

The Nature of Consciousness

Mark Rowlands

University College, Cork



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, VIC 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Mark Rowlands, 2001

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2001

Printed in the United Kingdom at the University Press, Cambridge

Typeface Plantin 10/12 pt. *System* L^AT_EX 2_ε [TB]

A catalogue record for this book is available from the British Library.

Library of Congress Cataloguing in Publication Data

Rowlands, Mark.

The Nature of Consciousness / Mark Rowlands.

p. cm.

Includes bibliographical references and index.

ISBN 0 521 80858 8

1. Consciousness. I. Title.

B808.9 .R69 2001

126 – dc21

2001035253

ISBN 0 521 80858 8

Contents

<i>Preface</i>	<i>page ix</i>
1 The problem of phenomenal consciousness	1
1 What is phenomenal consciousness?	1
2 The scope of ‘There is . . .’	4
3 What is the problem of phenomenal consciousness?	8
4 Explaining consciousness	13
5 Vertical strategies I: the mind–body problem	14
6 Vertical Strategies II: the mind–mind problem	16
7 Horizontal strategies	22
8 The shape of things to come	23
2 Consciousness and supervenience	26
1 Logical supervenience: ontological and epistemological interpretations	27
2 (Merely) natural supervenience	28
3 The incoherence of (Chalmers’ versions of) supervenience	29
4 Natural supervenience and weak supervenience	35
5 Natural supervenience as an epistemological concept	37
6 More on ‘reading off’	41
7 Logical supervenience and reductive explanation	48
3 The explanatory gap	51
1 Intuitions and arguments	52
2 Analysing the intuition	57
3 Truth and adequacy	58
4 Explanatory adequacy and epistemic satisfaction	59
5 Proto-epistemic satisfaction	64
6 Mechanistic explanations and correlations	68
7 Explaining consciousness	70
4 Consciousness and higher-order experience	75
1 HOR models of consciousness	75
2 The structure of HOP theories	79
3 Presuppositions of the HOP model	81
4 The independence condition	83
5 The explanatory primacy of vehicles	94

vi Contents

6	The primacy of transitive consciousness	97
7	What has gone wrong?	98
5	Consciousness and higher-order thoughts	101
1	HOT models	103
2	The problem of circularity	105
3	The problem of regress	113
6	The structure of consciousness	122
1	Introduction	122
2	Consciousness as object of consciousness: empirical apperception	123
3	Transcendental apperception	126
4	Consciousness as experiential act	133
5	What it is like	136
6	The ubiquity of objectualism	141
7	Summary	147
7	What it is like	148
1	Against objectualism	149
2	What it is like as a phenomenal particular	150
3	What it is like as a phenomenal property	152
4	What it is like as a representational property	157
5	For actualism	166
6	Phenomenology by proxy	168
7	Objections and replies	173
8	Summary	176
8	Against objectualism II: mistakes about the way things seem	178
1	Introduction	178
2	Three mistakes about experience	179
3	The objectualist gloss: qualia	181
4	Perceptual completion and neural filling in	183
5	Dennett's criticism of filling in	186
6	Change blindness and the richness of experience	187
7	Category (2) mistakes: how an experience seems and how it really is	189
8	Mistakes of category (3)	191
9	Why the way an experience seems cannot be explained as awareness of qualia	194
9	Consciousness and representation	197
1	Brentano's thesis	198
2	Consciousness as revealing and as revealed	201
3	Phenomenal revealing	204
4	Consciousness of and consciousness that	208
5	Representationism	209
6	Object representationism	211
7	Mode representationism	213
8	Actualism and representationism	214

Contents

vii

10	Consciousness and the natural order	216
1	What it is like and reductive explanation	216
2	Consciousness and materialism	219
3	Consciousness and causality	221
4	The epiphenomenalist suspicion	222
5	The standard problem of epiphenomenalism	226
6	The epiphenomenalist suspicion allayed	232
	<i>Bibliography</i>	236
	<i>Index</i>	242

1 The problem of phenomenal consciousness

Consciousness is perceived by many to provide the principal threat to materialist accounts of the mind. This threat has been developed, in somewhat different ways, by a lineage of writers from Nagel (1974) through Jackson (1982, 1986), Levine (1983, 1993) to McGinn (1989, 1991) and Chalmers (1996). While the precise nature of the threat posed by consciousness has tended to vary, the concept of consciousness perceived to underlie this threat has held relatively constant. It is *phenomenal* consciousness that is considered problematic. There are serious problems, if the authors of the above lineage are correct, involved in finding a place for phenomenal consciousness in the natural order. This book is concerned with these problems, with why they are problems, and with whether these problems admit of a solution.

1 What is phenomenal consciousness?

Any study of phenomenal consciousness faces an immediate problem. There is no perspicuous way of defining the associated concept. That is, there is no non-circular way of specifying the content of the concept of phenomenal consciousness that does not rely on concepts that are equally obscure. Attempts to explain its content, accordingly, tend to rely on a number of devices, linguistic and otherwise.

Examples

Attempts to explain what phenomenal consciousness is often proceed by way of examples: the way things look or sound, the way pain feels, and, more generally, the experiential properties of sensations, feelings and experiences. Sensations and feelings will include things such as pain, itches, tickles, orgasms, the feeling one gets just before one sneezes, the feeling one gets just after one has sneezed, the feeling of cold feet, and so on. When experiences are enlisted to provide an explanation of the concept of phenomenal consciousness, it is typically perceptual (and, to

a lesser extent, proprioceptive) experiences that are to the fore. These will include visual (colour, shape, size, brightness, darkness, depth, etc.), auditory (sounds of various degrees of complexity, decomposable into quantities such as pitch, timbre and the like), olfactory (newly mown grass, rotting fish, freshly baked bread, a paper mill, the sea, etc.), tactile (the feel of fur, velvet, cold steel, newly sanded wood, greasy hair, sand beneath one's toes) and gustatory (habanero sauce, ripe versus unripe apples, Hermitage La Chapelle 1988 versus my father's home-made wine, etc.) experiences.

The list could, obviously, be expanded indefinitely, both within each category and by the adding of new categories (emotions, imagery, conscious thought, etc.). But this is not necessary. One point is, perhaps, worth noting. There is often a tendency, particularly in the case of visual examples, to place undue emphasis on perceptually basic, or near basic, experiences: experiences of a patch of redness, and the like. But this, as Wittgenstein would put it, might provide a diet of philosophically one-sided examples. Often, the phenomenal character of an experience can depend on its *significance* for the experiencer, and this, at least ostensibly, cannot be reduced to the significance of a conglomeration of perceptually basic, or near basic, properties. I once saw Muhammad Ali at Nashville airport, and, believe me, this was an experience which very definitely had a phenomenal character, one which could not be reduced to the aggregation of significances of patches of colour, shape, contours, and the like. Nor is it clear that we *must* think of this as a combination of perceptual experience plus emotional response, with the richer phenomenal character lurking in the latter rather than the former. Or, if this strategy is available here, then it is not clear why it would not be available in the case of our experience of perceptually basic properties; and this would undermine the idea that visual experiences, as opposed to the emotional response they evoke, have a phenomenal character.

In any event, the idea that motivates these sorts of examples is simply that anyone who has had any of the above experiences will know that they *feel* or *seem* a certain way, that there is *something that it is like* to undergo them. This brings us to device no. 2.

Rough synonyms

The concept of phenomenal consciousness is sometimes explained, and I use the term loosely, by way of terms that are roughly synonymous with the original expression. Thus, phenomenally conscious states are ones which have, or are defined by, a *phenomenology*, which have a certain

qualitative feel or *qualitative character*. Such states are *experiential* ones, *subjective* ones. They are states that essentially possess *qualia*. Most importantly, perhaps, for any phenomenally conscious state, there is *something that it is like* to be, or to undergo, that state. ‘Fundamentally’, writes Thomas Nagel, ‘an organism has conscious mental states if and only if there is something that it is like to be that organism – something it is like *for the organism*’ (1974: 166).

Just do it

The third device embodies what we might call the *Nike*TM approach. Just do it. More precisely, one is invited to construct the circumstances that will produce in one states with a particular form of phenomenal consciousness. Sometimes, for example, one is invited to inflict mild bodily trauma on one’s person to reacquaint oneself with the content of talk of phenomenal consciousness (Searle 1997: 97–9). The possibilities here are, of course, endless.

I think we would be advised to treat these devices with some suspicion, and some of the grounds for this will be examined more closely later on. Fundamentally, however, what seems to unite all three types of device is that they are, essentially, devices of *ostension*; they are means of *pointing*, or attempting to point, at phenomenal consciousness. And we are all familiar with the problematic status of attempts to point at private, inner, qualities, such as phenomenal consciousness purports, or is commonly taken, to be. So, the assumption that these devices are collectively sufficient to fix the meaning, or delineate the content, of the concept of phenomenal consciousness is far from certain. Indeed, this is precisely one of the assumptions that those who are sceptical of phenomenal consciousness will reject (see, for example, Dennett 1997: 117–18).

If the devices, even collectively, do not show that we know what we are talking about when we talk about phenomenal consciousness, they do show something much weaker, but something perhaps robust enough to provide a stepping-off point for further investigation. What the devices, or more importantly, the widespread presumed efficacy of the devices, do show is that a large number of people *think* they know what they are talking about when they talk about phenomenal consciousness. Indeed, I am one of those people. In fact, the people who explicitly deny that they know what they are talking about when they talk about phenomenal consciousness (and most of them do still talk about phenomenal consciousness, if only to deny the coherence of the concept) are, in all probability, limited to those antecedently *in the grip* of some quite specific theory of mind. A completely unscientific survey of some of my drinking

acquaintances, for example – who, I think they will not mind me saying, are very definitely *not* in the grip of some quite specific theory of the mind – indicates that they at least *seem* to have no difficulty in understanding what I am talking about when I talk about the what it is like of experience. Or perhaps they are just being polite. Or trying to shut me up.

In any event, that we, or most of us, *think* we know what we are talking about when we talk about phenomenal consciousness, even if we are mistaken in this thought, is the place where this book begins. This, then, is a book for all those who *think* they know what they are talking about when they talk about phenomenal consciousness. If the collection of devices outlined above is not sufficient to convince you that you at least think you know what you are talking about when you, or someone else, talks about phenomenal consciousness, then there is probably nothing in this book for you.

In fact, I labour our inability to define phenomenal consciousness, or to specify in any standard and perspicuous way the content of this concept, for a quite specific reason. *This is an essential datum that any account of consciousness should explain.* Our inability on this score is not something to be treated with embarrassment, swept under the carpet, lip-serviced, or mentioned at the outset and then forgotten. Rather, it is a feature of our understanding of the concept that any adequate account of consciousness should address and, hopefully, explain. Approaches that are, broadly speaking, eliminativist about phenomenal consciousness will explain this by saying that there is no coherent concept there to specify, or that what is there is a jumbled mish-mash of conceptually variegated strands that cannot be rendered into any coherent whole. While I am not convinced that such an explanation would work, even on its own terms, this book is, in any event, realist, not eliminativist, about phenomenal consciousness, and, as such, has no recourse to such strategy. The seeming ineffability of the concept of phenomenal consciousness imposes a fairly pressing requirement on realist accounts. If phenomenal consciousness is real, and if the corresponding concept is coherent, or reasonably so, then we should be able to *eff* it. And, if we cannot do this, then we have to come up with some explanation of why the concept of phenomenal consciousness cannot be *effed*.

2 The scope of ‘There is ...’

To say that an organism is conscious is, Nagel claims in his seminal (1974) paper ‘What is it like to be a bat?’, to say that ‘there is something that it is like to be that organism – something it is like *for* the organism’ (166).

And the claim that *there is something that it is like* to undergo a conscious experience is now one of the most common ways of explaining the idea that experiences, and the organisms that undergo them, are phenomenally conscious. The claim, however, is open to a variety of interpretation, some of which can, I think, be reduced to questions of the scope of the existential quantifier.

One obvious construal of Nagel's claim is that there is some object of conscious acquaintance and that all bats are acquainted with this object, while there is a distinct object of acquaintance such that all humans are acquainted with it. More generally, there is a certain form of consciousness that associates with being human, a distinct one that associates with being a bat, and so on. Indeed, it is possible to adopt an even broader conception of the what it is like of conscious experience. Flanagan (1992: 87), for example, claims that there is something that it is like to be conscious. And, again, one way of understanding this is as the claim that there is some object of conscious acquaintance and that all conscious creatures are acquainted with this object.

It is possible, however, to narrow considerably the scope of this claim. Thus, one might claim that what it is like is associated not with being conscious in general, nor with being a particular species of conscious organism, but, rather, with types of experience. One construal of this claim would entail that for every type of conscious experience there is some object of conscious acquaintance such that a creature which undergoes this type of experience is acquainted with that object. One might narrow the scope even further and claim that what it is like associates only with particular tokens of types of experience. On this view, for example, while there is no one thing that it is like to be in pain, there is something that it is like to suffer a particular token of pain. In an important, but strangely neglected, passage, Wittgenstein gestures towards the latter construal:

Let us consider the experience of being guided, and ask ourselves: what does this experience consist in when for instance our *course* is guided? Imagine the following cases:

You are in a playing field with your eyes bandaged, and someone leads you by the hand, sometimes left, sometimes right; you have to be constantly ready for the tug of his hand, and must also take care not to stumble when he gives an unexpected tug.

Or again: someone leads you by the hand where you are unwilling to go, by force.

Or: you are guided by a partner in a dance; you make yourself as receptive as possible, in order to guess his intention and obey the slightest pressure.

Or: someone takes you for a walk; you are having a conversation; you go wherever he does.

Or: you walk along a field track, simply following it . . .

'But being guided is surely a particular experience!' – The answer to this is: you are now *thinking* of a particular experience of being guided. (1953: #172–3)

There is no one thing that it is like to undergo the experience of being guided, but, rather, this what it is like fragments into the what it is like of particular (i.e. token) experiences of being guided.

There is, in fact, no straightforward inconsistency between the view that the what it is like attaches, in the first instance, to experiential tokens, and Nagel's claim that there is something that it is like to be a bat (or human). There are at least two ways of rendering these claims consistent, one in terms of the idea of set membership, the other which appeals to higher-order properties of what it is like. According to the first strategy, to say that what it is like to be a bat is different from what it is like to be a human is to say (i) that for each (actual or possible) bat experience-token there is an associated what it is like, and for each (actual or possible) human experience-token there is an associated what it is like, but either (ii) the set of bat what it is likes does not overlap with the set of human what it is likes or (iii) the overlap between the two sets falls below a certain threshold. According to the second strategy, the what it is likes of bat experience-tokens instantiate a certain essential higher-order property B, while the what it is likes of human experience-tokens instantiate a certain higher-order property H, and B is distinct from H. That is, what it is like instantiates various higher-order properties, properties which vary from human to bat. On this view, what it is like attaches primarily to mental tokens and derivatively (in virtue of its higher-order properties) to organisms.

The claim that the what it is like of conscious experience attaches primarily either to experience-tokens (or to experience-types), however, does give rise to the following, more radical, possibility. The claim that there is something that it is like to undergo a token of one experience-type, say pain, might mean something distinct from the claim that there is something that it is like to undergo a token of a different type of mental state, for example, to token-instantiate (occurently) the belief that Ouagadougou is the capital of Burkina Faso. That is, it cannot be assumed at the outset that consciousness is a unitary property that attaches uniformly across all mental states.

The suspicion that it is not such a property can, in fact, be independently motivated by the following, well-known, considerations. Consider, first, the distinction between sensations and propositional attitudes. Propositional attitudes can certainly be associated with a phenomenology. There can be, in a given instance, something that it is like to have, say, a certain belief. However, propositional attitudes, it is commonly thought,

are not *defined* by a phenomenology, and their possession by a subject does not entail that this subject is presented with any phenomenology at all, let alone a particular phenomenology. However, this does not seem to be the case with at least some sensations. While, if Wittgenstein is correct, the phenomenology associated with an experience E may vary from one token of E to another, it seems that having some phenomenology or other, and indeed having a phenomenology constrained within certain reasonably definite limits, is essential to the tokening of at least some, and perhaps all, sensations. Even within the category of sensations there appear to be important differences. It is not only common, but also seemingly perfectly appropriate, to characterise the phenomenology of bodily sensations – pains, itches, orgasms, and so on – in terms of the notion of *feel*. With items such as perceptual experiences, however, the characterisation of their phenomenology in terms of the notion of *feel* sits a lot less comfortably. This is why the epithet ‘feels’ is, in the case of perceptual experiences, typically replaced by ‘seems’. If we do want to say that it *feels* a certain way to see a green wall, or Muhammad Ali, then it is far from clear that *feel* means the same thing in this context as it does in the case of sensations. But, of course, *feel* is often used as an alternative appellation for the what it is like of conscious experience, sensational, perceptual or otherwise. To say that there is something that it is like to undergo a conscious experience is often taken as equivalent to saying that having that experience feels a certain way. And if this is correct, then we cannot assume, *a priori*, that the existential quantifier in the claim ‘There is something that it is like to undergo X’ ranges across the same quantity for all Xs.

Therefore, we should be alive to the possibility that what it means for a mental state to be phenomenally conscious can vary from one category of mental state to another, perhaps from one type of mental state to another, perhaps even from one token mental state to another. Perhaps the concept of phenomenal consciousness is a fundamentally hybrid concept.¹ And, if this is so, we would look in vain for a unified account of in what phenomenal consciousness consists. At the very least, this is not something to be ruled out *a priori*.

In later chapters, when the real argument starts, I propose to avoid these potential difficulties by focusing on, and working with, certain very general features that any instances of phenomenal consciousness must,

¹ Of course, many have claimed that the concept of consciousness is a hybrid one. What they typically have in mind, roughly, is the idea that consciousness comes in many forms: phenomenal, introspective, self, monitoring, reportability, etc., etc. The present point, however, concerns only the category of phenomenal consciousness, and the possibility being mooted is that this is itself a hybrid category.

I shall argue, possess. Whether or not phenomenal consciousness turns out to be a conceptually or theoretically unified item, I shall try to show that anything that could possibly count as an instance of a phenomenally conscious state must have certain features, and it is upon these features that the arguments will be built.

3 What is the problem of phenomenal consciousness?

The above problems, unclarities, and cautionary notes notwithstanding, we perhaps (hopefully) have enough in the way of a preliminary characterisation of the concept of phenomenal consciousness to proceed to a preliminary (again) characterisation of the problem or problems it raises. Phenomenal consciousness is widely, though far from universally, accepted to create at least the appearance of a problem for materialism. Agreement on precisely what this problem is, or appears to be, however, is far less widespread. The intuition that there is at least the semblance of a problem, here, is commonly supported by the way of various intuition pumps.

1 *Abused scientists*

Mary has been forced to live her entire life in a black and white room and has never seen any colours before, except for black, white, and shades of grey (Jackson 1982, 1986). Filling in the details would be a rather fatuous exercise, but presumably her skin has also been treated with some pigment that makes it appear a shade of grey, which pigment has also transformed her irises appropriately, her hair has been dyed black, etc., etc. Despite her dysfunctional upbringing, Mary has become the world's leading neuroscientist, specialising in the neurophysiology of colour vision. She knows everything there is to know about the neural processes involved in the processing of visual information, about the psychophysics of optical processes, about the physics of environmental objects, and so on. However, despite this extensive knowledge, when she is let out of her black and white room for the first time, it seems plausible to suppose, she learns something new; she learns what it is like to experience colour. And, if this is correct, then this knowledge is neither something she possessed before nor something that could be constructed from the knowledge she possessed before.

2 *Zombies*

A zombie, in the philosophical as opposed to the Hollywood sense, is an individual that is physically and functionally human, but which lacks

conscious experience (Chalmers 1996; Kirk 1974, 1994). Thus, my zombie twin is physically identical to me and, we can suppose, is embedded in an identical environment. Moreover, he is functionally identical to me in that he is processing information in the same way, reacting in the same way as me to the same inputs, and so on. Nevertheless, he lacks phenomenal experience; he has no phenomenal consciousness. My zombie twin is not, it is generally accepted, a natural possibility (that is, he is incompatible with the laws of nature) but he is, it has been argued, a logical possibility.

3 *Deviants*

It is logically possible for there to be a world where qualia are inverted relative to the actual world (Shoemaker 1982; Chalmers 1996). My inverted twin is physically identical to me but has inverted conscious experiences. Thus, for example, where I have a red experience (i.e. an experience *as* of red) my inverted twin has a green experience (i.e. an experience *as* of green). That is, when he looks at a fire engine, he has an experience of the same qualitative colour character as I do when I look at grass. Again, my inverted twin may not be a natural possibility, but he is, it has been argued, a logical possibility.

4 *Demons*

Laplace's Demon is able to read off all non-basic facts from basic ones (Chalmers 1996). That is, the Demon knows every detail about the physics of the universe, the configuration and evolution of all the basic fields and particles that make up the spatiotemporal manifold. And from this knowledge, the Demon can read off, or infer, every other fact about the universe. Or, rather, *almost* every other fact. For, it has been argued, the Demon would not be able to read off facts about conscious experience (Chalmers 1996). Indeed, the Demon could not even work out, from its knowledge of the basic facts alone, that there is any conscious experience at all, let alone what it is.

A motley crew. Surely, it is only in recent discussions of consciousness – and perhaps some fairly questionable B-movies – that one could possibly find such a collection of characters. But the question is: what does all this mean? And this is a good question, one that subsequent chapters will spend some time trying to work out, and one that as yet has nothing even close to an accepted answer.

However, it is possible to broadly identify two axes along which potential answers may be developed. On the one hand, one can understand

the examples as establishing, or suggesting, an ontological or *metaphysical* conclusion that is, essentially, dualistic in character. Phenomenal experiences are distinct from, and not reducible to, any physical event, state or process. This conclusion is (or has at one time been) endorsed, on the basis of one or more of the above scenarios, by Jackson (1982, 1986) and Chalmers (1996). On the other hand, one can understand the examples as establishing, or suggesting, an *epistemological* conclusion. Roughly speaking, our knowledge of physical facts does not, in some way, add up to knowledge of conscious experience, and, consequently (perhaps) physical explanations do not, in some way, add up to explanations of consciousness. There is, as it is often put, an *explanatory gap* between consciousness and the physical. This conclusion has been endorsed by Levine (1983, 1993) and McGinn (1989, 1991, 1993) among others. Of course, those who endorse the metaphysical conclusion are also going to endorse the epistemological claim, and this is the case with Jackson and Chalmers. However, it is possible to endorse the epistemological claim alone.

In fact, there are, in my view, good reasons for endorsing the epistemological claim alone. All the above examples turn, ultimately, on a difference between phenomenal and physical *concepts*, and it is difficult to turn this into any substantive difference between phenomenal and physical *properties*. But it is the latter difference that is required to underwrite the metaphysical conclusion.

To see this, consider the knowledge argument. There are, in fact, various strategies available to the materialist should she want to resist the metaphysical interpretation of the significance of the knowledge argument. The one I favour is due to Brian Loar (1990). According to Loar, the materialist can allow that Mary acquires new information when she leaves the room, but she does so only under an *opaque* reading. Transparent construals of the information acquired by Mary would, in effect, beg the question against materialism. Drawing (legitimate) metaphysical conclusions from opaque contexts is never easy. And, given the opaque construal of what Mary learns, we can construct *prima facie* analogous cases, where a metaphysical conclusion manifestly does not follow from the premises. Thus, to borrow from Loar, Kate learns that the bottle before her contains $\text{CH}_3\text{CH}_2\text{OH}$. But, on an opaque reading, she does not know that the bottle contains alcohol. That is, she does not know that the bottle contains stuff called alcohol, or that the bottle contains the intoxicating component of wine and beer, the component that makes people drunk. Indeed, we can suppose that innocent Kate even lacks the ordinary concept of alcohol. Then, when she inadvisedly consumes the bottle's contents, she acquires new information: that the bottle contains

alcohol. If the knowledge argument, on the metaphysical construal, had a generally valid form, we could then infer from Kate's epistemic situation that alcohol is not identical with $\text{CH}_3\text{CH}_2\text{OH}$. And this, evidently, does not follow.

What seems to be going on here is that we have two distinct concepts associated with the same substance; one a theoretical-physical concept, the other what Loar calls a *recognitional* concept. The substance alcohol can be picked out both by way of theoretical description, and in terms of the properties by which one typically recognises it. However, the two types of concept are conceptually independent of each other, and this explains both why the above opaque reading of what Kate learns is possible and why this opaque reading does not yield a substantive metaphysical conclusion.

A recognitional concept has the form 'x is one of that kind'; i.e. they are type-demonstratives grounded in dispositions to classify, by way of perceptual discriminations, certain objects, events, and situations. Recognitional concepts, crucially, are typically conceptually independent of, and irreducible to, theoretical-physical concepts, even where both concepts, as in the above case, pick out the same property.

Loar argues that phenomenal concepts are essentially recognitional in character. Thus, materialism at the metaphysical level is underwritten by the claim that phenomenal and physical-functional concepts can pick out the same property, while the conceptual independence of these concepts is explained by the fact that recognitional and theoretical concepts are, in general, conceptually independent, and that the former cannot be reduced to the latter. Thus the epistemological reading of the knowledge argument is safeguarded and explained, and the metaphysical reading shown to be invalid.

This, of course, takes us only part of the way. It is not difficult to find a difference between the case of Kate and the case of Mary. Kate lacks knowledge of the contents of the bottle under a contingent description of it: stuff that gets you drunk. However, Mary's acquired information of what it is like to experience colour does not conceive it under a contingent mode of presentation. It is not as if she is conceiving of a property that presents itself *contingently* thus: it is like such and such to experience P. Being experienced in this way is essential to the property Mary conceives. Thus, when Mary later acquires new information (construed opaquely) the novelty of this information cannot be explained – as in the case of Kate – as her acquiring a new contingent mode of presentation of something she has known all along. This is why, according to its proponents, the knowledge argument can be valid on an opaque reading. There is no contingency in Mary's conception of the new phenomenal information

that explains it as a novel take on old facts. Therefore, we must suppose that she learns new facts *simpliciter*, and not new conceptions of old facts.

As Loar points out, however, there is an implicit assumption in this argument: a statement of property identity that links conceptually independent properties is true only if at least one concept picks out its associated property by way of a contingent mode of presentation of that property. Conversely, the underlying idea is that if two concepts both pick out the same property by way of its essential properties, neither mediated by contingent modes of presentation, then one ought to be able to see *a priori* – at least after optimal reflection – that they pick out the same property. If the two concepts pick out the same property by way of essential modes of presentation, then those concepts themselves must be logically connected.

However, Loar argues, convincingly in my view, that this assumption should be rejected. It rests on the idea that (i) if a concept picks out a property by way of an essential mode of presentation, then that concept must capture the *essence* of the property picked out, and (ii) if two concepts capture the essence of the same property, then there must exist constitutive conceptual connections between those concepts, such that one concept is derivable from the other *a priori*. However, when expressed in this way, it is fairly clear that these are equivocating uses of ‘capture the essence of’. On one use, it expresses a referential notion that comes to no more than ‘directly rigidly designate’. On the other, it means something like ‘be conceptually interderivable with some theoretical predicate that reveals the internal structure of’ the designated property. But the former does not imply the latter. Claims about rigid designation do not, in general, imply the conceptual interderivability of the designating concepts.

Once we allow that phenomenal and physical concepts can both (i) pick out a property by way of an essential mode of presentation, but (ii) still be conceptually independent of each other, then essentially the same deflationary strategy can be adopted with respect to the rest of the assorted cast listed above. The logical possibility of zombies, that is, need only be taken as indicative of the conceptual independence of phenomenal and physical-functional concepts, and not of any deeper metaphysical division. A similar account will be applicable to the case of the qualia-inverted deviants; their logical, as opposed to natural, possibility, need be indicative only of the logical independence of phenomenal from physico-functional concepts. And the failure of Laplace’s Demon to read off phenomenal facts from non-phenomenal ones, again, need only indicate the conceptual independence of phenomenal concepts from physical or functional ones.

Loar's account, of course, will not satisfy everyone. Indeed, despite my general sympathy to this line of reasoning, I think that Loar's claim that phenomenal concepts are recognitional ones needs to be severely qualified (see chapter 7). Nevertheless, I suspect that a story substantially similar to the one Loar tells can be made to work. And, for this reason, I am going to treat the problem of phenomenal consciousness as a primarily epistemological one. This may be incorrect. Perhaps consciousness provides a metaphysical problem also. If so, then so be it. If there is a genuine metaphysical problem, then it is outside the scope of this book. The book's subject is the epistemological problem posed for materialism by phenomenal consciousness: the existence, or apparent existence, of an explanatory gap between the phenomenal and the material. One thing is clear: if consciousness is not an epistemological problem, then it is not a metaphysical problem either.

4 Explaining consciousness

The problem of explaining phenomenal consciousness is the problem of explaining how consciousness can come from what is not conscious. And one can understand the idea of consciousness coming from what is not conscious either *causally* or *constitutively*. For various reasons I prefer the constitutive construal. Suppose, for example, we say that phenomenal consciousness is causally produced by brain activity (McGinn 1989, 1991; Searle 1992). Causal relations, as Hume taught us, involve distinct existences. So, if we talk of consciousness being causally produced by neural activity then there is a danger that we have already implicitly bought in to a metaphysical understanding of the problem: we have already implicitly assumed that consciousness is distinct from this neural activity. We can avoid this metaphysical temptation by regarding the causal relations by which consciousness is produced as diachronic, rather than synchronic. But then the production of consciousness by the brain has to be understood in terms of the idea that a phenomenal property instantiated at time t is produced by brain activity occurring at $t-1$. But this does not seem to be the correct model for understanding the production of consciousness by neural activity. What neural activity occurring at time $t-1$ actually causally produces is neural activity occurring at t . And then we still have the problem of explaining how consciousness is produced by this neural activity of time t . If we want to insist that this relation of production is a causal relation, then we fall right back into the metaphysical construal of the problem.

Intuitively, the relation of production we require seems to be more like the relation between the observable properties of water and its underlying

structure, and this (*pace* Searle 1992) is not a causal relation. Rather, the observable properties of water are, in some sense, *constituted* by the underlying molecular properties. Phenomenal consciousness, on this construal, is somehow constituted by neural activity, and the problem of consciousness is the problem of explaining how this could (possibly) be so. More generally, how can consciousness be constituted by what is not conscious?

While, for these and other reasons, I favour the constitutive rather than causal construal of the claim that consciousness is produced by what is not conscious, nothing much turns on this assumption. The arguments to be developed in the following chapters have, I think, an application broad enough to cover both constitutive and causal senses, and, accordingly, I shall usually employ the more general term *production* to subsume both constitutive and causal senses of the relation between consciousness and the material.

The focus of this book, then, is whether it is possible to provide an explanation of how phenomenal consciousness is produced by what is not conscious. We know, I shall suppose, that it is, in fact, produced by what is not conscious, that is why we are not concerned with the metaphysical construal of the problem. What we want is an explanation of how it is so produced.

In attempting to provide an explanation of phenomenal consciousness, it is possible to adopt two quite distinct strategies; one, as I shall put it, *vertical*, the other *horizontal*. Vertical strategies, roughly speaking, attempt to *build consciousness up* from what is not conscious. Horizontal strategies, again roughly, attempt the explanatory task by attempting to *pull consciousness out* into what is not conscious, i.e. the world. The next two sections deal with the former type of explanatory strategy, the one after that deals with the latter.

5 Vertical strategies I: the mind-body problem

To build consciousness up from what is not conscious is to show how various non-conscious processes can, collectively, constitute conscious activity. This strategy of, as we might call it, *phenomenal tectonics*, of constructing the phenomenal from the non-phenomenal, divides into two distinct approaches. On the one hand, we can try to build consciousness up from processes that are neither conscious nor mental. Our appeal, here, is likely to be to the brain, to neural activity broadly construed. On the other, we can try to construct consciousness from processes that are non-conscious but which are mental. The explanation here is likely to involve, quite centrally, higher-order mental states of some sort, states which are identified as not being essentially conscious, or, at the very

least, as not being phenomenally conscious. The former strategy requires solving the *mind-body* problem, the latter requires, in effect, solving the *mind-mind* problem.

As an example of the former strategy, consider the much-trumpeted hypothesis of Francis Crick and Christof Koch (1990) that 40 Hertz oscillations in the visual cortex and elsewhere may be the fundamental neural feature responsible for conscious experience. According to Crick and Koch, 40Hz oscillations play a crucial role in the *binding* of various sorts of information into a unified and coherent whole. Two different kinds of information about a visual scene – the shape and distance of an object, for example – may be represented quite separately, but Crick and Koch suggest that these separate neural representations may have a common oscillatory frequency and phase-cycle, allowing the information to be bound together by later processes and stored in working memory.

This provides a neurobiological model of how disparate information might be integrated in working memory. And it might, with suitable elaboration, be developed into an account of how information is integrated and brought to bear in the global control of behaviour. However, what it is not, or does not seem to be, is an explanation of phenomenal consciousness. Crick and Koch have, in fact, presented only an account of how a certain functional capacity – the capacity for integration of disparate information – is implemented in the brain. But this would be an explanation of phenomenal consciousness only if it could be shown, I think it is fair to say counterintuitively, that such consciousness could be reduced to a feature of, or function of, the capacity for binding. Much further argument is required; Crick and Koch have not presented such argument; and it is, indeed, difficult to see what such further argument might look like.

Similar limitations seem to affect Gerald Edelman's neurobiological theory of consciousness (Edelman 1989, 1992). The core of his theory is provided by the idea of re-entrant neural circuits which afford the conceptual categorisation of perceptual signals before they contribute to memory. On the basis of this, perceptual information interacts with internal states in various ways and give rise to 'primary consciousness'. The introduction of a new memory element of 'semantic bootstrapping' explains the generation of 'higher-order consciousness', and the concepts of the self, past and future. And this is linked to language production through Broca's and Wernicke's areas. Once again, and as Edelman in effect acknowledges, this is not an explanation of phenomenal consciousness. Rather, insofar as Edelman's theory is devoted to consciousness at all (as opposed to what appear to be its primary concerns with

perception, memory and language), what it seems it might explain are certain aspects of perceptual aspect-consciousness – the effects of perceptual processing on later processing operations and on the global control of behaviour – and aspects of self-consciousness, in particular, the origin of the concept of the self. Edelman's theory, then, is an account of certain forms of access-consciousness, not an account of phenomenal consciousness.

The limitations, or perceived limitations, of these approaches are indicative of a general problem with attempts to build consciousness up out of neural processes. It seems that no matter how much we know about the neural processes implicated in the production of consciousness – about oscillatory frequencies and phase-cycles in the visual cortex, the structure and function of re-entrant neural circuits, and so on – this goes no further than providing an explanation of certain *functional* capacities of the brain. Such knowledge might enable us to see how the brain binds disparate information into a unitary whole, how it underwrites the ability to categorise perceived events, and so on. But, glaringly, what it does not seem to provide is an explanation of phenomenal consciousness. In particular, such information does not enable us to see *how* the brain produces phenomenal consciousness.

Therefore, the standard objection to neural models of consciousness, and at present this is an objection I simply note not advocate, is that they do not bridge the perceived explanatory gap between the phenomenal and the physical. On the contrary, they merely reinforce that gap. This objection has been developed, in important and sophisticated forms, by Colin McGinn and David Chalmers. Both Chalmers and McGinn argue that phenomenal consciousness cannot be reductively explained in physical terms. Chalmers' argument is examined in chapter 2, McGinn's in chapter 3.

6 Vertical strategies II: the mind–mind problem

The second type of vertical strategy attempts to build phenomenal consciousness up out of states that are not conscious, or not essentially so, but which are, nonetheless, mental. The most influential recent forms of this strategy consist in the attempt to explain phenomenal consciousness in terms of *monitoring* or *introspective* consciousness.

Introspective consciousness, very broadly speaking, is the process either by which we can become aware *of* our internal states or, on some accounts (e.g. Shoemaker 1994; Dretske 1995), by which we become aware *that* we instantiate certain internal states. Reflecting the distinction between being aware *of* and being aware *that*, there are several, quite

different, models of in what such awareness consists. On a simple object perception model, for example, introspection will have essentially the same dyadic structure as perception, a structure constituted by a content bearing state whose directedness towards its intentional object is typically regarded as being effected by way of some sort of causal relation. On such a model, introspective abilities afford us access to mental events, states, processes, or objects. On other models, however, introspection affords us access to mental facts, where such access is commonly thought to be grounded in displaced perception of non-mental (typically, but not necessarily, environmental) events, states, processes, or objects. We learn *that* we instantiate certain mental properties (where a mental fact is conceived of as the instantiation of a mental property in a person at a time) in virtue of our displaced perception of non-mental objects (Shoemaker 1994; Dretske 1995). The principle is very much like discovering how much petrol is in the tank of one's car by way of perception of the fuel gauge.

Monitoring consciousness can, also, take at least two forms. According to the *higher-order experience* model associated with Armstrong (1968, 1981) and Lycan (1987, 1996), the consciousness of any given mental state *M* is to be explained in terms of the subject of that state having a quasi-perceptual experience of *M*. Both Armstrong and Lycan flesh out this general idea in terms of the notion of *internal scanning*. A subject's access to her mental states takes the form of an internal scanning, or monitoring, of those states by higher-order neural structures (that is, neural structures whose function is to register the activity occurring in other neural structures). And when a subject has such access to a given mental state, that mental state is a conscious one.

A related, but importantly distinct, account of monitoring consciousness is provided by the *higher-order thought* model associated with Rosenthal (1986, 1993) and also with Carruthers (1996). According to this account, access to one's mental states takes the form of higher-order thoughts about those mental states. On this view, very roughly, the consciousness of any given mental state *M* consists in the subject of *M* possessing a higher-order thought to the effect that he possesses, or instantiates, *M*.

Monitoring and introspective consciousness are often run together. This need not be a conflation, still less a confusion. It is possible to draw a distinction here, but it is far from clear that the distinction will correspond, in any salient way, to the use of the terms 'monitoring' and 'introspective' in the relevant literature: for this use varies. Nonetheless, if we do want to preserve a distinction between introspective and monitoring consciousness, and given the variability of the use of the terms in

the literature such a distinction is by no means obligatory, we can do so by regarding introspective consciousness as related to monitoring consciousness as genus is to species. Internal scanning, for example, is one way of understanding how an object perception model of introspection might be *implemented*. And an appeal to higher-order thoughts might be one way of understanding how introspection affords us access to mental facts, if we assume that in being aware of a thought we are thereby aware of its content. Thus, it is possible (though certainly not necessary) to view accounts of monitoring consciousness as causal or (probably more accurately) quasi-causal models of how the conceptual analyses proffered by accounts of introspection might be implemented.

Usually, however, monitoring models are presented not just as relatively determinate models of introspective consciousness, but as models of consciousness in general. Armstrong, for example, promotes his monitoring model as an account of 'consciousness in the most interesting sense of the word'. This is because Armstrong thinks that introspective consciousness has a peculiar centrality relative to other forms. And similar claims for the comprehensive scope of the monitoring model, claims resting on a similar faith in the centrality of introspection, can be found, to a greater or lesser extent, in most of the model's principal defenders. Such claims are often accompanied by a paring down (unacceptable to many) of what a model of consciousness can be expected to explain, or of what features can legitimately be thought of as essential to consciousness. Lycan (1996), for example, presents his version of the internal monitoring model as an account of consciousness. However, he also claims that there are certain aspects of consciousness – specifically its phenomenal character – that cannot be explained in terms of internal monitoring. He is unperturbed by this since he thinks that 'qualia problems and the nature of conscious awareness are mutually independent and indeed have little to do with each other' (1990: 756). Rosenthal (1990), on the other hand, is more equivocal. On the one hand, he explicitly separates consciousness from 'sensory quality', and says he is giving only a theory of the first. This suggests that phenomenal consciousness lies outside the scope of his account. On the other hand, he also says that a state is conscious when there is something that it is like to be in that state, which suggests that his subject is phenomenal consciousness after all. Carruthers (1996, 1998) asserts, reasonably explicitly, that his higher-order thought model is intended as an account of phenomenal consciousness also.

What unites introspective and monitoring consciousness, in all their forms, is that the corresponding concepts are all *functional* concepts. Indeed, not only are these concepts of consciousness all functional concepts, they all seem to be assimilable to a particular type of functional

concept, broadly understood. That is, they all seem understandable in terms of the notion of *access*. Introspective and monitoring consciousness, in whatever specific form they take, consist in access to one's mental states.

In an important paper, Block (1995) has distinguished between phenomenal consciousness and access-consciousness (*P-consciousness* and *A-consciousness* in his terminology). And Chalmers (1996) draws the essentially equivalent distinction between what he calls *consciousness* and *awareness*. In the spirit of Block, but the letter of Chalmers, we can characterise *access-consciousness* in the following way.

Access-consciousness: a subject, S, is access-conscious of some information, I, if and only if I is directly available for the global control of S's behaviour.

To talk of information being used in the *global* control of behaviour is just to say that this information is available to be brought to bear in a wide range of behavioural processes: verbal, motor, attentive, and the like (Chalmers 1996: 225). The motivation for the inclusion of *directness* in the above definition lies in the intuitive idea that states of consciousness, of whatever stripe, must be *occurrent* rather than *dispositional* in character.

The claims that phenomenal consciousness can be explained in terms of higher-order experiences or higher-order thoughts, then, are specific versions of a more general thesis: phenomenal consciousness can be explained in terms of access-consciousness. Both Block and Chalmers have provided reasons, in my view compelling reasons, for thinking that this thesis cannot be true. However, I shall focus on the specific versions of the thesis: the idea that phenomenal consciousness can be explained in terms of monitoring consciousness. Thus, chapter 4 examines higher-order experience accounts of consciousness; chapter 5 is concerned with the higher-order thought alternative.

However, it is well known that there are serious *prima facie* problems with the attempt to assimilate phenomenal consciousness to the possession of either higher-order experiences or higher-order thoughts. If the notion of a higher-order experience is (as in Armstrong 1981) explained in terms of the concept of internal scanning, then the problem seems to be that such scanning is not sufficient for phenomenal consciousness. As Rey (1983) has pointed out, ordinary laptop computers are capable of internal scanning, and it is not clear who would want to claim that they are conscious.² If, on the other hand, it is asserted that, unlike the case of

² Rey, in fact, advocates that we accept that internal scanning is sufficient for consciousness, if there is such a thing, and so he concludes that consciousness is a concept that includes and precludes laptop computers, and hence that the concept of consciousness is incoherent. Far more plausible, I think, is simply to reject the claim that internal scanning is sufficient for phenomenal consciousness. If so, we get no incoherence.

the laptop, the higher-order experiences must be conscious ones, then the account immediately runs into a problem of regress. On the other hand, interpreting the monitoring account in terms of higher-order thoughts seems *prima facie* equally problematic. Most obviously, the identification of phenomenal consciousness with the possession of higher-order thoughts shares the apparent over-intellectualism of the identification of phenomenal consciousness with self-consciousness. Dogs and human infants, it seems overwhelmingly likely, have phenomenally conscious states without thoughts to the effect that they have those states.

Of course, none of these considerations can, as yet, be taken as compelling. But they do raise a certain problem of *procedure*. The *prima facie* implausibilities associated with the idea that phenomenal consciousness can be explained in terms of either higher-order experience or higher-order thought models often leads to a certain type of fall-back strategy being embraced by defenders of such models. The strategy involves, essentially, a *paring down* of the explanandum: consciousness is often divested of those properties that are most problematic for higher-order representation models. Thus, as was mentioned earlier, Lycan (1990, 1996) seeks to divest at least the core notion of consciousness of its traditional association with qualia and phenomenal character. The core concept of consciousness, for Lycan, is awareness, and this can be accounted for in terms of the internal monitoring model associated with Armstrong. Consciousness, according to Lycan, is ‘the functioning of internal attention mechanisms directed upon lower-order psychological states and events’, and these attention mechanisms are devices that ‘have the job of relaying and/or co-ordinating information about ongoing psychological events and processes’ (1990: 755). Qualia, and the phenomenal character of experience in general, are to be explained by other means (according to Lycan, a functionalist-representationist account will do the trick).

The paring down of reduced properties is, of course, a standard part of the process of reduction and, in itself, is unobjectionable. However, in the case of phenomenal consciousness, this practice seems peculiarly problematic. In particular, it is necessary, but difficult, to steer a middle ground between twin dangers. The one danger is that of triviality, and this seems to threaten Lycan’s account in particular. The danger is that one pares down the concept of consciousness so much that what one is left with is simply the claim that internal monitoring is internal monitoring. The problem then is that we have all these other properties – in particular, phenomenal character – whose nature we still have to explain. In this case, the access-consciousness model has not bought us very much. Much of the hard work is left to do. Thus, we find that much of Lycan’s account of consciousness is provided not by his internal monitoring model, but

by his functionalist-representationist account of qualia, which he needs because he acknowledges that his internal monitoring model will not yield an explanation of phenomenal character.

The other danger is that we are forced to introduce distinctions into our discussion of consciousness that opponents will not accept. Consider, for example, Dennett's (1991) position. One strand of Dennett's argument consists in the claim that the idea that experiences have a phenomenal character is the result of a quasi-cognitive illusion brought about by our being in the grip of various illicit assumptions about the nature of the mind. This position, however, seems to force on us the claim that consciousness does not really have phenomenal character, it just *seems* as if it does. However, this involves foisting on the discussion the very appearance/reality distinction that defenders of phenomenal consciousness will reject. Phenomenal properties are precisely properties of *seeming*. So, if consciousness *seems* to have phenomenal properties, it thereby *does* have phenomenal properties (Strawson 1994). And their nature is, therefore, something that still requires explanation.

Either way, it seems there is still explanatory work to do subsequent to the reduction, or alleged reduction. Either the reduction requires that we pare down the concept of phenomenal consciousness so much that there are properties left over that the reduction does not incorporate. Then, all the hard work is still to be done. Or, we implicitly reintroduce, by way of an appearance/reality distinction, phenomenal consciousness in an unreduced form. Either way, we have not succeeded in reducing phenomenal consciousness.

If these ruminations indicate anything at all, then it is that the dialectical situation is rather complex. I (i) might object to higher-order representation accounts on the grounds that they do not explain features which (ii) their proponents' claim they do not have to explain but (iii) where I claim that their proponents' claim that they do not have to explain such features is illegitimate on the grounds that it commits them to triviality or illegitimate distinctions, but where (iv) their proponents will claim that there is nothing illegitimate about these distinctions and the triviality is only apparent, and so on. Clearly what is needed is a way of cutting through the dialectical complexity.

So, here it is. In chapters 4 and 5, when I develop the case against higher-order representation accounts of consciousness, I shall not even suppose that such models are in the business of explaining phenomenal consciousness. Rather, I shall assume only that they are in the business of explaining our access to our own mental states or to mental facts that we instantiate. And, then I shall argue that they cannot even do this. We do not, in these chapters, even need to get into the issue of *phenomenal*

consciousness; higher-order models cannot even explain *introspective* consciousness. And if there is one thing that remains clear amidst the dialectical cut and thrust, it is surely this: if higher-order models cannot even explain introspective consciousness – our access to mental states we possess or facts we instantiate – then they have no hope of explaining phenomenal consciousness.

7 Horizontal strategies

Horizontal strategies are characterised by the attempt to explain phenomenal consciousness not by building it up out of neural or functional components, but by, figuratively speaking, *pulling consciousness out* into the world. That is, very roughly, a horizontal strategy will try to show that the principal features of phenomenal consciousness are constituted not by features of neural or functional activity but, rather, by features of the world in which this activity is situated. The most common form of horizontal strategy is known as *representationism*. Very roughly, this is the view that the phenomenal character of an experience does not go beyond its representational content or, equivalently, that all phenomenal differences are representational differences. The phenomenal, that is, can, ultimately, be explained in terms of the representational.

Tye (1995) supplies a recent, sophisticated, version of representationism. According to Tye, the phenomenal *character* of an experience is identical with the phenomenal *content* of that experience, and phenomenal content is just a species of *intentional* or representational content. Specifically, phenomenal content is PANIC: poised, abstract, non-conceptual, intentional content.

The claim that the relevant contents are *poised* is the claim that they attach to the output representations of the relevant sensory modules and, thus, are in a position to make a direct impact on the belief/desire system. To say that the contents are in a position to impact the belief/desire system is not to claim that they actually do make such impact. Rather, it is to say that they supply the inputs for certain cognitive processes, ones which have the job of producing beliefs, or desires, directly from the appropriate perceptual representations if attention is properly focused (and the relevant concepts are possessed).

The claim that the relevant contents are *abstract* is the claim, roughly, that no particular concrete objects enter into these contents. This is required by the fact that different concrete objects can, phenomenally, look or feel exactly the same. The identity of the object presented to the subject of an experience, then, does not matter for the phenomenal content of that experience. Rather, the content depends on the general, phenomenal, features presented to the experience's subject.

The claim that phenomenal content be *non-conceptual* is the claim that these general features entering into the content of an experience need not be ones for which the experience's subject possesses matching concepts. It is possible to recognise, for example, far more different shades of colour than for which we possess stored representations. Perceptual discriminability outstrips our conceptual resources. Hence, phenomenal content is non-conceptual.

Tye's account provides one of the most sophisticated, and influential, forms of representationism. But what unites all forms of representationism is the idea that the phenomenal character – the what it is like – of a conscious experience is determined, indeed constituted, by the representational features of that experience. Since representational properties are not determined purely by what is occurring inside the head of an experiencing subject, representationism is committed to the view that phenomenal character is not constituted by processes occurring inside the head of experiencing subjects. The phenomenal character of an experience is constituted not just by what is going on inside the head of an experiencing subject, but also by what exists in the world in which that experiencing subject is situated.

There is no reason, of course, to regard vertical and horizontal strategies as mutually exclusive. They can be combined in a variety of ways. Lycan, for example, advocates a vertical approach to explaining (what he calls) awareness, and a horizontal, representationist, approach to explaining phenomenal character. More generally, it may turn out that a vertical approach is able to handle some features of phenomenal consciousness while a horizontal strategy is able to handle the rest. Or it may not so turn out. In any event, the horizontal, representationist, account of phenomenal character will be examined in chapter 9.

8 The shape of things to come

The book to follow can, nominally, be thought of as divided into two parts. Part 1, which consists of chapters 2–5, is concerned with vertical attempts to explain consciousness. Of these chapters, the first two examine the prospects of attempts to explain consciousness in physical terms. Or, more precisely, they examine two recent and (deservedly) influential attempts to show that these prospects are minimal or non-existent. Chapter 2 focuses on Chalmers' attempt to show that consciousness cannot be reductively explained in physical terms. Chapter 3 examines McGinn's case for the claim that there exists an unbridgeable explanatory gap between consciousness and the physical world.

My attitude to both positions is somewhat equivocal. I believe that both McGinn and Chalmers *might* be right, but I am not *convinced* that they are.

More specifically, I shall try to show that the arguments of both McGinn and Chalmers are far from conclusive. In so far as anything concrete emerges from chapters 2 and 3, then, it is simply that consciousness might be reductively explainable in physical terms.

Chapters 4 and 5, the remaining chapters of part 1, are concerned with attempts to explain phenomenal consciousness in terms of access-, specifically monitoring, consciousness. Chapter 4 examines the higher-order experience account of consciousness. In chapter 5, the focus is on higher-order thought models. I shall argue that both types of model fail as explanations of consciousness. They are not even adequate as models of introspective consciousness; and have no chance whatsoever of explaining phenomenal consciousness.

The nominal part 2 of this book comprises chapters 6–10. In these chapters, I shall develop a case against the possibility of explaining phenomenal consciousness in terms of what is not conscious, a case that applies equally against both vertical and horizontal explanatory strategies. In particular, I shall argue that the real reason why phenomenal consciousness is so problematic, from an explanatory point of view, has not been understood. The real reason, I shall argue, is this. The phenomenal aspects – the what it is like – of experience are not themselves objects of conscious awareness. They are not items *of* which we are aware in the having of an experience. Rather, they are items that constitute the taking of distinct, and non-phenomenal, items as the objects of experience. That is, the phenomenal aspects of experience are not items *of* which we are aware in the having of an experience, but (in a sense to be made clear) items *in virtue of* which, or *with* which, we are aware in the having of that experience. Alternatively, in a sense again to be made clear, phenomenal features are not *empirical* but *transcendental* features of experience. The bulk of the argument for these claims is to be found in chapters 6, 7 and 8.

This view of the phenomenal, it will be shown, has certain clear affinities with the representationist account of phenomenal character, in particular, the rejection of the view that phenomenal features are constituted purely by what is going on inside the head of an experiencing subject. However, in chapter 9, I shall draw attention to some of the important differences between this view and the representationist one. There, I shall argue that the transcendental status of phenomenal features of experience rules out the representationist attempt to explain the phenomenal in terms of the representational.

In chapter 10, the final chapter, I shall argue that the transcendental status of phenomenal properties or features is incompatible with any attempt to reductively explain the phenomenal in terms of the non-phenomenal.

The problem of phenomenal consciousness, the problem of *explaining* how phenomenal consciousness can come from what is not conscious, has no solution. We know consciousness is produced by what is not conscious, but we can never understand how. Chapter 10 also explores the wider question of the place of phenomenal consciousness in the natural order. It will be argued that the prospects for finding a place for consciousness in the natural order are not as bleak as the failure of reductive explanation might lead us to think.