in a sense, to see the Bursar's anger. Yet his anger doesn't have spatial properties so this looks like a problem for the argument for the Spatiality Condition. Or one might think about the fact that events are visible. Do events have spatial properties? Since they aren't shaped, solid, or extended in space this looks like another potential problem for the Spatiality Condition. It doesn't matter in this context that emotions and events might not qualify as 'objects'. The issue isn't whether *objects* must have spatial properties in order to be visible but whether *anything* that is visible must have spatial properties. And if the answer to this question is 'no' then we can't claim to have shown that the Bursar would be invisible if he lacked spatial properties.

In fact, the counterexamples aren't decisive. The Bursar isn't an event and events are in any case spatially located even though there are many other spatial properties that they plainly lack. As for the Bursar's anger, I haven't claimed that it is straightforwardly visible. My claim was that there is a *sense* in which it is visible, and that the sense in which it is visible is that it is manifested by bodily changes that are straightforwardly visible. Yet both the Bursar's body and his bodily changes are spatial. When one sees the Bursar's anger in his face, as we say, one sees his anger as spatially located and even as extended in space. On the other hand, if talk of the Bursar's anger as something with spatial properties is seen as metaphorical then talk of 'seeing' his anger must also be regarded as metaphorical. The reason is that it is built into our conception of what

it would be to literally see something that it involves the perception of at least some spatial properties. Since there is no such thing as non-spatial seeing there is no such thing as seeing something non-spatial. That is why it continues to seem compelling that Cartesian egos couldn't be seen and, by the same token, that the Bursar would be invisible if he lacked spatial properties.

In arguing in this way I haven't relied on claims about seeing that can only be established empirically. For example, I haven't relied on the claim that visibility requires the capacity to reflect light or the claim that only things with spatial properties can reflect light. The key to the argument for the Spatiality Condition is a conception of what makes seeing seeing that can be established by armchair reflection. That is why the Spatiality Condition is not just an enabling condition but an a priori enabling condition for knowing that the Bursar is angry by primary epistemic seeing, an enabling condition that can be established non-empirically. But this doesn't detract from the modesty of the Spatiality Condition. Crucially, it doesn't follow from the fact that the Bursar must have spatial properties to be visible that he must be a *corporeal* object to be visible. In theory, he would be visible even if, like a shadow, he has spatial properties without being a bodily thing. In practice, however, the obvious way for him to have spatial properties is for him to be embodied. And while the tightness of the connection between visibility and possession of spatial properties is, at least to some extent, open to debate, the extent to which this is so doesn't make it remotely plausible that one could see that the Bursar is angry by seeing the Bursar himself if the Bursar himself were non-spatial.

Turning to the Identity Condition, the idea is this: to see that the Bursar is angry one must have the concept of anger. Like other psychological concepts, the concept of anger is applicable in the very same sense to oneself and to others; 'angry' means the same thing in 'I am angry' and 'he is angry'. This means that to have the concept of anger one must be able to think of someone else's anger as a state of the very same type as one's own anger. To think of one's own anger and someone else's anger as states of the very same type is to grasp a sameness or identity relation. Grasp of this relation is therefore a background necessary condition for seeing that the Bursar is angry. And since we can know that this is so just by thinking about it the Identity Condition is not just an enabling condition but an a priori enabling condition for seeing that the Bursar is angry.

So far so good but now we run into a potential obstacle. When I ascribe anger to the Bursar I do so on the basis of observation yet I don't ascribe anger to myself on the basis of observation. Another way of putting this would be to say that self-ascriptions and other-ascriptions of anger have different 'assertibility-conditions'; what warrants my assertion that the Bursar is angry is fundamentally different from what warrants my assertion that I am angry. Given that this is so, how is it possible for me to understand that I am ascribing states of the same type in the two cases? How is grasp of this sameness relation possible?¹¹ This looks like an obstacle-dependent how-possible question, and I have already conceded that the fact that it arises at Level 3 looks like a problem for the suggestion that what goes on at this level can be sharply distinguished from what goes on at what I have been calling Level 2, the level of obstacle-removal. But before worrying about that let's just concentrate on the latest how-possible question on its own terms. As usual we have two options. We can either be dissipationists or we can accept that the obstacle is genuine and show how it can be overcome. Let's start by examining some dissipationist options before looking at the prospects for an obstacle-overcoming response.

The first dissipationist option would simply be to deny that self-ascriptions and other-ascriptions of anger have different assertibility-conditions or that I don't ascribe anger to myself on the same basis as I ascribe it to others. Someone who thinks this would have to think that I ascribe anger to myself on the basis of observation, and this isn't wholly implausible.¹² Observation of one's own behaviour or responses to the Bursar might give one grounds for judging that one is angry with him and perhaps there is even a sense in which one might be said to 'observe' one's own feelings of anger. But the sense in which one 'sees' that one is angry is still very different from the sense in which one sees that the Bursar is angry so the basic problem remains. And when one considers examples other than anger the problem seems even more acute. It's hard to believe, for example, that I ascribe pain to myself on the basis of observation of my own behaviour, or that I can 'observe' that I am in pain in anything like the sense in which I can observe that someone else is in pain.

¹¹ Cf. Strawson 1959: 110–15.

¹² Ryle is someone who thinks that our ways of finding out about our own minds are not fundamentally different from our ways of finding out about other minds. See Ryle 1949.

A different dissipationist option would be to claim that even a genuine difference in assertibility-conditions shouldn't make it difficult to think of one's own anger and the Bursar's anger as states of the same type. There is only a problem here if one assumes that meaning is tied to assertibility-conditions but that's an assumption the dissipationist doesn't feel compelled to accept. His idea is that the fact that I ascribe anger to myself and to the Bursar on different grounds isn't a good reason for thinking that 'angry' can't mean the same thing in 'I am angry' and 'He is angry'. By the same token, it isn't a good reason for thinking that his anger and mine are states of different types. Consequently, it doesn't stand in the way of thinking that they are states of the same type. What it is for the Bursar to be angry is just what it is for me to be angry regardless of any difference in the epistemology of self-ascription and other-ascriptions.

Although there is something right about this dissipationist response it can't be the end of the story. It's all very well saying that a difference in assertibility-conditions needn't amount to a difference in meaning but that doesn't explain away the intuition that there must be a difference in meaning in such circumstances; we need to remember that denying the existence of an obstacle isn't the same as dissipating it. In any case, it isn't as if the problem only arises if one has a prior commitment to an assertibility-conditions theory of meaning. There is an intuitive puzzle here rather than one that is exclusively generated by some dubious assumption about meaning. The dissipationist needs to show that the puzzle isn't genuine and that means explaining *how* the idea that the basis on which one ascribes anger to oneself is so different from the basis on which one ascribes it to others can fail to put significant pressure on the idea that it's the same kind of thing that is ascribed in the two cases. So far the dissipationist hasn't done much more than assert that these ideas *are* compatible.

In the light of these doubts about dissipationism perhaps this would be a good time to consider the prospects for an obstacle-overcoming account of how it is possible to satisfy the Identity Condition. The challenge is to explain how it is possible to get one's mind around the idea that the Bursar's anger and my anger are states of the same type despite the intuitive considerations that make this idea seem problematic. From the perspective of the perceptual model the way forward is clear. Suppose that I'm a good friend of the Bursar and am angered by the attack on him at the college meeting; in fact, I'm just as angered by the attack on him as he is. So what do I see when I look at the Bursar? I see that he

is angry but that is not all. I also see that he is just as angry as I am. In other words, I see that the state he is in is no different from the state I am in, just as I sometimes see that someone else is as frightened or nervous or depressed as I am. In such cases, the identity of mental state is a *presented* or *perceived* identity, and this is what makes it possible for me to think of the Bursar's state and mine as states of the same type.

This is a natural extension or application of the perceptual or observational model. As McGinn, who is no fan of this model, puts it:

If we allow that objects may be perceived to have the same property and that another's experience can (sometimes) be perceived ... , then I cannot see that there is any greater difficulty about the possibility of perceiving that another has the same experience as oneself: just consider being seasick with someone else or being in a group sauna. An adherent of the observational model of our knowledge of other minds would hold that in such cases one does come to know by 'direct perception' that the other has the same experience as oneself, and so has an experience whose content is correctly specified in this way. It is hard to see how this specific claim could be rejected without objecting to the observational model in general. (1984: 122)

Notice that the perceptual model is not claiming that the ability to see that the Bursar is in the same state as oneself is a necessary condition for being able to think of the two states as the same, any more than it is claiming that one must be able to see that the Bursar is angry in order to know that he is angry. The suggestion is only that perception is one route to grasp of the sameness relation. If I couldn't see that the Bursar is angry or that his reaction is the same as mine perhaps there would be other ways for me to come to know that he is angry or to understand that he feels what I feel. What we are looking for are not necessary conditions but ways of grasping a sameness relation. If perception can help me to grasp this relation then it can also provide me with *a* means of thinking of my state and the Bursar's state as the same, and therefore with *a* means of overcoming the obstacle to thinking that the two states are the same; it needn't be the only means.

To recap, the Identity Condition is a background necessary condition for seeing that the Bursar is angry. Whereas the correctness of the Spatiality Condition is at least open to debate the issue with the Identity Condition isn't whether it's true that in order to see that the Bursar is angry one must be able to think of his anger and one's own as states of the same type. What we want to know is what makes it possible to think of self-ascriptions and other-ascriptions of anger as ascriptions of the very same property. We now have an account of what makes this

possible for someone with the appropriate perceptual capacities. If I can see that the Bursar is angry I can see that he feels what I feel, and if I can see that he feels what I feel then I am well on the way to satisfying the Identity Condition on grasp of the concept of anger. All I have to do is to get from the perceived identity of his psychological state and mine to the thought that they are identical, and that doesn't seem an especially difficult thing to do.

The Spatiality and Identity Conditions are obviously not the only background necessary conditions for seeing that the Bursar is angry. For example, it's plausible that I can only see that he is angry because he manifests his anger. If, like one of Putnam's 'super-spartans', the Bursar never evinces anger I could never see that he is angry, though I might still come to know that he is angry in other ways.¹³ It's also important that the Bursar manifests his anger in ways that are recognizable as manifestations of *anger*. In practice this isn't a problem because the Bursar is a human being who expresses anger in familiar and intelligible ways. This is the relevance of Wittgenstein's remark that 'only of a living human being and what resembles (behaves like) a living human being can one say: it sees; is blind; hears; is deaf; is conscious or unconscious' (1978: # 281). What is wrong with this remark is the implication that it only makes sense to ascribe anger to such a being. What is much more plausible, however, is that only a living human being or what resembles a living human being in certain ways can be *seen* to be angry. That is why embodiment matters. It is not that the Bursar must have corporeal characteristics for me to see that he is angry because he must have corporeal characteristics for me to see *him* but because it wouldn't be possible for me or anyone else to see that he is *angry* if he lacked the right configuration of corporeal characteristics.

Having identified some of the background necessary conditions for knowing that the Bursar is angry by seeing that he is angry we are left with two issues. The first is whether we should be moderate or extreme anti-minimalists. We no longer have to worry about minimalism because that position has already been refuted. The minimalist denies that there are any a priori enabling conditions for acquiring knowledge of others' minds by visual means but this must be wrong if the Spatiality Condition

¹³ A community of super-spartans would be one in which 'the adults have the ability to successfully suppress *all* involuntary pain behaviour.... They do *not* wince, scream, flinch, sob, grit their teeth, clench their fists, exhibit beads of sweat, or otherwise act like people in pain or people suppressing the unconditioned responses associated with pain. However, they do feel pain, and they dislike it (just as we do)' (Putnam 1975: 332).

and the Identity Condition are both genuine examples of a priori enabling for seeing that the Bursar is angry. The best way of showing that these conditions can be established non-empirically is to establish them non-empirically, and that's what I have just done. The remaining question, therefore, is whether the identification of a priori enabling conditions is merely possible or whether it is necessary for the purposes of answering (HP_{om}).

The other unresolved issue concerns the suggestion that my account of the Identity Condition threatens to collapse the distinction between Level 2 and Level 3 in relation to (HP_{om}). As we have seen, what began as an attempt to identify a particular background necessary condition for seeing that the Bursar is angry rapidly transformed itself into an exercise in obstacle-removal. Since obstacle-removal is something that is supposed to happen at Level 2 the obvious implication is that Level 3 is an extension or continuation of Level 2 rather than a distinct level of enquiry. But if we accept this line of reasoning in this context then perhaps we should also accept it in relation to other how-possible questions, and that looks like a problem for the multi-levels framework that I have been recommending. It's now time to consider whether this problem is genuine and, if so, whether it is serious.

5.4 HOW MANY LEVELS?

Suppose we get as far as the idea that it's possible to know that the Bursar is angry by seeing that he is angry. By reflecting on the possibility of acquiring knowledge of other minds by visual means we seem to have come up with an answer to (HP_{omw}) and to (HP_{omt}). Then we start to worry about the viability of this answer because of the apparent obstacles that stand in the way of seeing that the Bursar, or anyone else, is angry. However, when we think a bit more about the alleged obstacles to this form of epistemic seeing we see that they can be overcome or dissipated. So where do we go from here? The moderate anti-minimalist thinks that while we *can* go on to ask what makes it possible to see that the Bursar is angry there is no *need* to ask this question. He thinks that we have already answered (HP_{omw}) and (HP_{omt}) at Level 2, and that the answer we have given is complete as it stands. This is what the extreme anti-minimalist disputes. His thought is that simply identifying means of knowing about other minds and removing obstacles to knowing about other minds by those means isn't enough; there is a further explanatory

gap that needs to be filled and it can only be filled by proceeding to Level 3.

How are we to understand the extreme anti-minimalist's notion of an 'explanatory gap'? We can start to get a fix on this notion by returning to the discussion of geometry in Chapter 1 since this provides the clearest illustration of what the extreme anti-minimalist has in mind. Let's assume for the sake of argument that Kant is right to regard construction in intuition as a means of acquiring geometrical knowledge, and that the singularity of constructed figures isn't an obstacle to this method of acquiring this kind of knowledge. The next question is: what makes it possible for imaginative construction to be a source of a priori knowledge, or indeed any knowledge, of the geometry of physical space? This is a naturally arising question because there appears to be a mismatch between Kant's account of the source of geometrical knowledge and our usual conception of its subject matter. If physical space, the subject matter of geometry, is non-mental, how can operations on mental diagrams be a source of geometrical knowledge? The only explanation, according to Kant, is that space is transcendently ideal; only the ideality of space can make it intelligible that construction in intuition is a source of geometrical knowledge, and that is a reason for thinking that space *is* ideal.

By representing the transcendental ideality of space as an enabling condition for the acquisition of geometrical knowledge Kant is effectively denying that space is wholly mind-independent. He is conceding that the mind-independence of physical space *would* make it impossible for intuitive construction to deliver knowledge of its geometry, and he deals with this problem by saying that space isn't mind-independent after all. This may seem a high price to pay for getting rid of the apparent mismatch between source and subject matter but Kant's point is that it is a price that has got to be paid. The sense in which it has got to be paid is that an account of how geometrical knowledge is possible that only talks about the role of construction in geometrical proof and stops there is incomplete. It leaves an obvious explanatory gap, the gap created by the apparent mismatch, and the only way of filling in this gap is to endorse transcendental idealism.

This looks like an argument for extreme anti-minimalism. What Kant is saying is that an account of what makes it possible to extract geometrical knowledge from operations on mental diagrams isn't an optional extra for his purposes, and that's just what one would expect an extreme anti-minimalist to say. On the other hand, extreme anti-minimalists

also think that the identification of enabling conditions is a Level 3 project rather than a Level 2 project, and that now looks doubtful. The enabling conditions that Kant focuses on are, in effect, obstacle-removing conditions. In other words, the fact that they are fulfilled removes what would otherwise be an obstacle to acquiring geometrical knowledge by construction in intuition, and that is why he thinks that it's not just possible but also necessary to say something about them. Yet obstacle-removal was supposed to be the business of Level 2 rather than Level 3. So if extreme anti-minimalism is defined by its commitment to a separate Level 3 then it isn't vindicated by Kant's discussion. Instead, there is now the prospect of Level 3 collapsing into Level 2. What justifies the demand for enabling conditions is the need to close an explanatory gap, and explanatory gaps that have been left open are really just obstacles that haven't been removed.

Compare this with the other minds case. Is there any sense in which someone who proposes that it's possible to know that the Bursar is angry by seeing that he is angry has failed to provide a complete explanation of the possibility of knowing that the Bursar is angry? Is there still an explanatory gap that needs to be filled? The first thing to say is that there isn't an obvious mismatch in this case between source and subject matter. Whereas the construction of mental diagrams doesn't *look* like the kind of thing that is capable of delivering knowledge of the geometry of non-mental reality, seeing that the Bursar is angry does look like a way of telling that he is angry. The interesting question is whether it's possible to see that the Bursar is angry, and that is a question that gets addressed at Level 2. As for the Level 3 enabling conditions that I have been discussing, the Spatiality Condition and the Identity Condition, they aren't obstacle-removing in the way that transcendental idealism is an obstacle-removing enabling condition for the acquisition of geometrical knowledge. In the other minds case, therefore, Level 3 is both optional and distinct from Level 2.

This amounts to a defence of moderate anti-minimalism in relation to (HP_{om}). If the Bursar lacked spatial properties it wouldn't be possible to see that he is angry but that doesn't mean that one hasn't fully explained how it is possible to know that he is angry if one has failed to mention the Spatiality Condition. Acknowledging this condition is no more essential for the purposes of explaining how it is possible to know that the Bursar is angry than referring to the existence of the Channel Tunnel is essential for the purposes of explaining how it is possible to get from London to Paris in less than three hours. If the

Bursar lacked spatial properties then that would be an obstacle to seeing that he is angry but that doesn't make the Spatiality Condition an obstacle-removing condition; it isn't motivated by, or a response to, some pre-existing, intuitive problem with the idea that epistemic seeing is a source of knowledge of other minds. While necessary conditions for seeing that the Bursar is angry are potentially obstacle-*generating*, in the sense that not meeting such a condition would prevent one from seeing that the Bursar is angry, this doesn't make them obstacle-*removing*. That is why, in this case, Level 3 doesn't collapse into Level 2.

In that case, what should we make of the fact that it can be difficult to understand how allegedly Level 3 conditions are fulfilled? For example, we saw it can be difficult to understand how self-ascriptions and other-ascriptions of anger can be ascriptions of the very same property and spent some time trying to overcome this obstacle. It doesn't follow, however, that discussion of the Identity Condition belongs at Level 2. The obstacles that get addressed at this level are directly means-related; they are obstacles to knowing about other minds by various specific means. In contrast, the obstacles that we might find ourselves having to deal with at Level 3 aren't directly means-related. Rather, they are obstacles to the satisfaction of the *enabling conditions* for knowing about other minds by various specific means. Although we can talk about obstacles in both cases they are different kinds of obstacle. In addition, there needn't be any obstacles to the satisfaction of Level 3 conditions. For example, there isn't anything that stands in the way of the Bursar evincing his anger yet the fact that he evinces his anger is an enabling condition for anyone else to see that he is angry. That is another difference between Level 2 and Level 3. Level 2 exists only in order to tackle specific obstacles to the acquisition of knowledge by specific means. Level 3 has some independent rationale; it has positive explanatory ambitions that go beyond the dissipating or overcoming of obstacles.

Having said all that, there is no need to go to the stake for the sake of maintaining a sharp distinction between Levels 2 and 3. We should think of the different levels of a multiple levels response to an epistemological how-possible question as different stages of a single, evolving enquiry. The idea is that the questions addressed at each level follow naturally from the questions addressed at the preceding level, and this doesn't require one to conceive of the different levels as rigid blocks. As far as (HP_{om}) is concerned, the fact is that there are lots of different ways of coming to know about other minds, and that the obstacles

to these different ways of knowing won't all be the same, any more than their background necessary conditions will all be the same. Among these ways of knowing is epistemic seeing. By helping us to understand how the relevant form of epistemic seeing is possible and what makes it possible the multiple levels framework provides us with a plausible answer to (HP_{omw}) and (HP_{omt}) . And that's another illustration of the value of this framework.

6

A Priori Knowledge

6.1 THE PROBLEM

So far in this book I have tried to explain how perceptual knowledge is possible and how knowledge of other minds is possible. In this final chapter I want to look at a third how-possible question, namely:

(HP_{apk}) How is a priori knowledge possible?

If, as I have been arguing, how-possible questions are obstacle-dependent the key to making any progress with (HP_{apk}) is to identify the obstacles to the acquisition or existence of a priori knowledge which give this question its bite. Having identified these obstacles we can then examine the prospects for a multi-levels response to (HP_{apk}). Predictably, the basic idea of this chapter is that such a response promises to cast at least as much light on (HP_{apk}) as on other epistemological how-possible questions. What I propose to do, therefore, is to identify means of coming to know things a priori, show how obstacles to the acquisition of a priori knowledge by the suggested means can be overcome or dissipated, and consider whether it is either possible or necessary to give a substantive account of what makes it possible to come to know things a priori by these means.

So what is the intuitive obstacle to the existence of a priori knowledge which philosophers who press (HP_{apk}) might have in mind? That obviously depends on how a priori knowledge is characterized. Kant famously characterizes it as knowledge that is 'absolutely independent of all experience' (B3) but how are we to understand the notions of 'experience' and of 'independence' from experience? To keep things as simple as possible, let's start by supposing that 'experience' in this context is sense experience, and that genuinely a priori knowledge is knowledge that is 'justificationally' independent from experience in this sense.¹ Specifically, a belief won't constitute knowledge unless it is

¹ This is Peacocke's way of putting things. See Peacocke 2000: 256.

adequately justified or warranted, and one knows that *p* a priori only if one's warrant or justification for believing that *p* is true is independent of sense experience. On this account, there can be a priori knowledge only if there can be a priori justification or warrant, and a justification or warrant is a priori 'if neither sense experience nor sense-perceptual beliefs are referred to or relied upon to contribute to the justificational force particular to that warrant' (Burge 1998: 3).

Unfortunately, this attempt to say what a priori knowledge is immediately runs into problems. For example, suppose that I am introspectively aware that I am currently thinking about the a priori and that I am thereby warranted in believing that I am thinking about the a priori. Is this kind of warrant perceptual? No, one might say, because introspection isn't a form of sense perception. Yet introspective warrant isn't a form of a priori warrant so it can't be sufficient for a warrant to be a priori that it isn't perceptual.² To block this objection one would either need to insist that introspection is a form of 'inner' perception or, perhaps more plausibly, that there is no good reason not to count introspective warrant as a priori.³ But now we run into another problem. Suppose that I am warranted in believing that I was thinking about the a priori yesterday by the fact that I can remember thinking about the a priori yesterday. This kind of memory warrant is non-perceptual but there is good reason not to count it as a priori. The good reason is that philosophers who have been attracted to the idea that some warrants are a priori have never intended the notion of a priori warrant to cover memory warrant.⁴

These considerations suggest that what it is and is not acceptable to count as an a priori warrant is partly a matter of tradition. It's unacceptable to count memory warrants as a priori not because they are perceptual but because it is built in to traditional ways of thinking about a priori warrant that memory warrant isn't a species of it. In contrast, it's much more difficult to decide whether it is unacceptable to count introspective warrants as a priori because this is something

² This criticism of Burge is due to Goldman. As he points out, 'introspection can give rise to warrant, but its type of warrant is neither perceptual nor a priori' (2002: 32).

³ Paul Boghossian is an example of someone who is prepared to count introspective warrants as a priori. He writes: 'I have always found it natural to regard a priori knowledge as encompassing knowledge that is based on no experience as well as knowledge that is based purely on *inner* experience' (1997: 363).

⁴ Actually the issue of whether memory warrants can be a priori is much more complex than this brief discussion suggests. Burge has some important things to say about this in his 1993 paper but they don't affect the point I am making here.

about which there has been much less agreement. The significance of this will become clearer as we go along. In the meantime, we can still conclude, at least on the basis on the discussion of memory warrant if not the discussion of introspective warrant, that ‘one cannot equate *a priori* warrant with non-perceptual warrant’ (Goldman 2002: 32). A priori warrant can’t be equated with non-perceptual warrant because memory warrant is neither perceptual nor a priori. This means that we still lack an adequate general characterization of a priori warrant, and, by extension, an adequate general characterization of a priori knowledge.

A further consideration is this. Although I have so far followed the usual practice of reading the Kantian notion of a priority as justificational a priority it’s not clear that this is the only or the best way of seeing things. For Kant, the fundamental contrast is between a priori and empirical knowledge. Empirical knowledge ‘arises out of experience’ (B1) and a priori knowledge is knowledge that doesn’t arise in this way. To say that a kind of knowledge does or doesn’t arise out of experience is to make a point about its *source*, about how it comes to be, so it seems that the emphasis on justification or warrant in the characterization of a priori knowledge is misplaced. Maybe the point about a priori knowledge is not that it is justificational independent of experience but that it has a non-experiential source. We can call this kind of independence *genetic* independence, so the present suggestion is that a priori knowledge is knowledge that is genetically rather than justificational independent of experience.

In fact, it’s a mistake to think that justificational and genetic considerations are unrelated.⁵ For present purposes, the ‘source’ of one’s belief that *p* is how one came to believe that *p*, and the extent to which one is justified or warranted in believing that *p* will depend, at least in part, on how one came to believe that *p*. So even if one thinks that one can’t know that *p* unless one believes that *p* and has a justification for believing that *p*, one can still agree with Peacocke that ‘the status of a belief as knowledge depends on how it is reached’ (2000: 264). How a belief is reached is relevant to its status as knowledge because it is relevant to its status as a warranted or justified belief. So, for example, if *S* believes that the cup is chipped because he can see that it is chipped that tells one something about the source of *S*’s belief and thereby also tells one something about its warrant. On this account, the

⁵ Alvin Goldman has always emphasized this point in his epistemological writings. See, for example, Goldman 1986.

distinction between a priori and empirical knowledge is fundamentally a distinction between 'different ways of coming to know that something is the case' (Peacocke 2000: 255). A priori ways of coming to know are non-experiential ways of coming to know and a priori knowledge is knowledge that has its source in an a priori way of coming to know. We can still say that a priori knowledge is justificational independent of experience but that is because this kind of independence is a reflection of its genetic independence.

With this account of a priori knowledge in the background let us now return to (HP_{apk}). What is the intuitive obstacle to the existence of a priori knowledge which might lead one to press this question? What, exactly, is the problem? This is how Paul Boghossian describes the situation:

Traditionally, three classes of statements have been thought to be objects of *a priori* knowledge: logical statements, exemplified by such truths as: either Brutus killed Caesar or he did not; mathematical statements, such as: $7 + 5 = 12$; and conceptual truths, for instance, all bachelors are unmarried. The problem has always been to explain how any statement could be known *a priori*. After all, if a statement is known *a priori*, then it must be true. And if it is true, then it must be factual, capable of being true or false. What could possibly entitle us to hold a factual statement true on *a priori* grounds? (1997: 334)

What we can extract from this passage are two ideas. The first is that the supposed objects of a priori knowledge are factual statements. The second is that the problem of a priori knowledge is fundamentally the problem of explaining how a priori knowledge of the truth of a factual statement is possible. But why is this a problem? If A is a factual statement, why is there anything problematic about the idea that one might be entitled to hold A true on a priori grounds or that there are non-experiential ways of coming to know that A is true? To see the force of (HP_{apk}) we need to understand *why* A's factuality is at least a potential obstacle to our coming to know that A is true other than on the basis of experience. As things stand, this looks like an unargued assumption.

The natural move at this point would be to turn to Hume. For Hume, the ultimate source of our knowledge of matters of fact is experience. Experience here includes sense experience, introspection, and memory. Causal reasoning can take us beyond the evidence of our memory and senses but Hume maintains that knowledge of causal relations is itself derived from experience and that causal reasoning must 'terminate in

some fact, which is present to your memory or senses' (1975: 46). But if experience, broadly construed, is the ultimate source of our knowledge of matters of fact then it is hard to see how a priori knowledge is possible, at least if we take it that a priori knowledge is knowledge of matters of fact.⁶ A priori knowledge is only possible if knowledge of matters of fact can have a non-experiential source, and this is precisely what Hume is calling into question. What we have here is therefore another version of what I have described in earlier chapters as the *problem of sources*. On this account, (HP_{apk}) is a good question because a priori knowledge lacks any obvious source.

Let's call the thesis that experience is the ultimate source of our knowledge of matters of fact Hume's Thesis. And let's call the thesis that a priori knowledge would have to be knowledge of matters of fact the Factuality Thesis. What we have seen is that (HP_{apk}) arises when Hume's Thesis is combined with the Factuality Thesis. As a result, we are now in a position to explain why the fact that some statement A is factual looks like standing in the way of our having a priori knowledge of its truth. A's factuality becomes an obstacle given Hume's thesis. But Hume's thesis is the core thesis of empiricism so another way of putting this would be to say that it is a prior commitment to empiricism which, at least in conjunction with the Factuality Thesis, makes the existence of a priori knowledge look so problematic. Without Hume's thesis, or something like it, in the background the question 'What could possibly entitle us to hold a factual statement true on a priori grounds?' would seem much less troublesome.

We now run into the following difficulty: the initial challenge was to identify an intuitive obstacle to the existence of a priori knowledge which might lead one to ask (HP_{apk}) but not every potential obstacle is an intuitive obstacle. The crucial question, therefore, is whether Hume's Thesis and the Factuality Thesis have any intuitive backing. If not, then the fact that someone who puts them together is going to find it difficult to account for the possibility of a priori knowledge is neither here nor there as far as the project of motivating (HP_{apk}) is concerned. After all, why worry about the fact that Hume's Thesis and the Factuality Thesis make it hard to understand how a priori knowledge is possible if one has no independent reason to accept either thesis? The next challenge, therefore, is to explain why these theses might strike one as compelling

⁶ This is the basis of Ayer's insistence that 'there can be no *a priori* knowledge of reality' (1946: 115).

despite the fact that they call the possibility of a priori knowledge into question.

The intuitive case for the Factuality Thesis is straightforward: if there is no such thing as a true statement which isn't factual then there couldn't be non-factual a priori knowledge.⁷ And one reason for thinking that there is no such thing as a true statement which isn't factual is that, as Boghossian remarks, there is no better answer to the question, 'what makes a statement true?' than 'the world' or 'the facts'.⁸ This line of thinking has been disputed by some empiricists but let's assume for the moment that it is successful, and that the Factuality Thesis is therefore in good shape. That leaves Hume's Thesis. What is the intuitive case for the empiricism to which this thesis gives expression? I want to suggest that what might lead one to think that experience is the ultimate source of our knowledge of matters of fact is a commitment to realism, and that the attractions of empiricism are partly the attractions of the realism which underpins it.

This might seem a surprising suggestion since empiricism is more commonly associated with idealism than with realism.⁹ Nevertheless, it is not difficult to understand why empiricism and realism might be regarded as natural allies. We have just seen that if A is a true statement then what makes it true is 'the world' or 'the facts'. But if, as realism claims, 'the world' is something that doesn't depend for its existence on human thought or perception then it is plausible that the only way of finding out whether A is true or false is by means of experience, by going and looking. The suggestion, then, is that experience is our sole source of information about the world beyond our minds because only experience, or inferences from experience, can put us in touch with what exists beyond our minds.¹⁰

One difficulty with this suggestion is that empiricists don't just think that experience must be the source of our knowledge of the world beyond our own minds. They also think that experience, in the form of introspection, must be the source of our knowledge of our own minds. Yet the operations and contents of our own minds aren't

⁷ In which case logical empiricists like Ayer must be mistaken when they claim that 'the propositions which we know to be valid independently of all experience, are so only in virtue of their lack of factual content' (1946: 115).

⁸ See Boghossian 1997: 336.

⁹ Indeed Berkeley seems to think that idealism is an inevitable consequence of empiricism.

¹⁰ See Markie 2004 for further discussion.

mind-independent so it can't be the non-mentality of matters of fact which explains why experience is the ultimate source of our knowledge of them; after all, some facts, such as the fact that I am at present thinking about the a priori, are mental facts. It's worth bearing in mind, however, that for empiricists like Locke statements about the operations and contents of one's own mind are no less factual than statements about the world beyond one's mind. Both kinds of statement are made true by how things are in the world. It's just that 'the world' includes one's own mind. That is why introspection, which is conceived of by the empiricists as a form of 'inner' looking and seeing, is allegedly the only means by which one can come to know whether statements about what is going on in one's own mind are true or false.

Let's continue to assume, therefore, that realism and empiricism are natural allies. Experience is the ultimate source of our knowledge of matters of fact, whether mental or non-mental, because genuinely factual statements are capable of being true or false, what makes them true or false is the way the world is, and only experience can tell us how things stand in the world. Much of what we know about the external world derives from ordinary sense perception whereas our knowledge of our own minds derives from 'inner' perception. In both cases, however, it is the brute factuality of matters of fact which explains why experience is the ultimate source of our knowledge of them. So the obstacle to the existence of a priori knowledge is no longer just empiricism. The obstacle is a combination of empiricism and realism. The realism motivates the empiricism and thereby makes it difficult to see how a priori knowledge of reality, of genuine matters of fact, is possible.

Faced with this obstacle, some might be inclined to conclude that a priori knowledge is impossible because the obstacle is insuperable. Let's call this response to (HP_{apk}) a *pessimistic* response. Faced with the question 'how is a priori knowledge possible?' pessimists say 'it isn't'.¹¹ Pessimism is therefore a form of scepticism. In contrast, *optimists* think that a priori knowledge is possible because they think that the alleged obstacle isn't insuperable. They believe that they can explain how a priori knowledge is possible by showing how the obstacle can be overcome or dissipated. I am an optimist. In previous chapters I have argued that there are a priori enabling conditions for the acquisition of perceptual knowledge and knowledge of other minds, and I have tried to identify some of these conditions. The project of identifying a priori enabling

¹¹ Quine is the paradigmatic pessimist. See, for example, Quine 1953.

conditions for the acquisition of various different kinds of knowledge by various different means would be doomed if the pessimist is right since a priori enabling conditions are just ones which can be known a priori. That is why, for better or worse, I am committed to optimism.

My brand of optimism tries to break the link between realism and empiricism. It rejects Hume's Thesis but not because it rejects realism. It agrees that a priori knowledge is knowledge that has its source in a non-experiential way of coming to know and is satisfied that a priori knowledge is possible because it is satisfied that there are non-experiential ways of coming to know about non-mental reality. Reflection is a non-experiential source of knowledge. I can come to know, just by engaging in the appropriate form of reflection, that nothing can be red all over and green all over at the same time. Reasoning, including philosophical reasoning, looks like another non-experiential source of knowledge. In Chapter 3, for example, I reasoned my way from a set of metaphysical and epistemological premisses to the conclusion that spatial perception is a background necessary condition for the acquisition of knowledge by means of epistemic perception. My argument for this Spatial Perception Requirement was therefore an a priori argument. A third source of a priori knowledge is calculation. When I come to know that $68 + 57 = 125$ by calculating the sum of 68 and 57 in my head my way of coming to know isn't experiential in any interesting sense of 'experiential'. That is why the resulting knowledge is a priori.

Three questions will need to be addressed with respect to each of these allegedly non-experiential ways of coming to know. One is whether it is really a way of coming to *know* anything. Another is whether I have succeeded in identifying sources of knowledge of *matters of fact*. According to Hume, for example, the knowledge that $68 + 57 = 125$ is knowledge of a 'relation of ideas' rather than knowledge of a matter of fact. A third question with respect to each proposed way of coming to know is whether it is really an a priori or non-experiential way of coming to know. I will be arguing that the answer to each of these questions is 'yes'. In my view the difficult question is not whether there are sources of a priori knowledge, since there clearly are, but rather: *what makes it possible* to acquire knowledge of 'the world' or 'the facts' by non-experiential means? This is the point at which it might be tempting to abandon realism but I'm going to argue that this is a temptation which can and should be resisted. To begin with, however, let's not worry about such deeper explanatory questions and concentrate instead

on making it plausible that reflection, reasoning, and calculation are non-experiential sources of knowledge.

6.2 REFLECTION, REASONING, AND CALCULATION

What is 'reflection'? To get a handle on this, suppose that in the course of thinking about the incompatibility of colours I ask myself whether something can be both red and green. I imagine a flag the left half of which is red and the right half of which is green. Would it be correct to judge that the flag is red? Yes, in a sense. Would it be correct to judge that the flag is green? Again, yes, in a sense. But if there is a sense in which the flag is red and a sense in which it is green then there is a sense in which something can be both red and green. On the other hand, it's clear that the flag wouldn't be uniformly red and green, so the next question to consider is whether something can be red and green *all over*. In an attempt to make some progress with this question I imagine a flag or some other surface that is red all over and then try to think of circumstances in which it would be correct to judge that the very same surface is green all over. Can I think of such circumstances? Yes, by thinking of the very same surface being red all over at one time and green all over at another time. This leads me, finally, to ask whether something can be red all over and green all over *at the same time*. And this is where I run into a block: I realize, on reflection, that circumstances in which it would be correct to judge that a particular surface is red all over at a given time are precisely circumstances in which it would not be correct to judge that the surface is green at that time. Equally, circumstances in which it would be correct to judge that a particular surface is green all over at a given time are precisely circumstances in which it would be correct to judge that the surface is not red at that time. So I conclude, on this basis, that nothing can be red all over and green all over at the same time.

In arriving at this conclusion I draw on my grasp of concepts like *red all over* and *green all over*, and on my ability to determine when these concepts do and when they don't apply. I come to know that nothing can be red all over and green all over at the same time just by thinking about what it would be correct to judge in various circumstances, and that is why it is appropriate to characterize my knowledge in this case as having its source in what Peacocke calls 'a priori reflection' (2000:

268). To say that 'reflection' is the source of my knowledge that nothing can be red all over and green all over at the same time is to make the point that my knowledge is grounded in my understanding of the proposition and in my ability to think.¹² And what justifies the labelling of 'reflection' as an a priori source of knowledge, and therefore as a source of a priori knowledge, is the fact that it isn't an experiential route to knowledge. If I come to know that nothing can be red all over and green all over at the same time just by reflecting on what it would be correct to judge in various circumstances then this is not a case in which my knowledge 'arises out of experience'. It is a case in which my knowledge arises out of thought and understanding.

None of this would be of much interest if the incompatibility of colours is a special case but there is no reason to think that this is so. There are many other kinds of knowledge that reflection can deliver, including philosophical knowledge. In Chapter 3, for example, I talked about Locke's Principle, which states that no two physical objects can occupy exactly the same region of space at exactly the same time. I suggested that this principle can be established by 'armchair reflection', and we can now see what this suggestion comes to: one can come to know that no two physical objects can occupy exactly the same region of space at exactly the same time by reflecting on the fact that the only circumstances in which it would be correct to judge that two physical objects occupy exactly the same region of space at exactly the same time are ones in which the two objects are of different kinds. The basis of this reflection is one's grasp of the concept of a physical object, so this looks like another case in which a piece of knowledge is grounded in a priori reflection. Indeed, it is a priori reflection which tells one that a priori knowledge is knowledge that has its source in an a priori or non-experiential way of coming to know. In other words, my account of the a priori is itself a priori.

For the moment, this is as much as I want to say about 'reflection', even though many important questions about this alleged source of a priori knowledge remain unanswered. Before coming back to these questions, let's consider some other non-experiential sources of knowledge to see how they fit into the overall picture. Suppose I know that Tony Blair is the Prime Minister and that the Prime Minister lives in Downing

¹² This is the point of Peacocke's description of a priori reflection as 'understanding-based reflection' (2000: 268). The argument of this paragraph owes a lot to Peacocke's discussion.

Street. Putting these pieces of knowledge together I conclude that Tony Blair lives in Downing Street. I come to know that Tony Blair lives in Downing Street by reasoning, by inferring that he lives in Downing Street from other things I know. So is my knowledge that Tony Blair lives in Downing Street, arrived at in this way, a priori? On the one hand, one might be tempted to say that coming to know that something is the case by reasoning isn't an experiential way of coming to know that it is that case. On the other hand, one might also be reluctant to describe my knowledge that Blair lives in Downing Street as a priori given that my knowledge of the premisses of the inference is presumably not a priori.

One possibility in a case like this would be to describe one's knowledge as partly a priori since reasoning is one of its sources. In Goldman's terminology, the conclusion 'has an element or component of a priori warrant, simply because there is one strand of its warrant that is a priori' (2002: 36). Cases in which one's knowledge of the conclusion of a piece of reasoning is straightforwardly a priori will be ones in which one's knowledge of the premisses is also a priori. This will be the case if reflection is the source of one's knowledge of the premisses. So, for example, if one is able to reason one's way to the Spatial Perception Requirement from metaphysical and epistemological premisses which one knows to be correct by a priori reflection then one's knowledge of the Spatial Perception Requirement is straightforwardly a priori. What makes it a priori is that it has two sources, reasoning and reflection, both of which are non-experiential.

That leaves calculation. I know, for example, that $68 + 57 = 125$. How do I know it? By calculating the sum of 68 and 57. Whereas there might be doubts about the respectability of what I have been calling 'reflection' there are usually no such doubts about the respectability of calculation as a source of knowledge. Competently calculating the sum of 68 and 57 is the best possible way of coming to know that $68 + 57 = 125$. Most of us are capable of acquiring mathematical knowledge in this way and there should be little temptation to assimilate knowledge whose source is calculation to knowledge whose source is experience. More cautiously, there should be little temptation to assimilate the two sources of knowledge in cases in which the calculations are performed without the assistance of a calculator or computer. If at least unassisted calculation is a non-experiential source of knowledge, then a priori knowledge is possible. Since unassisted calculation *is* a non-experiential source of knowledge, a priori knowledge is possible.

In the terminology of Chapter 1, what I have just provided is a Level 1 or 'means' response to (HP_{apk}). A Level 1 response to an epistemological how-possible question explains how whatever knowledge is at issue is possible by identifying means by which one can acquire it. The identification of reflection, reasoning, and calculation as various means by which a priori knowledge can be acquired makes it plausible that a priori knowledge is possible, just as the identification of epistemic seeing as a means by which knowledge of the external world can be acquired makes it plausible that knowledge of the external world is possible. Since the identification of reflection, reasoning, and calculation as sources of a priori knowledge carries no obvious commitment to idealism, or to thinking that there is anything wrong with the Factuality Thesis, it looks as though we have already gone a long way towards vindicating optimism. It's a fact about the mind-independent 'world' that nothing can be red all over and green all over at the same time but this doesn't mean that our knowledge of this fact can't be a priori. So it is Hume's Thesis rather than realism or the Factuality Thesis which needs to be given up if what I have been arguing is along the right lines.

But let's not be too hasty. The identification of means of knowing or of coming to know is rarely enough to satisfy philosophers who ask obstacle-dependent epistemological how-possible questions since the obstacles which led such philosophers to ask their questions also need to be tackled. After all, it's not as if epistemologists who worry that (HP_{apk}) has no satisfactory answer are unaware of the suggestion that reflection, reasoning, and calculation are sources of a priori knowledge. They are perfectly familiar with this suggestion and are also well aware that *if* these really are sources of a priori knowledge then it must be false that there are insuperable obstacles to the existence of this kind of knowledge. But they don't think that this observation gets one very far because they think that there *are* insuperable obstacles to the acquisition of a priori knowledge of matters of fact by reflecting, reasoning, or calculating. Since this is the basis of their pessimism, the only way of making further progress with this issue is to shift to Level 2, the level of obstacle-removal. The challenge at this level is to show that there is nothing that stands in the way of acquiring a priori knowledge of matters of fact by the proposed means, either because the alleged obstacles can be overcome or because they can be dissipated.

So what are the obstacles which need to be dealt with at Level 2? First, there are potential obstacles to the acquisition of *knowledge* by the proposed means. The worry is that when we think about what it

would be for a source to be a source of knowledge rather than mere belief we see that reflection, reasoning, and calculation can't be sources of knowledge. Second, there are potential obstacles to the acquisition of knowledge of *matters of fact* by the proposed means. The worry here is that reflection, reasoning, and calculation can't deliver knowledge of *matters of fact* even if they can deliver knowledge of 'relations of ideas' or some such. Finally, there are potential obstacles to the acquisition of a priori knowledge by the various means which I have identified. The suggestion, in other words, is that when we think a bit harder about what it would be for a piece of knowledge to be genuinely a priori we will come to see that reflection, reasoning, and calculation can't supply us with this kind of knowledge.

Let's take a closer look at these alleged obstacles, starting with obstacles to the acquisition of knowledge from the proposed sources. Why should one think that reflection, reasoning, and calculation can't be sources of knowledge? A relatively superficial answer to this question would be one which points out that these sources aren't error-proof. So, for example, however carefully one reflects on the incompatibility of colours one can't rule out the possibility that one is mistaken in concluding that nothing can be red all over and green all over at the same time; perhaps there is a way for something to be uniformly red and uniformly green at the same time which one has somehow managed to overlook. But if reflection isn't error-proof then what it provides one with can't be knowledge. Similarly, one can't rule out the possibility that one has miscalculated the sum of 68 and 57. Miscalculation is always a possibility, in which case calculating the sum of 68 and 57 isn't a way of coming to know its sum.

The problem with this argument is that it threatens to prove too much. To see why, think about perception. It's not as if perceptual errors aren't possible but we don't take this to show that perception isn't a source of knowledge. But if we don't take the possibility of misperception to show that perception isn't a source of knowledge why should we take the possibility of miscalculation to show that calculation isn't a source of knowledge? If we insist that genuine sources of knowledge must be error-proof then we aren't going to be left with many sources of knowledge. That is, of course, the point of scepticism, but the interesting question is whether there are any *special* problems with the idea that reflection, reasoning, and calculation are sources of knowledge. The argument from error is an argument for a generalized scepticism; it doesn't identify an obstacle to the acquisition

of knowledge by means of reflection, reasoning, and calculation which, if taken seriously, wouldn't also be an obstacle to the acquisition of knowledge by perceptual and other equally familiar means. If we suspect that we have misperceived the answer is to look again. Similarly, if we suspect that we have miscalculated the sum of 68 and 57, or overlooked a counterexample to the proposition that nothing can be red all over and green all over at the same time, the obvious remedy is to check the calculation or to think again. The sceptic is worried about errors which somehow remain undetected however many times one checks or thinks again, but there is no good reason to think that the elimination of errors which are incapable of being eliminated is necessary for knowledge.

It might seem that this response doesn't get to the heart of the matter. To see why not, consider the following principle:

A potential knowledge source K can yield knowledge for S only if S knows that K is reliable.

This is Stewart Cohen's 'KR Principle', and it quickly leads to problems for the suggestion that the various sources which I have identified are sources of knowledge.¹³ Take calculation as an example. KR implies that calculation can provide me with arithmetical knowledge, such as the knowledge that $68 + 57 = 125$, only if I know that calculation is reliable. To know that calculation is reliable I would need to know that the arithmetical beliefs it produces are largely true. But the only way of knowing that would be by means of calculation, and that would appear to make the whole procedure unacceptably circular. Assuming that the reliability of a particular epistemic source can't be established by relying on that very source the upshot is that I can't know that calculation is reliable. And if I can't know that, then calculation can't be a source of knowledge if KR is correct; calculating the sum of 68 and 57 can't be a way of coming to know its sum and therefore can't be a way of coming to know its sum a priori.

A quick response to this line of argument would be to dismiss it as a *reductio* of the KR Principle. If this principle really calls into question the capacity of calculation to deliver arithmetical knowledge then that just goes to show that there is something wrong with KR; calculation *is* a source of arithmetical knowledge so there must be something wrong with any principle which implies that it can't be. In a way this is right but we need to move more slowly. We need to bear in mind that the person

¹³ See Cohen 2002 for more on the KR Principle.

we are trying to convince is the pessimist, that is, someone who thinks that there are insuperable obstacles to the acquisition of knowledge by means of reflection, reasoning, or calculation. According to the pessimist, the KR principle is one such obstacle. If the only argument against this principle is that it threatens to make the acquisition of arithmetical knowledge by means of calculation impossible then this is unlikely to carry much weight with someone who is already convinced that calculation isn't a source of knowledge. Maybe this isn't a very sensible thing to think but it would help if there were some independent arguments against KR.

One such argument is that this principle is just too demanding. This is a point which externalists often make against higher-order epistemological requirements. After all, even the very young and the philosophically untrained can come to know that $68 + 57 = 125$ by working it out, just as they can acquire knowledge of the world around them by means of perception. Yet it is extremely implausible to suppose that they are even capable of forming higher-order beliefs about the reliability of their epistemic sources, let alone that these beliefs constitute knowledge of the reliability of their sources. A much more natural view is that at least some epistemic sources are basic, in the sense that they can 'deliver knowledge prior to one's knowing that the source is reliable' (Cohen 2002: 310). Clearly, not all sources are basic. For example, non-standard or unusual sources like clairvoyance and blindsight could only be non-basic sources of knowledge at best; they would only yield knowledge if one knows independently that they are reliable. But perception and calculation aren't like that. To say that they are basic sources of knowledge is to say that they are innocent until proved otherwise, and this means that they are not subject to KR.

This amounts to an obstacle-dissipating response to the suggestion that calculation and other non-experiential epistemic sources are not sources of knowledge. What is alleged to prevent them from being sources of knowledge is their failure to satisfy KR, and the proposed response to this allegation is the rejection of KR on several closely related grounds: it's too demanding, it doesn't apply to basic sources, and it makes it difficult to see how even simple calculations could deliver knowledge. Yet proponents of KR have what look like equally weighty arguments in support of this principle. Here is one: the conditions for knowledge include justified belief, and justified belief is epistemologically responsible belief. Since one is being epistemologically irresponsible in relying on epistemic sources whose reliability one hasn't

already established it follows that a potential knowledge source can yield justified beliefs, and therefore knowledge, only if one knows that the source in question is reliable. Indeed, even if there are some sources which are somehow not subject to KR, it isn't clear that all of the non-experiential sources which I have identified are among them. For example, it might be argued that as potential sources of knowledge go reflection is no more perspicuous than clairvoyance or blindsight, and that the case for regarding it as a non-basic source, and therefore as subject to KR, is just as strong as the case for regarding clairvoyance and blindsight as non-basic. What this brings out is that it is one thing to show that some knowledge sources are exempt from KR and another to show that a particular source is exempt, especially when the source in question is non-experiential.

What we now have is a stand-off in relation to KR. What the pessimist regards as an epistemological requirement which justifies his pessimism because reflection, reasoning, and calculation fail to meet it is seen by the optimist as bogus for the very same reason. So where do we go from here? Perhaps it would be worth pausing to consider whether we should really have conceded that reflection, reasoning, and calculation don't satisfy KR. If there is a satisfactory way of knowing that these sources are reliable then we would have the possibility of an obstacle-overcoming response to pessimism rather than an obstacle-dissipating response. But what would such a response look like? Let's start with calculation. Suppose that I calculate the sum of 68 and 57 in my head and come up with the answer '125'. Do I now know that $68 + 57 = 125$? KR says that I only know this if I know that my mental arithmetic is reliable. But this is something which it *is* surely possible for me to have established on independent grounds. Suppose that I have regularly checked my mental arithmetic using a calculator or by asking a mathematically competent friend, and have discovered that I rarely make mistakes. Don't I therefore count as knowing that my mental arithmetic is reliable? Yet I haven't established its reliability by doing more mental arithmetic, so this isn't a case in which I have established the reliability of an epistemic source by relying on that very source. Since I have used other means to check its reliability, there is no reason not to claim conformity with KR.

Reflection is trickier but there is still such a thing as knowing that I am more or less reliable when it comes to deciding the truth of a proposition just by reflection. For example, there might have been plenty of occasions on which I have thought that some proposition about the incompatibility of two properties is true because I have been

unable to think of how it could be false. If it has often happened that I or someone else has subsequently come up with a convincing counterexample then that would be a reason for thinking that my reflections aren't very reliable. If, on the other hand, they have stood the test of time and the scrutiny of others then I can know on this basis that they are reliable. Again, this wouldn't be a case in which I have established the reliability of an epistemic source by relying on that very source. Instead, I would have independent, non-circular grounds for concluding that my reflective beliefs are generally true.

Let's agree, then, that there is a sense in which I can know that at least some of my non-experiential knowledge sources are reliable. Perhaps my knowledge of the reliability of these sources is inductive and therefore empirical but this doesn't make these sources themselves empirical. Mental arithmetic is still a non-experiential source regardless of how I come to know that it is reliable. It's also worth pointing out that I need memory in order to establish that my mental arithmetic has a good track-record, and this raises a question about the reliability of memory. If one can't know that one's own memory is reliable without relying on memory then this might be a problem for someone who thinks that no epistemic sources are exempt from KR. The present question, however, is not whether it makes sense to apply KR across the board, that is, whether it makes sense to think of all epistemic sources as non-basic, but whether reflection and calculation can be regarded as non-basic sources of knowledge. They can be so regarded if knowledge of their reliability is possible, and I have just explained how this kind of reliability knowledge is possible given that memory is a source of knowledge.

Pessimists who justify their pessimism on the basis of KR might object that all of this is beside the point. According to such pessimists the question is not whether I can know that *my* calculations are reliable but whether it is possible to know that calculation *per se* is reliable. Similarly, the question is whether it is possible to know that reflection *per se* is reliable, not whether it is possible for me to know that my reflections are reliable. To see the significance of this move, take the case in which I use a calculator or ask a friend to check my mental arithmetic. How is the checking done? The obvious answer is: by calculating. So if the question was whether *calculation* is a reliable epistemic source then we are none the wiser. Using a calculator or asking someone else to check one's mental arithmetic can't be a way of establishing the reliability of calculation as such so we have still not shown how calculation can satisfy KR; the obstacle has yet to be overcome.

At this point, however, the case for pessimism is beginning to look distinctly shaky. The original thought was that a potential knowledge source K can yield knowledge for S only if S knows that K is reliable. In the case in which I work out that $68 + 57 = 125$ by doing the sum in my head what is K ? If K is *mental arithmetic* then its reliability can be independently checked. If K is *calculating* then it is more difficult to see how its reliability can be checked other than by just doing more calculations. But this leaves the optimist with considerable room for manoeuvre. He can insist that in my example the source is mental arithmetic, and therefore that its reliability can be checked independently. Or he can agree that the source is something more generic like 'calculating' but insist that the reliability of such a generic epistemic source can be established using that very source. This would open up the possibility of calculation being used to establish the reliability of calculation and reflection being used to establish the reliability of reflection.

If this is along the right lines then the lesson is that there isn't a straightforward argument from KR to pessimism. In particular, to make it plausible that reflection and calculation aren't sources of knowledge several supplementary assumptions are needed: the assumption that KR is well motivated, that it isn't too demanding, that epistemic sources can't provide one with knowledge of their own reliability, and that when one works out that $68 + 57 = 125$, or that nothing can be red all over and green all over at the same time, one lacks knowledge of the reliability of one's knowledge sources. Some of these assumptions are needed to block obstacle-overcoming responses to the suggestion that reasoning, reflection, and calculation can't be sources of knowledge while others are needed to block obstacle-dissipating responses. What we have seen, however, is that each of the pessimist's assumptions is questionable at least to some degree, and that is why optimism is still in the running. Yet KR is no longer being rejected simply on the basis that it implies that the non-experiential sources which I have been discussing can't be sources of knowledge. It's not that the fact that an epistemological requirement such as KR has potentially disastrous epistemological consequences isn't relevant when it comes to deciding its acceptability or otherwise. It clearly is relevant, but so are the independent worries about KR which I have identified.

This is as much as I propose to say in this chapter about the worry that reflection, reasoning, and calculation can't be sources of knowledge. What about the worry that they can't deliver knowledge of matters of

fact? What is supposed to prevent them from doing that? Here is one way of explaining the difficulty: the propositions which we can know to be true by means of reflection, reasoning, or calculation are all necessarily true. For example, it is a necessary truth that nothing can be red all over and green all over at the same time or that $68 + 57 = 125$. But necessary truths 'are entirely devoid of factual content' (Ayer 1946: 105), and this is the obstacle to the acquisition of a priori knowledge of matters of fact by means of reflection, reasoning, or calculation. They can't supply us with knowledge of contingent truths and therefore can't supply us with knowledge of matters of fact. So we still lack an answer to Boghossian's question. Recall that his question was: what could possibly entitle us to hold a factual statement true on a priori grounds? What has now emerged is that reflection, reasoning, and calculation can perhaps entitle us to hold certain statements true on a priori grounds or, as one might prefer to say, provide us with a priori knowledge of the truth of certain statements. But if the statements in question are non-factual then we still haven't explained what we were supposed to be explaining; we haven't explained how a priori knowledge of *matters of fact* is possible.

Why should anybody think that necessary truths are devoid of factual content? For logical empiricists like Ayer the point is that necessary truths are analytic, that is, true in virtue of meaning, analytic truths are disguised tautologies, and tautologies say nothing. Ayer gives the example of the proposition: 'either some ants are parasitic or none are.' Since this proposition 'provides no information about the behaviour of ants, or indeed, about any matter of fact' (1946: 105) it is devoid of factual content. Analytic propositions 'simply record our determination to use words in a certain fashion' (1946: 112), from which it follows that a priori knowledge of an analytically necessary proposition by means of reflection, reasoning, or calculation wouldn't amount to what Ayer calls '*a priori* knowledge of reality' (1946: 115). Only experience can provide us with knowledge of genuinely factual propositions, and that's the point of Hume's Thesis.

Optimists have little to fear from this line of argument since it fails to identify a genuine obstacle. For a start, it's not obvious that reflection, reasoning, and calculation can only deliver knowledge of necessary truths. As we have seen, I can know by reasoning from contingently true premisses that Tony Blair lives in Downing Street, yet it's certainly not a necessary truth that Tony Blair lives in Downing Street. What is necessarily true is that if Tony Blair is the Prime Minister and the Prime Minister lives in Downing Street then Tony Blair lives in Downing

Street. But the fact that this *conditional* is necessary obviously doesn't mean that knowledge of the *conclusion* of the inference is knowledge of a necessary truth. What about truths knowable by reflection and calculation? While 'Nothing can be red all over and green all over at the same time' and ' $68 + 57 = 125$ ' are both examples of necessary truths it is a further question whether they are analytic in Ayer's sense. Indeed, it is doubtful whether there are *any* analytic propositions in his sense since, as Quine points out, there is no such thing as true proposition whose truth does not 'hinge on reality' (1970: 10).¹⁴ This was the idea behind the intuitive argument for the Factuality Thesis, and Ayer's discussion does nothing to undermine this argument. Since neither contingent nor necessary truths can coherently be regarded as true by virtue of anything but 'traits of reality' (Quine 1976: 113) both kinds of truth are equally 'factual'. So if we can know a priori that nothing can be red all over and green all over at the same time, or that $68 + 57 = 125$, then it follows straightforwardly that a priori knowledge of matters of fact is possible.

That leaves the pessimist with just one option. Having been forced to concede that reflection, reasoning, and calculation are sources of knowledge, and that they provide us with knowledge of matters of fact, his only remaining hope is to try to make it plausible that they aren't sources of a priori knowledge. This seems a pretty tall order. I have argued that a priori knowledge is knowledge that has its source in a non-experiential way of coming to know, and that reflection, reasoning, and calculation are non-experiential sources of knowledge. This means that the knowledge which they make available must be a priori knowledge unless it is a mistake to regard them as non-experiential sources or as non-experiential ways of coming to know. But how can this possibly be a mistake? How can reflection, reasoning, and calculation be regarded as falling on the experiential side of the divide between experiential and non-experiential ways of coming to know?

From a Quinean perspective, there is a simple answer to these questions: for a source of knowledge or justified belief to be genuinely non-experiential it would have to be the case that its deliverances are incapable of being empirically undermined or defeated by future experience. A priori knowledge requires empirical indefeasibility yet even the results of reflection, reasoning, and calculation aren't indefeasible in this sense. That is why they aren't non-experiential sources of knowledge. It's true that there aren't other ways of coming to know which have a better

¹⁴ Cf. Boghossian 1997: 335–7.

claim to be regarded as non-experiential but that only means that there are no non-experiential sources of knowledge. And if there are no non-experiential sources of knowledge then there is no a priori knowledge.

On this account, the empirical indefeasibility requirement on non-experiential ways of coming to know is an insuperable obstacle to the existence of a priori knowledge. An obstacle-overcoming response to Quine's argument would therefore be one which shows that the results of reflection, reasoning, and calculation are indefeasible in the requisite sense. One might point out, for example, that it is extremely difficult to conceive of circumstances in which it would be rational to revise the belief that $68 + 57 = 125$. Yet this is something which Quine can take in his stride since, as Giaquinto observes, his view is not that 'for any current belief, however firmly held, we are currently able to imagine a set of circumstances in which it would be rational to reject it' (1996: 253). Quine's claim, against which it is hard to argue, is that such circumstances might obtain, whether or not we can currently conceive of them.

A better bet, therefore, would be to question the alleged indefeasibility requirement on non-experiential ways of coming to know and, by implication, on a priori knowledge. The obvious thought here is that a cognitive state's having a non-experiential source is one thing and its being empirically indefeasible is another. The latter is not a necessary condition for the former since it doesn't follow from the fact that a cognitive state can be defeated by experience that one came to be in that state on the basis of experience. For example, if one believes that $68 + 57 = 125$ because one has worked it out on one's head then the resulting knowledge is a priori because mental arithmetic is a non-experiential route to knowledge regardless of whether arithmetic is immune to revision. This is an obstacle-dissipating response to Quine's argument. There is no legitimate basis for denying that reflection, reasoning, and calculation are sources of a priori knowledge because Quine's indefeasibility requirement doesn't constitute a legitimate basis for denying that they are non-experiential sources of knowledge. It doesn't constitute a legitimate basis for denying that reflection, reasoning, and calculation are non-experiential sources of knowledge not because they meet the indefeasibility requirement but because it isn't a genuine requirement.

How can claims about what is and is not a genuine requirement on non-experiential sources be justified in the absence of a definition of the notion of an experiential source? This question is worth asking

because when I first introduced the distinction between experiential and non-experiential ways of knowing I didn't define 'experiential' or attempt to give a general account of the conditions under which a particular way of knowing falls on one or other side of the divide. The reason is that the notion of an experiential way of knowing can no more be defined than that of experience itself. But this doesn't mean that we have no idea whether it applies in particular cases. Intuitively, seeing or feeling that a particular cup is chipped are paradigmatically experiential ways of knowing that the cup is chipped, whereas calculating that $68 + 57 = 125$ is a paradigmatically non-experiential way of knowing that $68 + 57 = 125$. The appropriate classification of other cases depends on their similarities to these paradigms. So, for example, if we don't know whether to classify introspection as an experiential or a non-experiential source of knowledge that is because it seems to fall somewhere between the paradigms. For all that, we retain an intuitive grasp of the experiential/non-experiential distinction, and there is nothing in our intuitive grasp of this distinction which suggests that it boils down to the distinction between defeasible and indefeasible ways of knowing.

To sum up, I have explained how a priori knowledge is possible by identifying three non-experiential sources of knowledge. I have shown that there aren't insuperable obstacles to regarding these sources as sources of knowledge, as supplying us with knowledge of matters of fact, or as genuinely non-experiential. But if reflection, reasoning, and calculation provide us with non-experiential knowledge of matters of fact then Hume's Thesis is false; it is false that experience is our only source of knowledge of matters of fact. Yet at no point have I suggested that reflection, reasoning, or calculation can only provide us with knowledge of mental facts so I haven't committed myself to a narrowly idealist conception of the objects of a priori knowledge. It's a 'fact' that $67 + 58 = 125$ but not a mental fact.

Where does this leave the suggestion that realism and empiricism are natural allies? I mentioned this suggestion as a way of explaining the intuitive pull of empiricism; the idea was that the realist thesis that 'the world' is something that doesn't depend for its existence on human thought or perception makes it plausible that the only way of finding out about it is by experience, by looking and seeing. But we have now seen that calculating that $67 + 58 = 125$ is a non-experiential way of coming to know that $67 + 58 = 125$ even though the fact that $67 + 58 = 125$ doesn't obtain in virtue of human thought or

perception. This implies that the alliance between empiricism and realism is less close than empiricists sometimes suppose. One can be an empiricist without being a realist, as Berkeley showed, and it now appears that one can also be a realist without being an empiricist; one can think that the world doesn't depend for its existence on human thought or perception without thinking that experience is our only source of worldly knowledge.

The natural move at this point would be to turn to Level 3 of my multiple levels response to (HP_{apk}). Having identified several different means by which a priori knowledge is possible, and removed potential obstacles to the acquisition of a priori knowledge by these means, we might now consider what makes it possible to acquire a priori knowledge in the ways that I have been describing. One proposal is that reflection, reasoning, and calculation can only deliver a priori knowledge because the facts known by these means are, if not narrowly mental facts, at least ones which obtain in virtue of various human conventions or practices. This implies that it is a mistake to think that realism can dispense with Hume's Thesis. Since I have just been arguing that this is not a mistake now would probably be a good time to explain why this latest attempt to saddle realism with Hume's Thesis fails.

6.3 IDEALISM AND EXPLANATION

Kant is the best example of a philosopher who sees Hume's Thesis as an inevitable and unwelcome consequence of realism. Because he thinks that Hume's Thesis is plainly false he concludes that the only way out is to give up on realism, at least in its transcendental form.¹⁵ Indeed, this is his big argument for transcendental idealism. We have non-experiential knowledge of matters of fact, realism can't account for this, so realism has got to go unless it is the 'empirical' realism which Kant thinks is compatible with transcendental idealism. But why does Kant take the falsity of Hume's Thesis as given? Because he takes it as given that some of our knowledge, including mathematical knowledge, is synthetic a priori knowledge. If mathematical knowledge is synthetic then it is knowledge of matters of fact, and if it is a priori knowledge then it doesn't have its source in experience. It just can't be true, therefore, that experience is our only source of knowledge of matters of fact.

¹⁵ There is a helpful account of Kant's thinking in Craig 1987: 236–43.

As we saw in Chapter 1, Kant's positive account of our geometrical knowledge turns on the notion of 'construction in pure intuition'. To construct a geometrical figure in pure intuition is to visualize it or to 'draw' it in the imagination. Having done this, we are then in a position to acquire synthetic a priori geometrical knowledge by carrying out various operations on the constructed figure. What this brings out is that geometrical proofs, as Kant conceives of them, are fundamentally diagrammatic, but this raises an obvious question: how can construction in pure intuition tell us anything about the geometry of *physical* space as distinct from the geometry of the 'visual' space in which geometrical figures are constructed? This question is especially pressing for realism about space. As Dummett puts it on Kant's behalf, 'if we assume that space is a feature of objective reality independent of us, it is hard to see why "the nature of our intuitive abilities" should afford us any sure guide to its constitution' (1991: 150). Short of positing an intrinsic harmony between independent reality and the form of our intuitions the only alternative for realism is to refuse to offer any explanation of the epistemological power of construction in pure intuition.

This is the point at which transcendental idealism comes into its own since it purports to explain what transcendental realism can't explain. Transcendental idealism 'explains the matter by treating the ground of our knowledge as also being the ground for the truth of the propositions known, that is, by identifying that whereby we know them with that in virtue of which they are true' (Dummett 1991: 150). In Kant's terminology, it is only because space itself is an '*a priori* intuition' that 'has its seat in the subject' (B41) that constructing figures in pure intuition can be a way of acquiring synthetic a priori knowledge of its geometry. This can either be read as a Level 2 explanation or as a Level 3 explanation. Realism about space would constitute an insuperable obstacle to the acquisition of synthetic a priori geometrical knowledge by means of construction, and the only way of dealing with this obstacle is to abandon realism. At the same time, abandoning realism and adopting an idealist conception of space also puts us in a position to answer a what-makes-it-possible question. For we can now say that it is the ideality of space which makes it possible for construction in pure intuition to have the epistemological power which we know it has.

In essence, therefore, Kant's 'geometrical' argument for transcendental idealism is an explanatory argument. To recognize construction in pure intuition as a source of synthetic a priori knowledge is effectively to reject Hume's Thesis but it turns out that only idealism is capable

of explaining how construction in pure intuition can be a source of synthetic a priori knowledge. That is why, according to Kant, realism and empiricism go together, and the rejection of Hume's Thesis leads to idealism. More precisely, the rejection of Hume's Thesis leads to idealism given certain explanatory ambitions. We might, of course, have no interest in trying to explain what Kant wants to explain but there is a strong prima facie case for thinking that only idealism can explain the capacity of construction in pure intuition to deliver a priori knowledge of reality. If we are no longer persuaded by this argument it isn't because it isn't successful on its own terms but because we might not share Kant's conviction that we can discover the geometry of physical space by drawing and manipulating figures in the imagination.

In my discussion I have focused on reflection, reasoning, and calculation as sources of a priori knowledge rather than on construction in pure intuition. Nevertheless, the question which Kant raises about construction in pure intuition might also be raised about reflection, reasoning, and calculation: what makes it possible for these sources of knowledge to provide us with a priori knowledge of reality?¹⁶ Asking this question doesn't commit one to the view that it is a synthetic truth that $68 + 57 = 125$, or that nothing can be red all over and green all over at the same time. All one needs is the assumption that these are 'factual' propositions. Given that this is so, it looks as though Kant's explanatory question still arises. For example, if it is a fact about mathematical reality that $68 + 57 = 125$, how is it possible to acquire knowledge of this fact by doing calculations in one's head? And if it is a fact about these colours themselves that nothing can be red all over and green all over at the same time, how is it possible to acquire knowledge of this fact by reflecting on what it would be correct to judge in various circumstances? The problem, one might think, is that there is a kind of mismatch in these cases between what is known and the means by which it is known; paraphrasing Dummett, if numbers and colours are features of objective reality independent of us, it is hard to see why the exercise of intellectual abilities such as calculation and reflection should offer us any sure guide to their structure or relations.

This is the point at which some form of conventionalism might begin to look like an attractive option. In its most extreme form, conventionalism says that necessary truths like ' $68 + 57 = 125$ ' and

¹⁶ Peacocke is exercised by something like this question in his recent work on a priori knowledge. See, for example, 2000: 257.

'nothing can be red all over and green all over at the same time' are direct registers of linguistic conventions that we have adopted, and that our recognition of their truth or necessity is therefore 'a particular case of our knowledge of our own intentions' (Dummett 1966: 424). If our own intentions are knowable a priori then so is the truth of statements which are 'true by convention'. This is just an application of Kant's idealist dictum that 'we can know *a priori* of things only what we ourselves put into them' (Bxviii). Although the conventionalist doesn't regard what we can know by means of reflection or calculation as narrowly mind-dependent facts, the thesis that what we can know by these means are truths which obtain in virtue of human conventions is still a recognizably idealist thesis. According to the conventionalist, this thesis is required in order to explain the capacity of reflection and calculation to deliver a priori knowledge.

This argument for conventionalism has none of the virtues of Kant's explanatory argument for transcendental idealism. When Kant asks how construction in pure intuition can tell us anything about the geometry of physical space he is asking what seems to be a genuine question, and his own answer to this question closes what looks like a genuine explanatory gap. In contrast, why should the idea that it is an objective truth that $68 + 57 = 125$ make it difficult to understand how calculating can be a way of coming to know that $68 + 57 = 125$? Or why should the fact that colours are features of objective reality independent of us make it hard to understand how reflection can put us in a position to know that nothing can be red all over and green all over at the same time? There is no obvious mismatch in either case between source and subject matter, as there is between what Kant regards as the source of geometrical knowledge and what the realist regards as its subject matter. Even if one thinks that numbers are abstract objects, why shouldn't calculating be a way of coming to know arithmetical truths? In this case, the source and the subject matter of our knowledge appear to be made for each other; at any rate, there is no obvious reason for thinking that calculating couldn't be a way of coming to know that $68 + 57 = 125$, or that reflecting couldn't be a way of coming to know that nothing can be red all over and green all over at the same time.

If it is unclear that there is a genuine mismatch between the three sources of a priori knowledge which I have been discussing in this chapter and their subject matter it is even less clear how the conventionalist's solution to this alleged problem is supposed to work. The question was: how can doing a calculation in one's head be a way of coming to know

that $68 + 57 = 125$? The suggested answer is: because the statement ' $68 + 57 = 125$ ' is a register of a linguistic convention. But how is the fact that a statement is a register of a linguistic convention supposed to explain the capacity of mental arithmetic to deliver knowledge of its truth? The problem is that there is no intelligible connection between the *explanandum* and the proposed *explanans*. Perhaps it would be a different matter if the conventionalist were saying that working out in one's head that $68 + 57 = 125$ makes it true that $68 + 57 = 125$, so that the ground of our knowledge is also the ground for the truth of the proposition known. But since this is presumably not what the conventionalist is suggesting we are no closer to understanding how mental arithmetic can be a source of arithmetical knowledge. Indeed, conventionalism makes the situation worse not better. Given that the point of describing a statement as owing its truth to linguistic convention is usually to suggest that the statement isn't factual, there is no question of conventionalism being in the business of explaining how the fact that $68 + 57 = 125$ can be known by calculating. For the conventionalist, there is no such 'fact'.

The point of this discussion of idealism and conventionalism was to explain why it is a mistake to think that the rejection of Hume's Thesis leads to some form of idealism. One can reject Hume's Thesis without being an idealist because one can think that reflection, reasoning, and calculation are sources of a priori knowledge without thinking that they only provide us with a priori knowledge of aspects of reality for which 'we' are in some sense responsible. The thesis that it is a feature of objective reality independent of us that nothing can be red all over and green all over at the same time doesn't deprive us of the right to regard reflecting as a non-experiential way of coming to know that nothing can be red all over and green all over at the same time, any more than the thesis that it is a feature of mathematical reality that $68 + 57 = 125$ deprives us of the right to regard calculating as a non-experiential way of coming to know that $68 + 57 = 125$. There is therefore nothing wrong with a position which rejects Hume's Thesis without abandoning realism; optimism is still in good shape.

Where does this leave the Level 3 project of explaining what makes it possible for reflection, reasoning, and calculation to be sources of a priori knowledge? The answer to this question depends on why we think that an explanation is called for. If we think that one is called for because there is a mismatch between these supposed sources of a priori knowledge and their subject matter then we would effectively

be accepting that there is no difference between a Level 2 response to (HP_{apk}) and a Level 3 response. Explaining what makes it possible to acquire a priori knowledge by reflecting, reasoning, or calculating would be an exercise in obstacle-removal because a mismatch between source and subject matter is just an obstacle by another name. In that case, it wouldn't make sense to insist that there is a distinctive Level 3 perspective on (HP_{apk}) . Instead, we would be forced to concede that explaining what makes it possible to acquire a priori knowledge by certain specified means is the same as showing that there is nothing which makes this impossible.

But what if, as I have argued, there is no mismatch between source and subject matter, and therefore no reason to think that calculating isn't the kind of thing that is capable of revealing that $68 + 57 = 125$, or that reflecting isn't the kind of thing that is capable of revealing that nothing can be red all over and green all over at the same time? That shouldn't prevent us from asking what makes it possible to acquire a priori knowledge by reflecting or calculating. It's just that this question should no longer be seen as an invitation to overcome an intuitive obstacle to the acquisition of a priori knowledge by these means. Instead, since obstacle-removal is something that is supposed to have happened at Level 2, we are now free to think of ourselves as asking an independently intelligible *explanation-seeking* question rather than another obstacle-dependent question. What we should now be looking for, in other words, is a positive Level 3 explanation of the fact that it is possible to acquire a priori knowledge by reflecting, reasoning, or calculating, rather than a demonstration that it is not impossible to acquire a priori knowledge by engaging in such activities.

What would such an explanation look like? Suppose that we think of reflecting, reasoning, and calculating as cognitive activities. One question is: what are the background necessary conditions for the occurrence of these activities? Another question is: what are the background necessary conditions for these activities to be a source of a priori knowledge? Conditions of the first kind might be called 'type A enabling conditions' while conditions of the second kind might be called 'type B enabling conditions'. In these terms, one might think of the challenge of explaining what makes it possible to acquire a priori knowledge by engaging in certain specific cognitive activities as the challenge of identifying type A and type B enabling conditions for these cognitive activities. It might turn out in the end that there is no difference between type A and type B conditions but the basic idea is still that one explains

what makes something possible by identifying the background necessary conditions under which it is possible.

On the face of it, there are many such conditions. Perhaps one needs a specific physical or psychological constitution to be capable of reflecting or reasoning or calculating, and science might have something to tell us about such hardware requirements. But what can philosophy tell us about the conditions under which it is possible to come to know things by the various means which I have been discussing? Consider the following thought: even if I have a range of colour concepts, including the concepts *red* and *green*, I can't know whether the cup in front of me is red or green or some other colour just on the basis of my grasp of such concepts. Something more is needed, such as perception or the word of others. So, for example, I can come to know that the cup is green by seeing that it is green or being told that it is green. In contrast, to know that the cup can't be red all over and green all over at the same time all I need is a grasp of the concepts *red* and *green*, along with some elementary spatial and temporal concepts. As long as I have these concepts I can work out just by thinking about it, by reflecting, that nothing can be red all over and green all over at the same time. The resulting knowledge is an example of what Peacocke calls 'understanding-based *a priori* knowledge' (2000: 257) since it is grounded in, and made available by, my grasp of the relevant concepts. The same goes for my knowledge that $68 + 57 = 125$. What makes this a case of understanding-based knowledge is the fact that it is ultimately grounded in my grasp of the concept of addition; this is what makes it possible for me to come to know that $68 + 57 = 125$ by adding.

Understanding, or having a grasp of the relevant concepts, is both a type A and a type B enabling condition for the acquisition of a priori knowledge by reflection or calculation. It is a background necessary condition for the occurrence of such activities as reflecting or calculating, as well as a background necessary condition for these activities to be sources of a priori knowledge. What is more, these are a priori conditions since they have been brought to light by means of a priori philosophical reflection. In this sense, it is reflection itself which reveals the conditions under which reflection can be a source of a priori knowledge. But now there is a further question which might be raised. Having accepted that a priori knowledge that is arrived at by reflecting or calculating is understanding-based a priori knowledge we might go on to ask 'what it is about understanding that makes *a priori* knowledge possible' (Peacocke 2000: 257). The thought here is that it

isn't enough, for the purposes of providing a fully-fledged multi-levels response to (HP_{apk}), to point out that concept possession serves as a Level 3 enabling condition for the acquisition of a priori knowledge by the means identified at Level 1. It also needs to be explained *how* concept possession can play this role, that is, *what makes it possible* for the understanding to ground the acquisition of a *a priori* knowledge.

What should we make of this explanatory demand? This question reopens the debate between minimalism and anti-minimalism which first came up in Chapter 1. In the present context, a minimalist is someone who thinks that it is neither necessary nor possible to give a substantive account of what it is about understanding that makes a priori knowledge possible.¹⁷ It isn't necessary because (HP_{apk}) has already been answered once means of acquiring a priori knowledge have been identified and obstacles to the acquisition of a priori knowledge by the suggested means have been overcome or dissipated. And it isn't possible because it doesn't make sense to try to identify enabling conditions for enabling conditions. Indeed, from a minimalist perspective, any attempt to identify enabling conditions for the acquisition of a priori knowledge by reflection or calculation is already a step too far. It only makes matters worse to then seek to explain what makes it possible for understanding to make a priori knowledge possible.

The most effective way of undermining minimalism is to come up with just the kind of explanation which it claims to be impossible. So, for example, one might point out that to know that nothing can be red all over and green all over at the same time is to know something about what Laurence Bonjour calls 'the character of the extra-conceptual world' (1998: 18), and that understanding-based reflection can only put one in a position to know this fact about the extra-conceptual world because of the way that our colour concepts are tied to the actual colours of which they are concepts. In other words, it is only because the concept *red* is tied to the individuation of the colour red, and the concept *green* is tied to the individuation of the colour green, that reflection can yield the understanding-based a priori knowledge that nothing can be red all over and green all over at the same time. On this account, which is basically Peacocke's account, what it is about understanding that makes

¹⁷ Peacocke also talks about minimalism in his discussion. A minimalist in his sense 'regards the resource of what is primitively written into the identity of a concept as already a full explanation of the relation between the a priori and the concepts featuring on the content of a priori knowledge' (2000: 259).

understanding-based a priori knowledge of colour incompatibilities possible is the fact that one needs colour concepts in order to understand propositions about colour incompatibilities, and that colour *concepts* are closely related to colour *properties*.¹⁸

The basic idea, then, is that thought can tell us about reality not because, as idealism claims, reality is mind-dependent but because many of the concepts which are exercised in thought are world-dependent. This is externalism not idealism, and it adds weight to the suggestion that a substantive explanation of the possibility of a priori knowledge needn't be an idealist explanation. As long as the suggested account of the relation between colour concepts and colour properties works and has wider application it shows that minimalism is not the way to go. That leaves anti-minimalism. Extreme anti-minimalism is the view that we haven't fully answered (HP_{apk}) until we have identified background necessary conditions for the acquisition of a priori knowledge and explained what makes it possible for understanding to be such a condition. In contrast, moderate anti-minimalism is the view that (HP_{apk}) can be satisfactorily answered without going into this. As far as moderate anti-minimalism is concerned, questions about enabling conditions are ones to which substantive answers *can* be given but to which substantive answers *needn't* be given if what we are ultimately interested in doing is explaining how a priori knowledge is possible.

This seems right. It's true that there is no clear sense in which someone who has both identified means of coming to know things a priori and dealt with obstacles to the acquisition of a priori knowledge by the suggested means has failed to explain how a priori knowledge is possible. That is why moderate anti-minimalism is right to represent Level 3 explanations as optional in relation to (HP_{apk}). Explaining how a priori knowledge is possible is, first and foremost, a matter of identifying effective means by which such knowledge can be acquired, and the fact that one hasn't given a substantive account of what makes it possible to acquire a priori knowledge by such means needn't call their efficacy as means into question. In contrast, a failure to deal with obstacles to the acquisition of a priori by means of reflection, reasoning, or calculation would call their efficacy as means into question and so would amount to a failure to answer (HP_{apk}). This is the basis of the thought that a Level 2 response to (HP_{apk})

¹⁸ See Peacocke 2000 for much more on this.

is necessary in a way that a Level 3 response is not. But if there is this difference in importance between what happens at Level 2 and at Level 3 of a multiple levels response to (HP_{apk}) then that strengthens the case for treating these levels as distinct from one another.

To conclude, the conception of a priori knowledge which I have been recommending is, in a way, much more optimistic than Kant's conception. Optimism in the narrow sense is simply the view that a priori knowledge of matters of fact is possible. Kant is an optimist in this sense, since he rejects Hume's Thesis, but his optimism is a qualified optimism. As we have seen, he thinks that his rejection of Hume's Thesis commits him to idealism because he thinks that there is an important question about our a priori knowledge which only idealism can answer. My optimism is not qualified in this way. Although I have argued against Hume's Thesis on the basis that reflection, reasoning, and calculation can provide us with a priori knowledge of matters of fact I have done so within a realist framework. There is no need to abandon this framework because Kant's question is less important in relation to (HP_{apk}) than he thinks it is, and because realism does have an answer to this question; it can explain what makes it possible for the various sources which it identifies to be sources of a priori knowledge, and its explanation doesn't require the positing of an intrinsic harmony between mind and world. What it requires is a commitment to an independently plausible form of externalism.

Does this make my unqualified optimism a form of rationalism? It might seem that it does because rationalists are also unqualified optimists. They explain how a priori knowledge of matters of fact is possible by specifying means by which it is possible but deny that they are committed to idealism. Their suggestion is that *rational intuition* is the basic source of a priori knowledge, and that the resulting knowledge is knowledge of the layout of extra-conceptual, non-mental reality. But what is the difference between the proposal that rational intuition is the basic source of a priori knowledge and the proposal that reflection, reasoning, and calculation are all sources of a priori knowledge? The answer is that there are several differences between these proposals. I want to bring this chapter to an end by exploring some of these differences and explaining why it is unhelpful and potentially misleading to think of the multi-levels approach to (HP_{apk}) as a form of rationalism.

6.4 RATIONALISM

The basis of rationalism's identification of rational intuition or rational insight as the ultimate source of a priori knowledge is its conception of a priori knowledge as requiring a priori epistemic justification. According to one contemporary rationalist, a priori epistemic justification is a form of non-experiential justification which 'occurs when the mind directly or intuitively sees or grasps or apprehends (or perhaps merely seems to itself to see or grasp or apprehend) a necessary fact about the nature or structure of reality' (BonJour 1998: 15–16). Consider, once again, the claim that nothing can be red all over and green all over at the same time. Given that I understand the ingredients of this proposition,

I am able to see or grasp or apprehend in a seemingly direct and unmediated way that the claim in question cannot fail to be true—that the natures of redness and greenness are such as to preclude their being jointly realized. It is this direct insight into the necessity of the claim that seems, at least *prima facie*, to justify my accepting it as true. (BonJour 1998: 101)

'Direct' in this context means 'non-discursive'.¹⁹ When one 'sees', in the relevant sense, that a proposition must be true one's apprehension isn't mediated by a series of steps. While this doesn't make rational intuition or rational insight any less 'rational' it does imply its irreducibility to discursive reasoning. On this account, rational insight is a *sui generis* source of a priori justification, and therefore a *sui generis* source of a priori knowledge.

The emphasis on the directness of rational insight makes sense of the perceptual metaphor. When I see, in the ordinary sense of 'see', that the cup is chipped I don't reason my way to the conclusion that the cup is chipped; I just see that it is chipped. No conscious inference or reasoning is involved, which is why the resulting perceptual knowledge is non-inferential. But what about reflecting, reasoning, and calculating? Are these non-discursive sources of knowledge in the way that, by analogy with ordinary seeing, rational insight is supposed to be a non-discursive source of knowledge? It seems not. Even in the case of moderately complex arithmetical propositions, talk of 'just seeing' is out of place. I don't 'just see' that $68 + 57 = 125$, I have to work it out. Since this means that conscious reasoning is involved, I don't have

¹⁹ This is clear from BonJour's discussion. See, for example, BonJour 1998: 104.

any rational insight into the truth or necessity of the proposition if such insights are supposed to be 'direct'. The same goes for the claim that nothing can be red all over and green all over at the same time. If this is something which I have to work out then I do not see in a 'direct and unmediated way' that the claim cannot fail to be true.

A further consideration is this: rational insight is supposed to be intuitive insight into *necessity*, the idea being that if I am justified in believing that a certain proposition *must be* true then I am justified in believing that it *is* true. One potentially unwelcome consequence of rationalism is that it leaves no room for contingent a priori knowledge. Assuming that 'seeing that a proposition P is necessarily true' is factive, I can't see that P is necessarily true unless it is necessarily true. This won't bother rationalists who think that there is no such thing as contingent a priori knowledge but the problem goes deeper than that. Take the case in which I calculate that $68 + 57 = 125$. Let's agree that coming to know that $68 + 57 = 125$ is coming to know a necessary truth. Does it follow that, when I come to know that $68 + 57 = 125$ by working it out, I come to know *that* $68 + 57 = 125$ is a necessary truth? No. When I add 68 to 57 what I discover is that their sum *is* 125. While I might become convinced, on further reflection, that this is a necessary truth, this is not something which I need to know or believe in order to know that $68 + 57 = 125$. Yet, as long as I know that $68 + 57 = 125$ by calculating, the resulting knowledge is a priori. This is an important difference between calculation and so-called 'rational intuition' as non-experiential sources of knowledge. Calculating can tell me that something is so without telling me that it must be so. Rationally intuiting can only tell me that something is so by telling me that it must be so.

What about coming to know by means of a priori reflection that nothing can be red all over and green all over at the same time? Here it might seem that talk of an intuitive insight into necessity or impossibility is more appropriate since the proposition known has modal content. It's worth pointing out, however, that it can be true that nothing can be both f and g even if it isn't a necessary truth that nothing can be both f and g. For example, it is true, but not necessarily true, that it isn't possible for someone to be a member of the British House of Commons and the House of Lords at the same time. The incompatibility of red and green is different because it is not just true but necessarily true that nothing can be red all over and green all over at the same time. Nevertheless, it is a further question whether coming to know by a

priori reflection that nothing can be red all over and green all over at the same time is coming to know *that* this is necessarily true. If not, then reflecting is like calculating in this respect; it needn't provide one with modal knowledge in the way that rational intuition is supposed to provide one with modal knowledge.

This is not to say that there aren't other differences between reflecting and calculating as sources of a priori knowledge. Both are forms of reasoning but calculating is a more specialized form of reasoning than reflecting. In addition, not all reasoning is calculating or reflecting. I don't *calculate* that Tony Blair lives in Downing Street on the basis that he is the Prime Minister and that the Prime Minister lives in Downing Street. If reflection is involved in this kind of reasoning it is not the same as the understanding-based reflection by means of which one comes to know that nothing can be red all over and green all over at the same time. This points to another important difference between the multi-levels approach and rationalism. The multi-levels approach emphasizes the multiplicity and variety of sources of a priori knowledge. Rationalism, at least in its simplest form, is more one-dimensional; it wants to represent a priori knowledge as having just one basic source.

These differences between the multi-levels approach and rationalism have a bearing on the standard objection to rationalism. The basic problem with rationalism, it is often claimed, is that it posits 'a special evidence-gathering faculty of *intuition*, distinct from the standard five senses, which allows us to arrive at justified beliefs about the necessary properties of the world' (Boghossian 1997: 334). This is a problem because this faculty 'has never been described in satisfactory terms' (ibid.). The multi-levels approach doesn't face this objection because what it posits is a variety of different non-experiential ways or means of coming to know rather than a single, evidence-gathering faculty. While there is more to be said about each of these means, the issue is not that reflecting, reasoning, and calculating 'have never been described in satisfactory terms'. We might, for example, ask what makes it possible to know that $68 + 57 = 125$ by calculating but the suggestion that calculating *is* a way of coming to know that $68 + 57 = 125$ lacks even the superficial oddity of the suggestion that rationally intuiting the necessity of the proposition is a way of coming to know that $68 + 57 = 125$.

It should be clear by now that the multi-levels account of a priori knowledge isn't a straightforward form of rationalism. But it's not a form of empiricism either, at least if empiricism is defined by a commitment

to Hume's Thesis. In fact, the best policy from a multi-levels perspective is simply to avoid being drawn into the somewhat sterile debate between rationalism and empiricism. What matters in the end is not whether one classifies oneself as a 'rationalist' or 'empiricist' but whether one is optimistic or pessimistic about the possibility of a priori knowledge. If one thinks that a priori knowledge is possible, the next question is: how is it possible? Saying that it is possible by means of rational intuition doesn't get one very far, and makes (HP_{apk}) look far more difficult than it actually is. If what I have been arguing in this chapter is correct, the right answer to (HP_{apk}) is much less exotic than the rationalist's answer: a priori knowledge is, just like empirical knowledge, possible by many different means, including reflection, reasoning, and calculation. And what makes reflection, reasoning, and calculation sources of a priori knowledge is simply the fact that they are a priori sources of knowledge, that is, non-experiential ways of coming to know that something is the case.

This page intentionally left blank

Bibliography

- Allison, H. (1983), *Kant's Transcendental Idealism* (New Haven: Yale University Press).
- Ameriks, K. (2003), 'Kant's Transcendental Deduction as a Regressive Argument', in K. Ameriks, *Interpreting Kant's Critiques* (Oxford: Oxford University Press): 51–66.
- Austin, J. L. (1979), 'Other Minds', in J. L. Austin, *Philosophical Papers*, 3rd edn., ed. J. O. Urmson and G. J. Warnock (Oxford: Oxford University Press): 76–116.
- Avramides, A. (2001), *Other Minds* (London: Routledge).
- Ayer, A. J. (1946), *Language, Truth and Logic*, 2nd edn. (Harmondsworth: Penguin).
- Ayers, M. R. (1974), 'Individuals without Sortals', *Canadian Journal of Philosophy*, 4: 113–48.
- Bennett, J. (1966), *Kant's Analytic* (Cambridge: Cambridge University Press).
- Baldwin, T. (1990), *G. E. Moore* (London, Routledge).
- Boghossian, P. (1997), 'Analyticity', in B. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language* (Oxford: Blackwell): 331–68.
- (2003), 'Blind Reasoning', *Aristotelian Society*, supp. vol. 77: 225–48.
- BonJour, L. (1998), *In Defence of Pure Reason* (Cambridge: Cambridge University Press).
- Burge, T. (1993), 'Content Preservation', *Philosophical Review*, 102: 457–88.
- (1998), 'Computer Proof, Apriori Knowledge, and Other Minds', in J. Tomberlin (ed.), *Philosophical Perspectives*, 12 (Malden, Mass.: Blackwell): 1–37.
- Campbell, J. (1997), 'Sense, Reference and Selective Attention', *Aristotelian Society*, supp. vol. 71: 55–74.
- (2002), *Reference and Consciousness* (Oxford: Clarendon Press).
- Cassam, Q. (1998), 'Mind, Knowledge and Reality: Themes from Kant', in A. O'Hear (ed.), *Current Issues in Philosophy of Mind* (Cambridge: Cambridge University Press): 321–48.
- (1999), 'Self-Directed Transcendental Arguments', in R. Stern (ed.), *Transcendental Arguments: Problems and Prospects* (Oxford: Oxford University Press): 83–110.

- Cassam, Q. (2004), 'Introspection, Perception, and Epistemic Privilege', *The Monist*, 87: 255–74.
- (2005), 'Space and Objective Experience', in J. Bermúdez (ed.), *Thought, Reference, and Experience: Themes from the Philosophy of Gareth Evans* (Oxford: Clarendon Press): 258–89.
- Cohen, S. (2002), 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research*, 65: 309–30.
- Collins, A. (1999), *Possible Experience: Understanding Kant's Critique of Pure Reason* (Berkeley and Los Angeles: University of California Press).
- Craig, E. J. (1987), *The Mind of God and the Works of Man* (Oxford: Oxford University Press).
- Dray, W. (1957), *Laws and Explanation in History* (Oxford: Oxford University Press).
- Dretske, F. (1969), *Seeing and Knowing* (London: Routledge & Kegan Paul).
- (2000), 'Simple Seeing', in F. Dretske, *Perception, Knowledge and Belief* (Cambridge: Cambridge University Press): 97–112.
- Dummett, M. (1966), 'Wittgenstein's Philosophy of Mathematics', in G. Pitcher (ed.), *Wittgenstein: The Philosophical Investigations* (London: Macmillan): 420–47.
- (1973), *Frege: Philosophy of Language* (London: Duckworth).
- (1991), 'Frege and Kant on Geometry', in M. Dummett, *Frege and Other Philosophers* (Oxford: Clarendon Press): 126–57.
- Evans, G. (1980), 'Things without the Mind: A Commentary upon Chapter Two of Strawson's *Individuals*', in Z. Van Straaten (ed.), *Philosophical Subjects: Essays Presented to P. F. Strawson* (Oxford: Clarendon Press): 76–116.
- (1982), *The Varieties of Reference*, ed. J. McDowell (Oxford: Oxford University Press).
- Falkenstein, L. (1995), *Kant's Intuitionism: A Commentary on the Transcendental Aesthetic* (Toronto: University of Toronto Press).
- Fodor, J. (1981), 'The Present Status of the Innateness Controversy', in J. Fodor, *Representations: Philosophical Essays on the Foundations of Cognitive Science* (Brighton: Harvester Press): 257–316.
- Frege, G. (1967), 'The Thought: A Logical Inquiry', in P. F. Strawson (ed.), *Philosophical Logic* (Oxford: Oxford University Press): 17–38.
- Friedman, M. (1991), 'Regulative and Constitutive', *Southern Journal of Philosophy*, 30: 73–102.
- Giaquinto, M. (1996), 'Non-Analytic Conceptual Knowledge', *Mind*, 105: 249–68.
- Goldman, A. (1986), *Epistemology and Cognition* (Cambridge, Mass.: Harvard University Press).
- (2002), *Pathways to Knowledge: Private and Public* (Oxford: Oxford University Press).

- Guyer, P. (1987), *Kant and the Claims of Knowledge* (Cambridge: Cambridge University Press).
- (1992), 'The Transcendental Deduction of the Categories', in P. Guyer (ed.), *The Cambridge Companion to Kant* (Cambridge: Cambridge University Press): 123–60.
- Hampshire, S. (1979), 'Some Difficulties in Knowing', in T. Honderich and M. Burnyeat (eds.), *Philosophy as it is* (Harmondsworth: Penguin): 281–308.
- Hatfield, G. (1990), *The Natural and the Normative: Theories of Spatial Perception from Kant to Helmholtz* (Cambridge, Mass.: MIT Press).
- Henrich, D. (1989), 'Kant's Notion of a Deduction and the Methodological Background of the First Critique', in E. Förster (ed.), *Kant's Transcendental Deductions* (Stanford, Calif.: Stanford University Press): 29–46.
- Hill, C. (1991), *Sensations: A Defense of Type Materialism* (Cambridge: Cambridge University Press).
- Hume, D. (1975), 'An Enquiry Concerning Human Understanding', in L. A. Selby-Bigge (ed.), *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, 3rd edn., rev. P. H. Nidditch (Oxford: Clarendon Press).
- James, W. (1884), 'What is an Emotion?', *Mind*, 9: 188–205.
- Kant, I. (1932), *Critique of Pure Reason*, trans. N. Kemp Smith (London: Macmillan).
- (1977), *Prolegomena to Any Future Metaphysics*, original trans. P. Carus, rev. J. Ellington (Indianapolis: Hackett Publishing Company).
- (1992), 'The Jäsche Logic', in I. Kant, *Lectures on Logic*, trans. J. Michael Young (Cambridge: Cambridge University Press): 517–640.
- Kitcher, P. (1982), 'How Kant Almost Wrote "Two Dogmas of Empiricism" (And Why He Didn't)', in J. N. Mohanty and Robert W. Shahan (eds.), *Essays on Kant's Critique of Pure Reason* (Norman: University of Oklahoma Press): 217–49.
- Kornblith, H. (2002), *Knowledge and its Place in Nature* (Oxford: Clarendon Press).
- Locke, J. (1975), *An Essay Concerning Human Understanding*, ed. P. H. Nidditch (Oxford: Clarendon Press).
- Longuenesse, B. (1998), *Kant and the Capacity to Judge: Sensibility and Discursivity in the Transcendental Analytic of the Critique of Pure Reason*, trans. C. T. Wolfe (Princeton: Princeton University Press).
- McDowell, J. (1998a), *Meaning, Knowledge, and Reality* (Cambridge, Mass.: Harvard University Press).
- (1998b), 'Singular Thought and the Extent of Inner Space', in McDowell (1998a): 228–59.
- (1998c), 'Criteria, Defeasibility, and Knowledge', in McDowell (1998a): 369–94.

- McDowell, J. (1998d), 'Having the World in View: Sellars, Kant, and Intentionality', *Journal of Philosophy*, 95: 431–91.
- McGinn, C. (1984), 'What is the Problem of Other Minds?', *Aristotelian Society*, supp. vol. 58: 119–37.
- (2004), 'Inverted First-Person Authority', *The Monist*, 87: 237–54.
- Mackie, J. L. (1976), *Problems from Locke* (Oxford: Clarendon Press).
- Markie, P. (2004), 'Rationalism vs. Empiricism', in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2004 edition), URL: <http://plato.stanford.edu/archives/fall2004/entries/rationalism-empiricism>.
- Mill, J. S. (1891), *A System of Logic* (London: Longmans, Green and Co.).
- Moore, G. E. (1922), *Philosophical Studies* (London: Routledge).
- (1953), *Some Main Problems of Philosophy* (London: Allen & Unwin).
- Nozick, R. (1981), *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press).
- Nudds, M. (2001), 'Experiencing the Production of Sounds', *European Journal of Philosophy*, 9: 210–29.
- Peacocke, C. (1992), 'Scenarios, Concepts and Perception', in T. Crane (ed.), *The Contents of Experience: Essays on Perception* (Cambridge: Cambridge University Press): 105–35.
- (1993), 'Intuitive Mechanics, Psychological Reality, and the Idea of a Material Object', in N. Eilan, R. McCarthy, and B. Brewer (eds.), *Spatial Representation* (Oxford: Basil Blackwell): 162–76.
- (1999), *Being Known* (Oxford: Oxford University Press).
- (2000), 'Explaining the A Priori: The Programme of Moderate Rationalism', in P. Boghossian and C. Peacocke (eds.), *New Essays on the A Priori* (Oxford: Clarendon Press): 255–85.
- Pickard, H. (2003), 'Emotions and the Problem of Other Minds', in A. Hatzimoysis (ed.), *Philosophy and the Emotions* (Cambridge: Cambridge University Press): 87–104.
- Prior, A. N. (1949), 'Determinables, Determinates, and Determinants', *Mind*, 57: 1–20.
- Putnam, H. (1975), 'Brains and Behaviour', in H. Putnam, *Mind, Language and Reality: Philosophical Papers*, vol. ii (Cambridge: Cambridge University Press): 325–41.
- Quine, W. V. (1953), 'Two Dogmas of Empiricism', in W. V. Quine, *From a Logical Point of View* (Cambridge, Mass.: Harvard University Press): 20–46.
- (1969), 'Epistemology Naturalized', in W. V. Quine, *Ontological Relativity and Other Essays* (New York: Columbia University Press): 69–90.
- (1970), *Philosophy of Logic* (Englewood Cliffs, N. J.: Prentice Hall).
- (1976), 'Carnap and Logical Truth', in W. V. Quine, *The Ways of Paradox* (Cambridge, Mass.: Harvard University Press): 107–32.
- Ryle, G. (1949), *The Concept of Mind* (London: Hutchinson).

- Searle, J. (1983), *Intentionality: An Essay in the Philosophy of Mind* (Cambridge: Cambridge University Press).
- Sellars, W. (1956), 'Empiricism and the Philosophy of Mind', in H. Feigl and M. Scriven (eds.), *Minnesota Studies in the Philosophy of Science*, vol. i (Minneapolis: Minnesota University Press): 253–329.
- Spelke, E., and Van de Walle, G. (1993), 'Perceiving and Reasoning about Objects: Insights from Infants', in N. Eilan, R. McCarthy, and B. Brewer (eds.), *Spatial Representation* (Oxford: Basil Blackwell): 132–61.
- Stern, R. (2000), *Transcendental Arguments and Scepticism: Answering the Question of Justification* (Oxford: Clarendon Press).
- Strawson, P. F. (1959), *Individuals: An Essay in Descriptive Metaphysics* (London: Methuen).
- (1966), *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason* (London: Methuen).
- (1997a), *Entity and Identity and Other Essays* (Oxford: Oxford University Press).
- (1997b), 'Entity and Identity', in Strawson (1997a): 21–51.
- (1997c), 'Kant's New Foundations of Metaphysics', in Strawson (1997a): 232–43.
- (1997d), 'The Problem of Realism and the A Priori', in Strawson (1997a): 244–51.
- Stroud, B. (1984), *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press).
- (2000a), *Understanding Human Knowledge* (Oxford: Oxford University Press).
- (2000b), 'Scepticism and the Possibility of Knowledge', in Stroud (2000a): 1–8.
- (2000c), 'Transcendental Arguments', in Stroud (2000a): 9–25.
- (2000d), 'Understanding Human Knowledge in General', in Stroud (2000a): 99–121.
- (2000e), 'Epistemological Reflection on Knowledge of the External World', in Stroud (2000a): 122–38.
- (2004), 'Perceptual Knowledge and Epistemological Satisfaction', in J. Greco (ed.), *Ernest Sosa and his Critics* (Oxford: Blackwell Publishing): 165–73.
- Unger, P. (1971), 'A Defense of Scepticism', *Philosophical Review*, 80: 198–219.
- Warren, D. (1998), 'Kant and the Apriority of Space', *Philosophical Review*, 107: 179–224.
- Wiggins, D. (1980), *Sameness and Substance* (Oxford: Blackwell).
- (1997), 'Sortal Concepts: A Reply to Xu', *Mind and Language*, 12: 413–21.
- (2001), *Sameness and Substance Renewed* (Cambridge: Cambridge University Press).

- Williams, M. (1996), *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism* (Princeton: Princeton University Press).
- Williamson, T. (2000), *Knowledge and its Limits* (Oxford: Oxford University Press).
- (2003), 'Understanding and Inference', *Aristotelian Society*, supp. vol. 77: 249–93.
- Wittgenstein, L. (1978), *Philosophical Investigations*, trans. G. E. M. Anscombe (Oxford: Basil Blackwell).

Index

- Acquisition Argument 141–8
Allison, H. 40 n., 55 n., 89 n., 96, 107
Ameriks, K. 57
analyticity 11–12, 206–7
anti-minimalism, *see* minimalism
a priori knowledge 22–4, 188–223
 Kant on 10–22, 188–91, 210–12
 obstacles to 191–6, 199–210
 minimalism about 214–19
assertibility-conditions 179–82
auditory perception 100–6, 111–15,
 116–22
Austin, J. L. 5, 29–30, 82, 156, 165
Avramides, A. 156 n., 158 n.
Ayer, A. J. 192 n., 193 n., 206
- background condition 17–18
 see also enabling conditions; necessary
 conditions
Baldwin, T. 33 n.
basic knowledge 202–4
Bennett, J. 141 n., 144 n., 147 n.
Berkeley, G. 193 n., 210
blindsight 202–3
bodily subjects 175–8, 182
 see also subjects, spatial
Boghossian, P. 73 n., 189 n., 191, 193,
 207 n., 222
BonJour, L. 217, 220
Burge, T. 17–18, 189
- calculation 195–210, 212–23
Campbell, J. 86 n., 103, 120, 135
Cassam, Q. 70 n., 78 n., 108 n., 157 n.
Categorical Thinking
 Requirement 88–9, 129–54
 necessity of 131–52
 foundations of 152–3
 a priority of 153
 minimalism about 153–4
categories 43, 67–9, 70–80, 131
clairvoyance 83, 202–3
Cohen, S. 201–2
Collins, A. 51 n., 53
concept acquisition, Kant's views
 on 43, 142–6
- concepts
 abstractionism about 142–8
 acquisition of, *see* Acquisition
 Argument
 categorical, *see* categories
 characterising 130–2, 140–1,
 146–8
 determinable 130–1
 empirical 41–3, 71–7
 holism about 42–3, 138–40
 legitimacy of 70–80
 sortal 130–2, 140–1, 144–8
 see also Categorical Thinking
 Requirement; categories;
 Generality Constraint
conceptual analysis 11–13, 83–4
Conceptual Argument 141, 148–52
conclusive reasons 45
contingent a priori 206–7, 221
conventionalism 212–14
corporeal characteristics 175–6, 182
correctness conditions 111
Craig, E. 210 n.
- differential certainty, argument
 from 33–4
differentiation
 auditory 100–4, 111–14
 object–object 93–107
 schema-based 103–5
 sortal dependency of 133–5
 subject–object 107–16
 visual 39–41
Dray, W. 1 n., 5 n.
dreaming 30–2
Dretske, F. 17 n., 27, 29 n., 36, 40 n.,
 45, 62, 87 n., 92 n., 94–5, 130 n.,
 133, 158 n., 160 n., 161–2, 164 n.
dualism 173–6
Dummett, M. 73 n., 211–13
- empirical knowledge 22–4, 81–4
empiricism 192–5, 209–10, 222–3
enabling conditions 9–10, 16–22,
 86–90, 106, 127–8, 186–7
 strongly vs. weakly a priori 39

- epistemic perception 27–50, 86–90,
140–54
 enabling conditions of 34–50
 primary vs. secondary 36
 see also Categorial Thinking
 Requirement; Spatial
 Perception Requirement
- epistemic perception, argument
 from 140–52
- error, argument from 200–1
- Evans, G. 108–9, 148–9
- evil, problem of 2–3
- experience, Kant on 11–13, 52,
54–5
- explanation 16–22, 38–9, 42, 80–4
 see also minimalism
- existence unperceived 108–9
- externalism 218
- external world, knowledge of 4–10
- factive mental state operators
 (FMSO) 44
- Factuality Thesis 191–3, 206–7
- Falkenstein, L. 124 n.
- Fodor, J. 147 n., 148
- forms of intuition 67–9, 79–80,
123–5, 211–12
- freedom, possibility of 1–3
- Frege, G. 153
- Friedman, M. 75
- Generality Constraint 69 n., 149
- generality, problem of 62–5
- geometry, Kant's views on 10–22,
55–6, 184–5, 210–12
- Giaquinto, M. 208
- Ginsborg, H. 77 n.
- Goldman, A. 5 n., 6 n., 7 n., 83 n., 189
n., 190
- Guyer, P. 90 n., 137 n.
- hallucination 30–2
- Hatfield, G. 51 n., 67
- Henrich, D. 67 n.
- Hill, C. 158 n., 166 n.
- how-possible questions 1–51
 Kant's version of 10–22
 obstacle dependence of 2
 see also Means Response; multi-levels
 response
- how-questions 5–6, 82–3
- Hume's Thesis 191–5, 210–14
- idealism 89–91, 124–5, 210–12,
218–19
 see also transcendental idealism
- Identity Condition 171–2, 178–83
- incompatibility of colours 196–7,
216–18, 221–2
- indefeasibility 207–8
- introspection 189–90, 193–4
- intuitions, Kantian 136–7
- James, W. 162 n.
- Kant, I. 10–15, 18–22, 24–5, 40–3,
46–7, 49, 51–60, 65–79, 87–8,
96, 107, 111, 113, 123–4, 131,
136–9, 141–8, 150, 184–5, 188,
190, 210–13, 219
- Kitcher, P. 74
- Kornblith, H. 84
- KR principle 201–5
- location, *see* Spatial Perception
 Requirement
- Locke, J. 194
- Locke's Principle 126, 197
- Longuenesse, B. 43 n., 141–3
- McDowell, J. 16 n., 29 n., 30–1, 137,
144 n., 151 n., 158 n.
- McGinn, C. 157 n., 181
- Mackie, J. L. 146 n.
- Markie, P. 193n.
- material objects 36–8, 100–12,
119–22, 124–6
- means 13–14, 99–100, 186
 see also generality, problem of;
 necessary conditions
- Means Response 6–7, 83–5
- memory 17–18, 189–90, 204
- Mill, J. S. 158 n.
- minimalism 8–10, 86–90, 126–8,
183–7, 217
 Kant on 15, 19–22, 46–50
 about perceptual knowledge 36–50
- Moorean view 8, 15, 32–4, 169
- Müller-Lyer illusion 98 n., 110–11
- multi-levels response 9–10, 20–1, 51,
186–7
- naturalism 19 n., 38–9, 42, 84
- necessary conditions 13, 17, 38, 52–3,
58–62, 80–4, 181
 see also enabling conditions

- necessary truths 11–12, 206–7
 non-discursive knowledge 220–1
 Nozick, R. 1, 2 n., 9 n.
 Nudds, M. 112 n., 117 n.
- object perception 36–41, 40, 91–2,
 132–40
see also simple seeing
 object perception, argument
 from 132–40
 obstacle dependence 2, 16, 32–4
 other minds, knowledge of 155–87
 inferential model of 158–60
 minimalism about 171, 182–7
 obstacles to 161–71
 perceptual model of 158–72, 180–2
- Peacocke, C. 6 n., 9 n., 23 n., 111, 125
 n., 172, 188 n., 190–1, 196, 197
 n., 212 n., 216–18
- perception
 and inference 159–60
 and judgement 98–9, 109–10
see also perceptual knowledge
 perceptual content 110–11
 perceptual differentiation, *see*
 differentiation
 perceptual knowledge 5–10, 22–50,
 57
 minimalism about 34–50
 obstacles to 22–36
see also epistemic perception
 Pickard, H. 162 n.
 predication, *see* Identity Condition
 primary qualities 120–1
 Primary Qualities Argument 121–3,
 124–6
 Prior, A. N. 131 n.
 Putnam, H. 182
- Quine, W. V. 19 n., 194 n., 207–8
- rational intuition 219–23
 rationalism 9 n., 219–23
 realism 89–90, 124–5, 193–5,
 209–14
 reasoning 195–210, 212–23
 receptivity 108–9
 reflection 195–210, 212–23
 Refutation of Idealism 54–5
 responsibility, epistemic 202–3
 Ryle, G. 179 n.
- scepticism 3, 33–4, 170
see also Moorean view
 Schematism 14–15
 Searle, J. 17 n.
 self-ascription/other-ascription 175–7
see also assertibility-conditions;
 Identity Condition
 self-knowledge 32, 54–5
 Sellars, W. 137 n.
 simple seeing 27, 132–40
 sounds 101–2, 104, 111–14
 spatial perception 90–1
 Kant on 40–1
see also Spatial Perception
 Requirement
 Spatial Perception
 Requirement 88–128
 necessity of 88, 91–3
 direct argument for 116–23
 indirect argument for 93–116
 a priority of 88, 125–6
 foundations of 89, 123–5
 minimalism about 126–8
 spatial properties 90–1
 Spatiality Condition 171–8, 181–3
 Spelke, E. 99 n.
 Stern, R. 54
 Strawson, P. F. 64 n., 69 n., 90 n., 148
 n., 149, 158 n., 172, 174–7, 179
 n.
 Stroud, B. 1 n., 4 n., 5 n., 6, 7 n., 19 n.,
 24–6, 27 n., 30 n., 54 n., 158 n.
 subjects, spatial 172–4
 super-spartans 182
 synthesis 136–7, 142–4
 synthetic a priori knowledge 10–22, 56
- tactile perception 95–6, 117, 120–2
 transcendental arguments 51–84
 anti-sceptical 54–7
 explanatory 80–4
 revelatory 66–71
 validatory 70–80
 how-possible questions, relation
 to 51–62
 multi-levels approach, relation to 55
 regressive 57–62
 self-directed vs. world-directed 70
 Transcendental Deduction of the
 Categories 65–84
 transcendental idealism 18, 79–80,
 123–5, 184–5, 210–12

transcendental ideas 75–7

Unger, P. 4 n.

universality, problem of 14–15

van de Walle, G. 99 n.

visibility 177–8

see also Spatiality Condition

Warren, D. 40 n., 72, 91 n., 97–100

what-is-it questions 81–4

what-makes-it-possible questions 9, 86

see also enabling conditions

Wiggins, D. 93 n., 130 n., 131 n., 133

n., 134, 135 n.

Williams, M. 4 n.

Williamson, T. 27 n., 29 n., 44, 73 n.,

83

Wittgenstein, L. 182

zombie-elimination

requirement 165–70