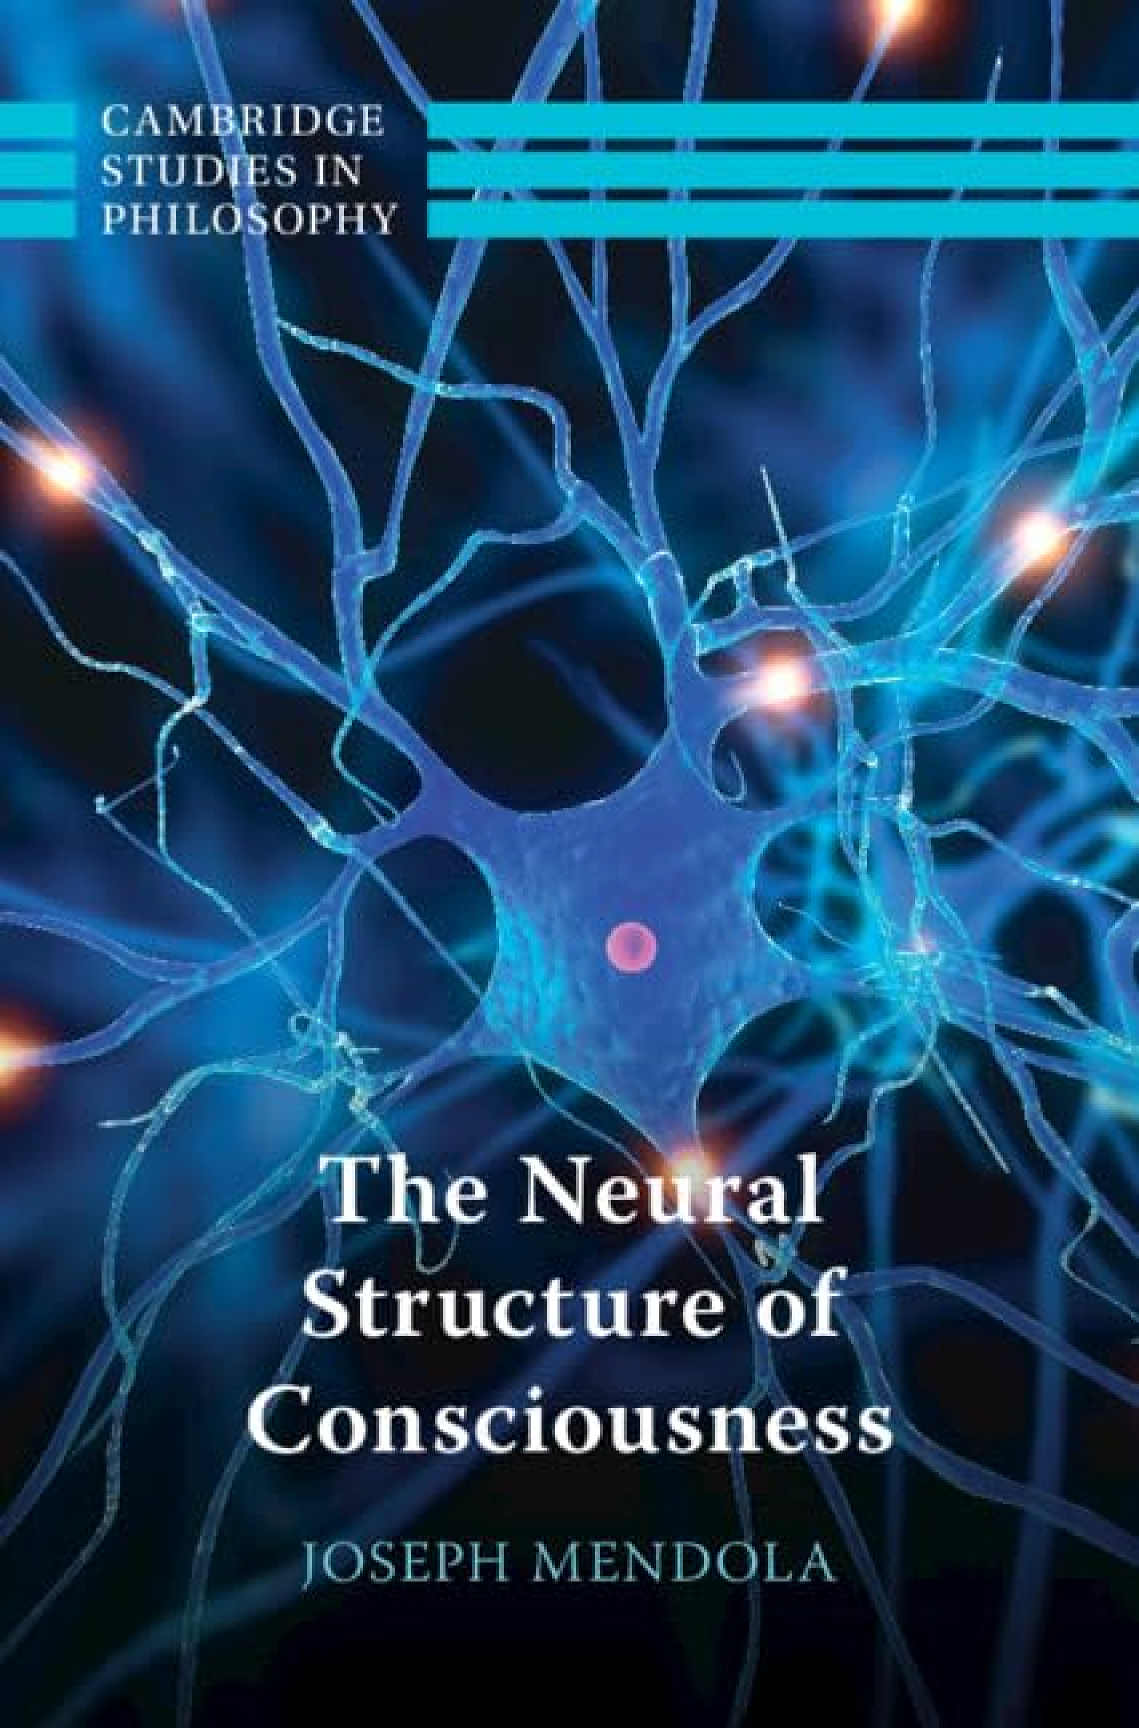

CAMBRIDGE
STUDIES IN
PHILOSOPHY

The Neural Structure of Consciousness

JOSEPH MENDOLA

THE NEURAL STRUCTURE OF CONSCIOUSNESS

Consciousness is an intriguing mystery, of which standard accounts all have well-known difficulties. This book examines the central question about consciousness – that is, the question of how phenomenal features of our experience are related to physical features of our nervous system. Using the way in which we experience color as a central case, it develops a novel account of how consciousness is constituted by our neural structure, and so presents a new physicalist and internalist solution to the hard problem of phenomenal consciousness, with respect specifically to sensory qualia. The necessary background in philosophy and sensory neurophysiology is provided for the reader throughout. The book will appeal to a range of readers interested in the problems of consciousness.

JOSEPH MENDOLA is Professor of Philosophy at the University of Nebraska–Lincoln. His previous books include *Goodness and Justice* (2006), *Human Interests* (2014), and *Experience and Possibility* (2021).

https://doi.org/10.1017/9781009514774

CAMBRIDGE STUDIES IN PHILOSOPHY

Cambridge Studies in Philosophy is the cornerstone of Cambridge University Press's list in mainstream, high-level analytic philosophy. It serves as a forum for a broad range of monographs on the cutting edge of epistemology, the philosophy of language and mind, ethics, and metaphysics.

Recent Titles

DAVID LEWIS *Papers on Ethics and Social Philosophy*
FRED DRETSKE *Perception, Knowledge, and Belief*
LYNNE RUDDER BAKER *Persons and Bodies*
ROSANNA KEEFE *Theories of Vagueness*
JOHN GRECO *Putting Skeptics in Their Place*
RUTH GARRETT MILLIKAN *On Clear and Confused Ideas*
DERK PEREBOOM *Living Without Free Will*
BRIAN ELLIS *Scientific Essentialism*
ALAN H. GOLDMAN *Practical Rules*
CHRISTOPHER HILL *Thought and World*
ANDREW NEWMAN *The Correspondence Theory of Truth*
ISHTIYAQUE HAJI *Deontic Morality and Control*
WAYNE A. DAVIS *Meaning, Expression and Thought*
PETER RAILTON *Facts, Values, and Norms*
JANE HEAL *Mind, Reason and Imagination*
JONATHAN KVANVIG *The Value of Knowledge and the Pursuit of Understanding*
ANDREW MELNYK *A Physicalist Manifesto*
WILLIAM S. ROBINSON *Understanding Phenomenal Consciousness*

https://doi.org/10.1017/9781009514774

D. M. ARMSTRONG *Truth and Truthmakers*
KEITH FRANKISH *Mind and Supermind*
MICHAEL SMITH *Ethics and the A Priori*
NOAH LEMOS *Common Sense*
JOSHUA GERT *Brute Rationality*
ALEXANDER R. PRUSS *The Principle of Sufficient Reason*
FOLKE TERSMAN *Moral Disagreement*
JOSEPH MENDOLA *Goodness and Justice*
DAVID COPP *Morality in a Natural World*
LYNNE RUDDER BAKER *The Metaphysics of Everyday Life*
MICHAEL J. ZIMMERMANN *Living with Uncertainty*
SANFORD GOLDBERG *Anti-Individualism*
CRAWFORD L. ELDER *Familiar Objects and Their Shadows*
JAMIN ASAY *The Primitivist Theory of Truth*
RICHARD FUMERTON *Knowledge, Thought and the Case for Dualism*
CHRISTOPH KELP AND MONA SIMION *Sharing Knowledge*
ARTŪRS LOGINS *Normative Reasons*
MONA SIMION *Resistance to Evidence*
JOSEPH MENDOLA *The Neural Structure of Consciousness*

THE NEURAL STRUCTURE OF CONSCIOUSNESS

JOSEPH MENDOLA
University of Nebraska–Lincoln

CAMBRIDGE
UNIVERSITY PRESS

https://doi.org/10.1017/9781009514774

CAMBRIDGE
UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom
One Liberty Plaza, 20th Floor, New York, NY 10006, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India
103 Penang Road, #05-06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781009514798

DOI: 10.1017/9781009514774

When citing this work, please include a reference to the DOI 10.1017/9781009514774

First published 2025

A catalogue record for this publication is available from the British Library

A Cataloging-in-Publication data record for this book is available from the Library of Congress

ISBN 978-1-009-51479-8 Hardback

For my grandchildren,
Aria and Aven Herrera

https://doi.org/10.1017/9781009514774

Contents

https://doi.org/10.1017/9781009514774

Acknowledgments

I am grateful for help from Aaron Bronfman, Tom Carson, Joey Dante, Jerry Fodor, Zack Garrett, Shane George, Rachel Goodman, Lily Griffin, Reina Hayaki, David Henderson, Harry Ide, Franklin Mathews Jebaraj, Brian McLaughlin, Abi Sears, Christopher Stratman, Adam Thompson, Vileru Tivexi, J. D. Trout, and three anonymous referees. And special thanks to Hilary Gaskin and Sandy Goldberg for their advice and support.

https://doi.org/10.1017/9781009514774

CHAPTER 1

Introduction

The topic of this book is phenomenal consciousness. Its thesis is that phenomenal consciousness can be explained by appeal to the modal structure of our neurophysiology. This provides a solution to the so-called hard problem of consciousness.[1] This chapter is an introduction to that problem and to the solution proposed.

1

Consciousness seems to many a great and evident mystery, something that may even require of its possessors a special metaphysical status as dualist or spiritual beings, outside the natural order otherwise well described by the physical sciences. To Aristotle, there seemed a great divide between the living and the nonliving. But to many today, there seems a great divide between living plants and sentient, conscious animals.

One reason for this historical difference is that Aristotle believed we live in a much more familiar and cozy world than we do. He thought that things out in the world were very much as we experience them to be, for instance that there are irreducibly green trees and irreducibly red apples. Color on his view was a fundamental feature of reality, not constituted by anything else. Any omniscient god would see pretty much the colors we see; any creator must have spangled color intentionally throughout the world. This is not to say that Aristotle thought our experience was always veridical. As a lemon recedes from us, it in some sense looks smaller. But this distortion in our experience was supposed to be due to the nature of the omission and reception of sensible species of colors that fly off the lemon and somehow converge on the eye, or alternatively was supposed, as Plato thought, to be due to differences in fanciful extramissive rays that emerge from the eye. But the color on the objects and the color we

---

[1] Chalmers (1995).

https://doi.org/10.1017/9781009514774.001

perceive were supposed to be much the same. If bats or the color blind have a different experience of the color of the receding lemon,[2] that is only, it was thought, because their perception misses or distorts things in the world that ordinary humans can see. If in the dark the lemon looks black, that is only because the sensible species cannot make it through the air unless the air is energized by light, so that the true color is blocked from us by darkness. Although some complications like these were recognized, according to Aristotle there is basically no gap between the way we experience the world and the way it is.

But we know better. We inherit a sophisticated science predicated in part on Galileo's thought that colors are mere names, or at least on the recognition that the colors we experience are not simple irreducible aspects of reality. There are other ways than Galileo's inflammatory way to think of this. Descartes thought that our perception of colors rather involved a kind of confusion of what was really there. Other near contemporaries, like Locke, thought instead that objective colors of lemons should be conceived as powers in objects to cause experiences in us. In any case, there is experience on one hand and there is the world on the other, and aspects of colors we most immediately experience, of which we are most immediately aware, of what we might call "phenomenal colors," are held to be merely in experience and not in the world itself. They are analogous in some respects to perceptual illusions. Some features of our sensory experience are, on this conception, misleading about the world itself.

We now generally think there may well be a gap between, on one hand, any objective features of reality that correspond to experienced colors and, on the other hand, the so-called qualia such color experiences involve, relevant to differences in what it is like for us to experience red, green, and yellow. But creatures for whom it is like something to experience things are creatures with consciousness. Hence we have more trouble with consciousness than Aristotle recognized. The problems of consciousness and of qualia are in this way closely entwined.

There are several reasons why we need an explanation of our experience that is not so commonsensical as Aristotle's, several reasons to think many of the properties and relations that concrete objects seem to present in our experience are, at least in large respects, not out there in reality. First, there are dreamers and hallucinating madmen and deluded perceivers, who see things that aren't there. Second, the world that our physics reveals is radically different from that our experience seems to present to us. In a

---

[2] Color blindness was in fact a nineteenth-century discovery.

strange world of curved space-time and quantum mechanical amplitudes, this problem is even more dramatic for us than it was for Galileo. But third, and most important from our perspective, there are very many detailed ways in which our sensory experience is not plausibly veridical even under the best circumstances, even when we are awake and sane and the light is good and microphysics is not in question. We will be accumulating many cases of this third sort as we proceed. But, for purposes of initial orientation, consider one cluster involving color.[3]

There are those who still say with Aristotle that the phenomenal colors we experience, the red of some umbrella and ball we seem to see, are in fact present in reality. And such objects do have in reality certain physical properties, certain spectral reflectances, which is to say certain tendencies to reflect various wavelengths of light in various proportions, that correspond reasonably well to specific perceived surface colors. When an object is perceived to have a certain hue when in white light by normal humans, then it is within a certain limited range of spectral reflectances in reality, a specific range of fully objective physical properties. So perhaps we might identify the colors we see with surface spectral reflectances of a certain range.[4]But this cannot suffice to deliver all the features of the kinds of colors we seem to experience, all the structure of similarities and differences among experienced colors. Even contemporary physicalists about color, who identify colors out in the world with surface spectral reflectance ranges, do not characteristically attempt to defend the implausible view that the entire structure of the color properties that we seem to see, the full nature of phenomenal colors, is present in objective surface spectral reflectances.

There are many relevant examples. Lights are seen as colored but involve no surface spectral reflectance. The sky is blue and has no surface. There are differences in spectral reflectance that are not visible to normal humans. Perhaps we might say that some of these are imperceptible differences in color, which is strange enough. Even so, there are metamers, which involve very different surface spectral reflectances, but yet present the same colors to those with normal human color vision under normal conditions of illumination, while yet they can be seen to be different in odd conditions of illumination. What's more, bluish purples have quite different physical bases than some reds, but seem to be quite similar colors. And there is the structure of unique hues in color experience – the specific red that is neither orangish nor purplish but rather pure red, and also analogous unique yellows, blues, and greens. These are privileged

---

[3] Hardin (1993).

[4] Hilbert (1987).

phenomenal colors that correspond to no objective privilege in reality. Furthermore, because infrared and ultraviolet radiation exist but are invisible to humans, there is nothing about surface spectral reflectances that explains why there are just the colors that we see.[5]

Phenomenal color – the color with all the structure that we see – is not in fact out there in reality. No plausible account of the color we experience can have the general form favored by Aristotle and untutored common sense. So one aspect of the hard problem of consciousness is how to deliver our color qualia, the specific what it's like of our color experience, in our at least mostly physical world. The peculiarities of humans plausibly matter to this. The specific way it is like for us to experience some flying colorful bug is plausibly quite different from the very foreign way it is for some bat to experience the bug through its sonar.

I will attempt to show here that our human neurophysiology of color vision accounts for our human color qualia when that neurophysiology is understood in the particular way I develop, which I call the modal structural way. There are of course other senses, and there is also intuitively much more to the hard problem of consciousness than just the problem of sensory qualia. There is a second key difficulty about consciousness to which we should attend, and which requires other aspects of my proposal. Even if some mechanical device could in some way represent a bug in just the way our qualia, or the different qualia of a bat, suggest, it wouldn't follow immediately that it was conscious, that it was like something to be that device.

And on this issue, I think that Aristotle, with his focus on animal life as the arena in which consciousness arises, has something still to teach us.[6] Animals sense the world and act on the basis of what they sense, and their perception and action is mediated by their neurophysiology. It is specifically in such a locus of animal life that I believe that our neurophysiology of vision constitutes our experience of phenomenal color. Whether or not animal life with a neurophysiology is strictly necessary for phenomenal consciousness of any sort, still such life includes all the obvious cases, and it plausibly constrains at least our human form of it. Modal structure, as we will see, is also involved in this aspect of our phenomenal consciousness.

That's the very short version. Before I elaborate the solution I propose to the hard problem of consciousness, explain modal structure, and sketch the plan of the book, it will be useful to further refine our topic and

---

[5] Hardin (1993); Palmer (1999).

[6] This elaborates a theme of Godfrey-Smith (2016b) and E. Thompson (2007).

problem, and introduce some standard terminology in which it is now customarily posed.

2

Humans have some relatively unique capacities, and what it is like to be us is in some respects inflected by these capacities. I will be concerned here with what explains what it is like to have our human experience. But I will be specifically concerned with our perceptual or sensory experiences in particular, which is the sort of thing that other animals also plausibly possess. I will not be concerned for instance with how we experience our own thoughts and emotions and desires, which may require something fancier. Only human sensory experience is our target.

I will of course be yet more specifically concerned with *conscious* sensory experience in particular, and not for instance any unconscious analogue. And still more exactly, I will be concerned with the sort of *phenomenal consciousness* involved in conscious sensory experience. And while we will attempt to understand the *qualia* involved in our sensory experience, that is only in one sense of that ambiguous term, which is closely linked with phenomenal consciousness in particular. So let me now explain better what this jargon means.

Things are said to be conscious in several different ways. We say that animals are conscious or aware of their environment in a way that rocks and plants are not. We say that animals who are asleep or knocked out are not at the time conscious, in a perhaps more demanding sense. So there is at least one sense in which any awake animal who has sensations or perceptions of an environment is conscious.

But we pursue here a kind of consciousness that is more specific and seems more demanding than merely that. We are interested in the *conscious* sensations of these awake animals, of which they are in some sense aware. The horrible sensation felt by a squirrel when a limb is ripped or bitten off is something of which they are plausibly aware in this way, whereas they might even while awake have an unconscious perception of some subtle conspecific pheromone of which they are not aware in this vivid way.[7]

On the other hand, there are related kinds of consciousness that are more demanding still, at least on the surface, than what we pursue here. Sometimes we are quite explicitly conscious or aware of our own mental states in a very articulate way, when we report them to ourselves, in words

---

[7] We will eventually see reason to doubt this common distinction.

we say out loud or to ourselves. Sometimes we also seem to enjoy a special kind of introspective access to their nature, and sometimes we are even explicitly aware of ourselves as thinking and seeing selves, with these particular mental states. But these are kinds of consciousness or self-consciousness that are fancier, at least on the surface, than our focus. They are plausibly beyond the capacity of the squirrel.

Since Freud and the development of cognitive science, we generally believe that we have states of awareness of things in the world of which we are not and cannot be aware in any explicit and introspective way.[8] But even some of which we are aware of in a way relevant to our focus, we are not aware of in the fancy ways noted in the last paragraph. Think again of the pains of an unsophisticated animal, say some robin or cardinal. It seems aware of them in some relevant sense, even though it cannot reflect and introspect in any very fancy way upon them, even though it cannot talk about what it feels or ruminate about the fact that it feels them. That is the kind of thing we are after.

It is like something to be the pained bird or squirrel. That is the kind of consciousness that we seek to understand. Pain of course may seem a sensation very unlike the sensation of color. But my point is that we are interested in a kind of human conscious experience of color and other sensory properties that has the same rough character as a bird's or squirrel's conscious experience of its pain.

There may be philosophical arguments linking consciousness in one of the less demanding or more demanding senses noted to consciousness of the sort we are centrally concerned with. But on the surface, what we seek is in between. There is something it is like for an animal to have sensory experiences of certain sorts, for instance to have certain views of bright yellow objects. That is our focus.[9] It is possible that there are unconscious perceptions of environmental features by awake animals that lack this vivid sort of consciousness. And, on the other hand, it seems that creatures can have the intermediate sort of consciousness we seek, say creatures in pain or sensing yellow things, even though they do not possess, and indeed are incapable of, explicit introspective awareness of their own mental states as such, or of themselves as thinking and seeing selves, and even though they can't talk.

---

[8] Newell and Shanks (2014) is a skeptical review of psychological evidence for unconscious influences on decision-making. But Wiltshire (2014) is an examination of unconscious motives in Jane Austen!

[9] Perhaps there is a way it is like to believe, or even to be me, which involves no particular qualia. See Kim (2011: 273–277). But we are focused on what it is like to have a sensory experience.

This in-between sort of consciousness on which we are focused, enjoyed by awake animals when it is like something for them to feel a pain or see a yellow thing, is called “phenomenal consciousness.”[10] Differences in what it’s like, differences in phenomenal consciousness, are called differences in “qualia,” in the sense of that term I will use.[11] As I use the term “qualia,” there is no implication that qualia are properties of sense data, that they are nonrepresentational properties of experience, that they are nonphysical, or even that they are nonveridical aspects of experience although I have already argued that in some respects they are. And the term is sometimes taken to imply those things. But I deploy a sense of the term that does not imply presumptions on these issues, that is focused on the details of what it’s like however it is in fact constituted.

So now I can put our topic in this more exact way: We are trying to account for the qualia involved in our sensory experience, for that type of human phenomenal consciousness, particularly in a world that is, as we will see, very unlike the one we seem to experience.

But to focus and situate our topic still more closely, it will be useful to introduce a little more standard terminology.

Today, many of us think Galileo was basically right about the world. Nevertheless, quite a few philosophers still favor the sort of commonsense conception of perception and the world that Aristotle fostered. It is natural to wonder whether this is a coherent set of views, whether Galileo really fits with Aristotle in this way. But in any case, we inherit the complex problem of fitting this together and working it out. And in the current climate of opinion, some distinctions must be kept in mind.

There is an apparent division between, on one hand, the phenomenal consciousness or qualia involved in sensory experience, for instance color qualia, and what is called the “representational content” of sensory experience, on the other hand. There is this apparent division even when we are restricting our concern to the phenomenally conscious sensory experiences of awake animals. The representational contents of these experiences are what they are about out in reality, and are relevant to whether they are veridical or not. In a Galilean world, that is unlikely to be Aristotelian colors. But we are interested here in the qualia, in the what-it’s-like, in the relevant sorts of phenomenal consciousness, in the sort of Aristotelian color that we at least naively seem to experience. The general idea behind the distinction at hand is that it is like something for a human to experience a bright green car in good light, but that that what it’s like, those

---

[10] Block (1995).

[11] This is the original sense specified by Peirce (1982).

qualia, that phenomenal consciousness, may well not be exhausted or fixed or even affected by the way that experience represents the green car in the world to be, by the representational content of that experience. Something like this distinction is required if we are to insist that our sensory experience represents the world more or less accurately, and that the way qualia are present to us in consciousness can lead us to mistake the true Galilean world for a fanciful Aristotelian world.

Such a distinction between qualia and representational content in sensory systems is not uncontroversial. In fact, I myself believe that the what-it's-like and the representational content of a sensory experience are closely related, in the sense that the qualia fix the representational content. And a much larger number of philosophers believe that there is a dependency that runs the other way, that the qualia are fixed by what our senses truly represent, and so the distinction collapses in the other direction, despite the difficulties for such a view I have already noted. Still, the relevant point for now is that neither of these views can be presumed in our current philosophical context. If anything, a distinction between qualia and representational sensory content is currently presumed. So while I stress that my use of the word "qualia" will itself make no presumptions on this issue at all, that it is focused on the what it's like of sensory experience, whether or not that should be distinguished from sensory content, still, nevertheless, there are many well-known cases that are supposed to support the now popular distinction between qualia and sensory content, cases that we will need to bear in mind as proceed. The cases can be divided loosely into the fanciful and the real.

Many hold that there might be so-called color inverts, those who have the same sort of qualia and conscious sensory experience when they see, for instance, red things that color normal people have when they see green things, with corresponding inversions throughout color experience. And it is also often held that color inverts still have sensory experiences with the same veridicality conditions, with the same representational contents, as the color normal. The inverts only differ from the color normal in qualia, in the what-it's-like, not in sensory content, on such views. Beyond color inverts, there might be those on Inverted Earth, who are presumed to be individuals who generally have the same qualia but different sensory contents.[12] And philosophical zombies are alleged to be physical duplicates of us who lack any what it's like, whatever they say, and so experience sensory representational contents but no what it's like. There are also

---

[12] Block (1990).

analogous stories about fading or dancing qualia that do not affect sensory contents.[13]

Maybe these fanciful stories are relevant in some indirect way to what is actual. But in any case, there are also real cases where there is a gap between qualia and the world, and so between qualia and at least any *veridical* sensory content. First, and most important from our perspective, there are numerous phenomena involving human sensory experience that are analogous to those relevant to color experience that I have already mentioned, for instance the unique hues. We will consider many such cases as we proceed. Second, some people who are color-blind must have a different what-it's-like in at least certain cases from the color normal. It is not plausible that those who are wholly red-green color-blind could miss such a vivid difference in their qualia if it was really there for them. But if we are prepared to call their perceptions veridical, then the what-it's-like must be distinct from their sensory content. Even if we aren't prepared to call their experiences veridical, color experience provides other analogous cases. Even different color normal people disagree somewhat about which objects in the world have unique hues. We will return to such cases as well.

Third, there are mundane cases like this: Even if a tree doesn't really look smaller as it recedes from you, there is some difference in what it's like to see a distant or close-up tree of the same objective size.[14]

To repeat, we are interested here specifically in the what-it's-like of sensory experience, in sensory qualia and phenomenal consciousness, which, for purposes of discussion, we are granting may be distinct from sensory content, although we are not presuming that. Still, we will consider views that collapse representational sensory content and the what-it's-like. Moreover, as I said, I myself think one idiosyncratic version of the collapse view is correct. I think that the world is very unlike the way we experience it to be, because I think the way we experience it to be is closely tied up with our qualia, while our qualia and the real world are not in close alignment. But I will not presume that idiosyncratic conception here, though we will be gathering some grounds to believe it.

The main point for the moment is that, despite these various differences, more or less everyone in the current climate can agree that the qualia of conscious sensory experience require some explanation that it isn't easy to provide. The old Aristotelian story doesn't work, at least in that respect. That much should be reasonably uncontroversial. So that will be our organizing focus despite what I admit is my desire to eventually deploy

---

[13] Chalmers (1996).

[14] Peacocke (1983).

such qualia in other and more ambitious ways, to explain the representational content of sensory experience, and despite the fact that we will have to attend somewhat to the popular views of those who think the dependency runs in the opposite direction. We will be trying to solve this crucial portion, relevant to sensation, of the hard problem of consciousness.

There is another restriction on our focus of concern worth noting up front. There are debates about the richness of sensory content and hence about relevant associated qualia. On one traditional conception, the contents of sensory perception are limited to the traditional characteristic objects of the senses, for instance, in the case of vision, to shapes and sizes and colors and motions, and to smells and tastes and sounds and tactile feels in the case of other senses. But a competitor view is that fancier properties such as being an oak tree, or even a written word of English that means tree, might be present in the most immediate way in the perceptual experience of someone who is trained in the requisite way, in a way that affects their relevant phenomenology.[15] As will eventually be evident, I have considerable sympathy with this more expansive second view, and think that the general model we will develop here might be extended profitably to cover its favorite cases. Indeed, it is arguable that the ultimate viability of my proposal may depend on that extension. But without presuming to resolve this dispute, we will focus here principally on the traditional sensory properties.

3

So that is our problem. Now consider possible solutions. While our major focus will be on building a positive plausible account of human sensory phenomenal consciousness, it is worth mentioning for purposes of orientation who I think is wrong about it. Among philosophers, I think almost everyone is. While I can't aspire to disprove all standing accounts in any brief space, it is worth mentioning obvious difficulties of the major standing views for which there are no obvious solutions, by way of motivating another sort of attempt. If you don't feel the need for such motivation, which of necessity will get down into the weeds a little, you can skip to Section 4.

There are two clusters of solutions that have some contemporary currency. Some deploy naturalistic accounts of representation, and others resort to dualism or analogs like panpsychism. Begin with the first cluster.

---

[15] Siegel (2010).

Talk of representation and information is ubiquitous in contemporary biology and psychology, and also in contemporary discussions of sensory qualia and sensory content.[16] There are thought to be mental representations, neural analogs of words, which represent certain things because they carry information about those things, paradigmatically because they are caused by those things in such a way as to assure covariation.[17] That's the basic idea, but there are necessary refinements. Representation plausibly requires something slightly more complicated than actual causal covariation, because of the following worries: When something is caused, it generally has more than one cause, in a chain. What's more, false representations and nonveridical experiences are not usually caused by what they are about. The standard reply to these standard worries specifies that the relevant cause, the one that constitutes for instance the sensory content of a sensory representation, is the "normal" cause of that mental representation. When false, or when it has many causes, a mental representation is caused by other than its normal causes, but it is only the normal cause that counts as what is represented, as its "content," about which the mental representation is said to carry information. Then we face the question of what a normal cause is. And there are various standing accounts of normal causes. There are teleological accounts, which focus on what causes something should have, in some sense of "should" allegedly delivered by evolutionary or learning history.[18] Or perhaps a normal cause is a statistically normal cause.[19] Or perhaps it is specified by asymmetric dependence so that the normal cause is the cause on which all the other causes depend while it depends on no other.[20]

I argue at length elsewhere that none of these proposals can deliver mental content of any sort,[21] and I can't repeat all that here. So I will grant here for purposes of argument what I do not believe, that the currently dominant informational accounts of representation can be made viable, in the sense that they can provide some plausible account of normal causes that undercuts standard objections. And I will make no objection to their use in biology or psychology. And I will even grant that they might deliver sensory content of some sort.

Nevertheless, focus closely on whether, even granting all this, they can help much with sensory qualia and phenomenal consciousness. There is

---

[16] Dretske (1981) is seminal.

[17] There are variations on this idea. Perhaps a mental word must be decoded or read to represent.

[18] Millikan (1984). [19] Harman (1973). [20] Fodor (1990).

[21] Mendola (2008: 103–172).

certainly reason to be initially doubtful. As we have already noted, many features of sensory qualia do not correspond to objective features of external objects out there in reality, so cannot be normal causes. We see objects as unique red, and as having purples similar to reds, but that sort of structure is absent in any physical correlate of color out there on the objects. The tree looks in some sense larger when it is merely closer, but it isn't. In fact, cases like these are bound up in the very notion of what sensory qualia are. And we will be gradually collecting many cases of this sort. A blunt way to put this criticism is that qualia are nonveridical, whereas informationally based representation must be dominantly veridical. There are no things in the world corresponding to many qualia for mental words to informationally represent.

Nevertheless, information-based accounts of sensory qualia have a variety of ingenious resources to deploy towards defusing these initial worries. For instance, so-called representationalist views maintain that, despite the surface cases we have considered that suggest a distinction between the qualia associated with sensory states and their sensory contents, still in fact sensory contents constituted by causally based information do fix the relevant qualia. So despite the initial difficulties of representationalist accounts of phenomenal color and other sensory qualia, you may want to consider some of these resourceful details.

There are two rough forms such a representationalist reply may assume. The first focuses on more objective properties of external objects, and the second on more subjective properties of objects, that involve their relations to the perceiver.

Tye's account is a well-developed objective version. We might distinguish two periods of Tye's work on color: Tye A and Tye B. Tye A, in 1995, deployed a physicalist account of color experience apparently modeled on suggestions of Hilbert about color.[22] There are cone cells in the retina that are maximally responsive to short, medium, and long wavelengths of visible light. And our perception of color depends on light being reflected from objects. So "the color of a surface is an ordered triple of the reflectances of the surface with respect to light in these three wavelength bands . . ., where the reflectance of a surface at a given wavelength is its disposition to reflect a certain percentage of light at that wavelength."[23] But the problem is that there are standard objections to this as an account of our experience of phenomenal color, of color qualia. Metamers, you will recall, have different spectral reflectances but look the same in normal

---

[22] Hilbert (1987).

[23] Tye (1995b: 146).

light. Tye A thought that this worry can be handled by specifying relevantly wide bandwidths.[24] Though this is overly optimistic, let it pass. More importantly, there are relations of similarity between colors that color objectivism of this sort seems to miss. The standard hues form a similarity circle, but light frequencies do not. Tye A says, in response, that we are to think of color space as three dimensional, with each dimension corresponding to the surface reflectances at one of the three wavelength bands, and to think of the relevant triples of reflectances as coordinates in this space, with the hues marking out a closed circular loop in color space.[25] But this is fanciful. It is true that any test light can be matched by three primary color lights whose intensities scale a three-dimensional space in which ordinary hues mark something like a circle. But this is not a color appearance space that captures similarity of phenomenal colors.[26] Another standard worry about such an account of color qualia is that there is a phenomenological distinction between unique hues like pure red, blue, yellow, and green, and binary hues like orange, which seems yellowish and reddish. Tye A says this: "As for the binary-unitary distinction, it can be preserved as a basic truth about color mixing. Orange, for example, is the color you get when you mix red and yellow pigments; but red is not the color you get when you mix purple and yellow pigments."[27] But green can be got by mixing yellow and blue and still there is unique green, so this story doesn't work.

Tye has moved on.[28] "Opponent processing" is a feature of our neurophysiology of vision to which we will return. According to Tye B in 2000, an object is pure red if it has a surface spectral reflectance that normally produces opponent-processing that is distinctive of the experience of pure red.[29] And an orange object generates forms of opponent-processing characteristic of being seen as yellowish and also forms characteristic of being seen as reddish.[30] On such a view, color qualia present us with these subjective features of human visual processing. This account is much less objective than that of Tye A, and faces characteristic difficulties for subjective accounts to which we will shortly turn. There is no explicit account of the color similarity circle in Tye B, but it is presumably supposed to be connected to this proposal, and hence similarly subjective. And there remains a general problem for even such a very limitedly objective account of color qualia. The apparent colors of objects shift somewhat as the light changes. "The colors of objects typically do not

---

[24] Ibid., 147.
[25] Ibid., 148.
[26] Byrne and Hilbert (2010: 261–264).
[27] Tye (1995b: 148).
[28] Tye (2000).
[29] Ibid., 160–161.
[30] Ibid., 163–164.

change when they are moved from outdoors to a setting illuminated by incandescent lights, for example," according to Tye.[31] But of course there is some difference in our phenomenal experience introduced by this shift. Tye has no account of this.

It is not a surprise that a developed representationalism tends in a subjectivist direction, since qualia are frequently nonveridical artifacts of perception. A more uniformly subjectivist representationalism is developed by Hill. He identifies qualia with properties of external objects that are more subjective than most of those that Tye deploys. He calls the properties in question "appearance properties" or "A-properties." Qualia, on his view, are A-properties.[32] But which properties in particular are A-properties?[33] Color physicalists often identify objective colors of solid objects with certain surface spectral reflectances. But since color appearances change with the light, this is too objective for A-properties. Perhaps, then, we should say that when an object presents a red appearance to you, the proper focus is what is going on in your retina, that the "object is projecting a retinal image that involves appreciably more activity in long wavelength cones than in middle wavelength cones, and about the same amount of activity in short wavelength cones as in long wavelength cones and middle wavelength cones combined,"[34] that they have a "red projection property" of this sort. But this neglects the effects of color constancy in our visual processing. It is too subjective. So Hill believes that we should say instead "an object counts as A-red if (i) it has a red projection property, (ii) appropriate constancy transformations have been applied to this property, (iii) the output of the transformation falls within a certain range."[35]

But the problems for such a subjectivism, in the context of a representationalist, information-based account of qualia, are manifest and debilitating. These fancy complicated A-properties seriously misrepresent the relatively simple and straightforward qualitative content of our sensory experience. Color properties seem to objectively qualify external objects themselves; experience of such properties does not seem to be about our own cones even in part. What's more, it is hard to believe that information-based accounts of representation can deliver such complicated and recherché contents.[36]

---

[31] Tye (1995b: 146).

[32] However, in an account of pain qualia, these complexities fade in favor of a more objectivist account. See Hill (2009: 169–191).

[33] For reasons to think they aren't properties, see Egan (2006).

[34] Hill (2009: 168 fn 31).

[35] Ibid.

[36] I do not claim that I have adequately argued here that all representationalist accounts of all sensory qualia are implausible. Neander (2017: 171–174) is perhaps the best case.

Hill has a resource in response to at least the first of these worries. He distinguishes between visual qualia as represented and visual qualia as they are in themselves.[37] Visual qualia are appearances, but appearances that appear to one in a somewhat misleading way.[38] But if it is the appearance of an appearance that really matters to intuitive visual phenomenology, we have in fact moved into the next class of views we need to consider.

On these other views, it is the representation *of* mental states, for instance states of sensory experience, that is supposed to deliver phenomenal consciousness. According to these higher-order representational accounts of consciousness, an experience is conscious when it is represented by a higher-order mental state, a mental state about a mental state. To avoid a regress, the higher-order state need not itself be conscious. There are different forms of this view. In the opinion of Armstrong and Lycan, what is involved is a kind of higher-order perception.[39] In the opinion of Rosenthal, it is higher-order thought.[40] It might also be that phenomenal consciousness merely requires of a state the capacity to be accessed in one of these higher-order ways rather than actual access of that sort.[41] That would be a dispositionalist higher-order account.

There are significant worries about both higher-order thought and perception, or about dispositions to them, as resources necessary for phenomenal consciousness, since young children and animals intuitively have phenomenal consciousness without such apparently fancy mental states. But focus on the main issue for us, whether they can plausibly deliver sensory qualia. It is striking that no higher-order view has so far provided a reasonably well-developed account of the different natures of different first-order sensory qualia, which is our central concern.[42] How might higher-order representation conceivably help to deliver first-order qualia that differ from the representational contents of first-order perception? On the surface, sensory qualia seem to involve features of the world itself, and it is not obvious how representation of our mental states of experiencing can deliver this aspect of the qualia without presuming that those mental states of experiencing themselves already involve the very qualia in question. But there seem to be roughly three possibilities. First, what is represented at the higher level might be the content of the

---

[37] Hill (2009: 146). [38] Ibid. [39] Armstrong (1968); Lycan (1987, 1996).

[40] Rosenthal (2005). [41] Carruthers (2000).

[42] There have been suggestions by Rosenthal, but they are not specially integrated into his higher-order thought view of consciousness.

lower-level state. It might seem that this possibility would not allow the higher-order representation to do any work beyond representationalism itself in determining the relevant sensory qualia, although perhaps a kind of partial or misrepresentation of first-order sensory qualia might sometimes be a useful resource. Second, what is represented at the higher level might not be the content but the intrinsic nature of the first-order representational state. Or third, what is represented might be the representation by that first-order state of its first-order content. The first model might help with metamers, where there is a difference in what is represented out in the world but not in phenomenal consciousness. It is, however, hard to see how it can deliver the structure of unique hues, or adjust our sense of the content of sensation as much for instance as Hill generally requires, because orthodox information-based accounts of representation are not suited to deliver significant and normal misperceptions or analogous errors in representational states. They can abstract in representation from what is there, but have serious difficulty in adding to it in any general way. But it might seem that in that type of case, our second model might help. It might be that there are intrinsic differences among the mental states that are representations of unique hues as opposed to other hues. But it is still a puzzle how representations of internal mental states can be as of the world itself, how one could see, for instance, differences in neural processing as differences in worldly colors. And our experience of colors seems to mix together two different sorts of factors that the first two models handle differently, without there being any obvious way to understand how in our color qualia they might be conjoined. And most significant, such a focus on the intrinsic properties of the first-order representing states suggests that that is really the key to qualia differences, not how they are represented at a higher-order. But perhaps the third and itself mixed model of what higher-order representation represents can evade this worry, and also explain how the factors characteristic of the first and second model might be mixed. Perhaps what is represented at the higher level is the representing by the first-order state of its first-order content. But it is very difficult to see how a causal-covariation story about representation could deliver such a complex and modally complicated content to the higher-order state in question.

Our first two attempts to deploy information-based representation to deliver sensory qualia seem to face insuperable difficulties. Third attempt: Some views of consciousness with current defenders hold that it involves mental self-representation. The idea is that conscious states of sensory experience involve self-representing states. But it is hard to see how this would help with the particular worries about representation and qualia that engage us. And in Kriegel's relatively well-developed version, self-representation is not supposed to determine the differing natures of specific sensory qualia, which are instead given a representationalist account.[43]

We have been focused on accounts by philosophers in which representation and information play crucial roles. But there are also neuroscientific accounts to which such notions are crucial. Still, these somewhat analogous proposals about the nature of phenomenal consciousness also don't seem very well positioned to tell us what constitutes the differences among sensory qualia. For instance, the global workspace theory is that there is a kind of central clearinghouse into which information from various neural subsystems is gathered, in which it is subject to unified processing and from which it may be broadcast throughout the nervous system, and holds that being accessed by that central neural clearinghouse is what constitutes consciousness.[44] The integrated information theory holds that it is a kind of integration of informational representation that constitutes consciousness.[45] I think we should be cautious about what it means for the relevant sort of information to be gathered, broadcast, or integrated in these ways. But, suspending complications regarding normal causes, it would be possible to transmit causally based representation down any chain of representers that is locked together in suitable covariation, and so such causal influences might be gathered and broadcast in a way that makes some sense of these ideas. So, maybe, such a proposal can help isolate consciousness of some sort. But still, it is hard to see how such locked covariation could help make any special sense of qualia differences and similarities, how it provides sufficient help for us in that crucial role. These neuroscientific accounts inherit the difficulties for information-based views we have already discussed.

We have been focusing on the difficulties for representation-based views of phenomenal consciousness in delivering human sensory qualia. But another relevant objection is that at least the philosophical versions fail to explain why only animals are intuitively phenomenally conscious, since representation is ubiquitous and consciousness intuitively is not. Even evolutionary teleology encompasses plants. So let's move on to our second cluster of familiar views, which in some forms have at least a crude, brute-force answer to this worry.

---

[43] Kriegel (2009).

[44] Baars (1988, 1997). For alternative "cognitive" models of consciousness, see McGovern and Baars (2007).

[45] Tononi (2007, 2008).

This cluster deploys ghostly, spiritual, dualist stuff, which is not constituted by the neurophysiological or even by the physical, at least according to standard understandings of those things. So it might be thought that phenomenal consciousness in general is, or at least requires, such a general type of dualist stuff, and it might be thought that different qualia are constituted by different forms of it. When dualist stuff is proposed, it is customarily proposed to fill both of these roles. There are two basic forms of dualism relevant to phenomenal consciousness, although some partisans of the first do not accept the dualist label. And at least the second form might plausibly allow that only animals enjoy phenomenal consciousness.

The first form is panpsychism. It is the view that everything has a phenomenal character, down to the smallest bits of matter, and that the phenomenal aspects of the bits of matter that make up your physical aspect, or some of it, somehow sum up to constitute your phenomenal consciousness.[46] Everything is phenomenally conscious on such a view. That last point is an intuitive objection to the view. But such a view still might provide an interesting story about qualia differences.

Still, there are serious problems with panpsychism in its crucial role as an account of qualia similarities and differences. Since we know our physical aspect is constituted out of a limited number of types of molecules with a fairly limited chemical basis, and also that atoms are constituted by protons, neutrons, and electrons, it would seem that the phenomenal aspects of protons, neutrons, and electrons would, by necessity, play at least a relatively basic role in such a story about our qualia differences. Or perhaps we should proceed on down to quarks. But there are debilitating difficulties regarding how the relevant summation of phenomenal aspects of such parts into those of relevant wholes would work. This is called "the combination problem."[47] Simple addition seems obviously inadequate, and there are no obvious coherent and metaphysically unproblematic alternatives.[48] Goff has distinguished various aspects of this problem:[49] The "subject-summing problem" is how little subjects like electrons and quarks can sum up to specific big subjects like you and me. The "structural-mismatch problem" is how and why the structure of our phenomenal consciousness matches the structure of our psychology in particular and not our more detailed physical structure. And the "palette problem" is how the variegated differences and similarities among sensory

---

[46] G. Strawson (2006).

[47] For discussion, see Seager (2016: 305–330).

[48] For discussion of emergentist and infusionist alternatives, see Seager (2016).

[49] Goff (2017: 165–219).

qualia can be introduced by any sort of summation of basic phenomenal aspects, when experiences of such different qualia are constituted by the same limited sorts of microphysical particles with the same limited types of phenomenal character. That relatively fundamental phenomenal character would be so ubiquitous as to be apparently irrelevant to intuitive qualia differences of sensory qualia. The palette problem is especially damning in our context. As far as I can see, there is no remotely plausible solution available.

So a second type of dualist story about phenomenal consciousness, a more traditional form, may seem preferable to the first. Such a view says that there is an irreducible dualist phenomenal aspect for only certain sorts of physical structures, which of course might be only animals. And a traditional dualist is able to posit exactly the dualist aspect that delivers the details of our qualia.

But this view still faces severe difficulties. Why think that in all the vast universe some extra fundamental metaphysics is required to account merely, at least for all we know, for one psychological feature of a small portion of the biomass that exists on a very thin layer around the third planet of a mediocre star on the edge of a mediocre galaxy? We aren't that important. Nor are all the earthly animals. And why does it have just the structure it has, which is correlated with our neurophysiology in just the detailed way that it is, so we can talk accurately about our phenomenal consciousness and it matches our behavior in the various ways that it does, when our talk and behavior seems obviously directly explained by our neurophysiology. There are no other fundamental correlational laws in the world as we understand it that have the same kind of inelegant complexity as the correlational laws required to explain this.

There are motivations for dualism beyond metaphysical pride and theological commitments. There are dualist arguments that physicalist resources such as I will deploy here can be seen to be incapable of providing an adequate account of phenomenal consciousness. We will return to these arguments at the end of the book, when my positive proposal is available for scrutiny. And it may be that the various surface difficulties I have reviewed in this section for standing accounts of phenomenal consciousness are merely that, that they can be suitably answered by fans of dualism or naturalist representation as accounts of phenomenal consciousness. But these difficulties suffice at least to provide motivation to look for something else, a physicalist account of phenomenal consciousness that deploys a mechanism other than information-based representation. That is my goal.

4

My positive account of phenomenal consciousness in human sensation has three key features. Its first two features can be indicated quite briefly, although their proper development and explanation will require Chapters 2 and 3. We will begin here with these. Its third feature is perhaps most unfamiliar and characteristic, and will require a bit more initial explanation.

While our central focus is the phenomenal consciousness involved in human *sensory* experience, still one good intuitive test for something being a thing that it is like something to be, a being with phenomenal consciousness, is that it feels pains or itches, even if those aren't always considered, like vision and hearing, cases of sensation. It is certainly like something to feel a pain or itch. And it is quite natural to think that, in the actual living creatures who are the entities we are most intuitively certain feel pains or itches, it is the way that their nervous systems actively mediate between stimulations of certain sorts and actions that somehow constitutes such experiences. We don't intuitively expect such sentience in a tree stump or a rock, although they may register information in analogous ways and might even conceivably have dualist substance attached.

And a part of my proposal about our sensory phenomenal consciousness, in fact two of its three main parts, rests on considerations like that, on a certain understanding of the sort of living neurophysiology that mediates in us between stimulation and action, an understanding that focuses on what I will call the modal structure and dynamics of all that, and also on an analogous understanding of the phenomenal consciousness that it plausibly constitutes. We will develop first a certain modally inflected understanding of our nature as living beings with unified nervous systems, and, second, a certain modally inflected understanding of our phenomenal consciousness as a whole, and see, pushing at once up from below and the neural facts and down from above and the phenomenological facts, or, if you prefer, in from the third-person facts and out from the first-person facts, how the two sides fit together, how we should expect such a live human animal to possess a human phenomenological consciousness.

Those are the first two key features of my proposal. But you may well wonder what I mean by a modally inflected understanding of these things. And perhaps the best way to provide an initial, rough understanding of that, as well as of my overall positive proposal, is to focus in a more extended way on its third key aspect, my modal structural proposal about sensory qualia details and differences. Briefly put, but a little less briefly than before, the account of phenomenal consciousness to be developed here is that the modal structure of our sensory neurophysiology explains our conscious sensory qualia, largely because the actual modal structure of that neurophysiology explains the apparent modal structure of those qualia. Those details and differences of sensory neurophysiology and of sensory qualia match.

So let me sketch these ideas. Begin with the central notion of this book, the idea of modal structure.

Modality involves what is possible and necessary, what might be, and what must be. I mean something very general and inclusive by "modal structure." I count any modal aspect of anything as one kind of modal structure. Different things involve very different sorts of modal structure, but are rarely without some modal aspect that is significant to what they are. Some tree house might be painted in a variety of colors. Perhaps through some special technology or magic, it might even fly or shrink. But it could not be a red corpuscle or Notre Dame de Paris. It could not survive certain alterations, but it could survive others, and that modal nature is part of what such a thing is. I also count any dispute about the nature of what makes modal claims true as a dispute about modal structure.

Consider this case: A blood cell is some vivid red, or, if such tiny things aren't really red, a cherry or child's ball is so. Presume for the moment such a color is just as Aristotle assumed. It is then a fundamental ontological entity, which yet has essential relations and differences to other colors, colors the ball might sequentially exhibit. I also count such essential similarities and differences of that color, similarities and differences it must have to be the color it is, as relevant to its modal structure. They are part of what I mean by "structure" in that phrase.

That general characterization of modal structure may well remain opaque, and crucial aspects of the notion will gradually unfold throughout the book. But perhaps the best way to convey the flavor of what I mean by this notion in a helpful preliminary way is to focus on some examples of *shared* modal structure.

Presume a perfectly scaled desktop model of some town, deployed in some world's fair or war game, or a hugely enlarged model of some living human animal. Each model bears a very special sort of isomorphism with its original. It is not just that to each element in the model there corresponds some element of the original, with properties and relations except spatial distances in fact strictly duplicated, but that even essential similarities between properties and relations are duplicated. One building in the desktop model is just as similar in color to another building in the model as the original of the first is to the original of the second. And the heart of the body facsimile also consists of chambers arranged in the same way but on a smaller scale.

But there is an important aspect of shared modal structure not captured by this sort of isomorphism. A static model submarine with a frozen propeller or rudder lacks something other than size that its original has, and not merely a certain trajectory of change over time. Even an appropriately moving and developing model town or submarine, which tracks the history of its original, lacks something important and relevant to shared modal structure. It would be better if the rudder *could* move in other characteristic ways even at each moment of its history, so that there were alternative possible states of the model submarine at each moment that corresponded to alternative possible states of its original at the relevant moment, preserving patterns of similarity as well. Shared modal structure requires not just an isomorphism between the actual and the actual, but among the merely possible and the merely possible. And notice that any merely possible color variation in a red barn in the model town matches a corresponding possible color variation in its original. There are even appropriate relations of similarity among possible variations in the respective lengths of the relevant barn doors, though the model has in reality a much smaller door than the original.

There isn't much variation between the exact models and originals in these cases. As far as it can be managed, just proportional variation in size is involved. But modal structure can also be shared in much more latitudinarian ways. Imagine a creature with experience that is inverted on many axes relative to our human experience, although they look at the same world as we do. For instance, what seems on the left to us seems on the right to them, and not because they are facing us. Their experience is as ours in a mirror wherever they look. They intend to act into such a mirror-inverted world, we would say in error, but things come out alright on the outside, because of the systematicity of their mental inversions of left and right. They seem to themselves to live in a world that is mirror-inverted relative to what we see, but they never encounter conflicts of perception or belief that reveal that, and their actions suffer what we take to be a similar inversion relative to their expectations and intentions, so that everything works out fine. This mirror inversion is not the only inversion in their experience. Their color experience is also inverted relative to ours, with red exchanged with green, yellow with blue, and black with white, and with corresponding changes all over the similarity space of phenomenal colors, such that similarity relations in their experience of the colors of things preserve the similarity relations that are present in our experience.[50] There are also corresponding inversions in their other forms of sensory experience, in tastes and smells and the like.

Though many philosophers disagree about this, I think that a part of what you and I mean when we say "red" has much to do with our own experience of phenomenal red. If these inverted creatures call "red" the things out there in the world that we call "red," they will apparently mean something different by that. They will arguably mean what we mean by "green." But regardless of my idiosyncrasies on this topic, there is clearly a very high degree of shared structure between what we think and say and experience and what they think and say and experience.

Now imagine something more, that they are dead right, that the world is just how they experience it to be, and that we are wrong, that our experience is inverted relative to reality in all these ways. Aristotle thought that ordinary human experience of green captured something out there in reality. But imagine instead that these creatures get it right in that way, and we get it wrong. Consider our experience and thoughts under such a condition. I think in one strict sense our experiences and related thoughts are pretty clearly false. We are wrong about what things in the world are like. We think things are green when they are red. But it is also oversimple to say just that, even if you accept my unpopular view of the meaning of color terms. There is a very serious isomorphism between what we think on one hand and what the other creatures think and also reality on the other.

This is an important kind of shared modal structure. It is not merely that there is a pairing between any of our given claims about reality and some one of the other creatures' true claims, with for instance red of some shade in our experience paired with green of some shade in theirs. There is more structure of reality than that which even we capture accurately. The ball is a certain red and we think it is a certain green, but more, the ball might have been, or might become, a slightly more orangey red, and we think in a corresponding way it might have been, or might become, a slightly more blueish green. It might be further to the left, but we think it might be further to the right. It is not merely that to any true claim they make there corresponds a somewhat different thought we would have under the same circumstances that bears a recognizable isomorphism to theirs. Rather, beyond that, various possible claims we might make about

---

[50] We will eventually see ways in which this simple model of color inversions fails.

things preserve a kind of actual modal structure possessed by reality and by the fully veridical claims of the other creatures as well. Our own various possible claims are related to other possible claims we might make in what is in fact, despite our inversions relative to reality, a veridical way. The whole structure of alternative possibilities that the other creatures recognize corresponds as a body in structure to ours. A slight deviation in color in their experience maps to a similarly small deviation according to us, and to every change in color that something might undergo in our experience, there is a corresponding change it might undergo in theirs. Possibilities within each system have equivalent similarities amongst themselves.

This is another example of what I mean by a shared modal structure. But notice that here, unlike the case of the toy submarine, it is the *content* of our thought and its *apparent* modal structure that matches the true modal structure of reality. So now consider a further step. Maybe sometimes the apparent modal structure of someone's experience fails to meet the true modal structure of reality. Just as you might hallucinate a dragon when there is no dragon in reality, you might see something with such modal structure and it may not really be present in reality. Reality and the content of experience might not share modal structure. I think we have already seen reason to think this happens in the case of phenomenal color, although why this is so will require further explanation.

There might be even more radical differences that yet preserve modal structure. Shared modal structure may not require that lengths go to lengths, though perhaps of different extent or in different directions, or colors to different colors. We might switch lengths with colors, as long as similarity relations among different lengths in the first system correspond to similarity relations among corresponding colors in the other, and preserve the relevant sort of modal structure. But let's continue to focus on the case of color.

5

Such a notion as modal structure may seem too general to be useful. But what matters here are certain specific sorts of modal structure. Consider the way in which some yellow lemon is presented in our experience, and pretend for the moment that that experience is entirely veridical, that we live in an Aristotelian world. That particular lemon, its nature as a lemon, and its shape and color are presented in experience as having certain modal characters, I believe, which underwrite some characteristically Aristotelian claims about them. For instance, that color might be found on some other object, say, an umbrella, and the lemon might survive turning brown, while that piece of fruit arguably would not survive loss of the property of being a lemon, and its color arguably could not exist unless had by something. That is part of the modal structure of that object and that color, at least as they appear in our experience, I claim.[51] Now focus specifically on the essential nature of the exact yellow hue of one bit of the lemon's surface, as it appears in our experience. In other words, focus on that phenomenal color. Its nature involves, I believe, essential relations of similarity and difference to other exact phenomenal colors, other exact colors that may not color anything even in an Aristotelian world and so are in that sense merely possible. For instance, it is part of what that yellow is that it has a certain relation of similarity to unique yellow, and other similarity relations to various specific greens and oranges. The yellow of the lemon assures, if it really exists, that unique yellow is possible, and its being is entwined with that possibility.

Of course, it may well be that this complex modal nature of the phenomenal yellow in question is not really out there on the relevant bit of any real-world lemon. Indeed, perhaps, that is just one more reason to think that such aspects of our experience are not veridical, but are mere qualia. However, my basic idea about color qualia is to explain the apparent modal structure of such nonveridical phenomenal color qualia by appeal to the real modal structure of our neurophysiology of color vision that underlies, that indeed constitutes, that experience.

Focus on both sides of this equation, the qualia as experienced and its neural basis. And begin with the first, continuing to presume for expository purposes that Aristotle was right, that things out in the world really have the phenomenal colors that our experience seems naively to exhibit to us.

Some chair is a certain phenomenal red, a certain scarlet. But it isn't unique red, the red that is neither orangish nor purplish. Nothing in even Aristotelian reality, we will presume, is unique red, which is analogous in that way to Hume's missing shade of blue. Yet that scarlet red of the chair has essential relations of similarity to unique red. How can that be if nothing is unique red, if unique red lacks that familiar sort of existence? There two ways to make the metaphysics work, I believe, two ways to make sense of the unique red in question. There might be a Platonic transcendental universal of unique red, which exists even though nothing has it, to which our instanced red bears some essential relation. Or it might

---

[51] Mendola (2021) develops this conception.

be that there is somehow a subsisting but not existing unique red, which is entwined in being with the instanced, existing red, yet itself merely subsists without existing. Plato or Meinong:[52] take your pick. You may well object to both sorts of fancy and questionable metaphysics. But remember we are just talking now about the metaphysics of things that appear in our experience, and experience is not necessarily veridical. We are just talking about the putative metaphysics of what appears as it appears. I argue at length in *Experience and Possibility* that the second model of entwined subsistence and existence is more coherent and better fits the phenomenology of color,[53] but either model can work for our purposes. What matters most is that the experienced phenomenal color is, if it exists in reality as Aristotle presumed, entwined in being with colors that things might but don't have, with the merely possible colors of things.

Let me elaborate my own preferred model of how this would work in a bit more detail, partly because it is unfamiliar and yet I think phenomenologically more accurate, but mostly because it provides a better guide to analogous claims about our visual neurophysiology that will be important to us here. In a world of Aristotelian color, I claim: (i) the actual and the merely possible are deeply entwined and inseparable in being in certain ways, and (ii) the merely possible subsists without existing, has a different and more shadowy type of being than what is actual.

There are less and more complicated ways to elaborate such a structure, but focus on one relatively simple modal structure of this type that I will call "the superworld," which can serve our purposes here. Please remember that I am not going to ask you to believe in such a metaphysical monster. Rather, I claim that our experience involves the appearance of such a structure and that this mere appearance is an important clue to what actual neurophysiology constitutes us to have such experience, and how it does that.

The superworld has two key features. First, the merely possible is neither actual nor existent. It has rather a different type or mode of being, which I am calling "subsistence." There is a true metaphysical difference between the merely possible and the actual, but it is not a mere locational difference captured by some indexical term, nor captured by the difference between the abstract and the concrete, and it does not involve a special property of actuality had only by some existent things, something superadded on to existence, as some philosophers have proposed. On this conception, in such a superworld, what exists is actual and what is actual

---

[52] Although this is not "subsistence" in Meinong's sense. See note 54.

[53] Mendola (2021).

exists. But the merely possible involves a different sort or lesser degree of being than what is existent and actual.[54]

There is also a second key aspect of a superworld. The so-called possible worlds are various alternative ways that the universe might be as a whole and in detail. One way to put this second aspect is that possible worlds are not distinct individuals, as is traditionally claimed. Another is that facts about possibility are more local than traditional possible worlds models suggest, in a way that implies we cannot easily disentwine modal facts into distinct worlds. Though this a nonstandard view, it is natural and intuitive, at least when we consider certain cases. It really is not very intuitive to think that many of the modal features of my bike, for instance that it might be green though it is actually purple, involve a whole different universe, a different possible world out there somewhere which happens to have in it that bike being a different color. All that is intuitively relevant seems located right here with the actual bicycle. If the bike has a disposition, say to squeak when the peddles turn, or a power, say to refract light in certain ways even though now the light is off, then the bike itself and its actual properties, things here in actuality, seem to extend into what is merely possible, not only in something like the way that particulars extend on traditional conceptions from the present back into the past, or even in which scarlet might exist in more than one possible world, but in an even deeper way. The very being here in actuality of some properties and particulars, it seems, extends out into various alternative ways things might be. And this happens in such a manner that various alternative ways the world might possibly be, including the actual way it is, are not really wholly distinct from one another in the way that traditional conceptions of possible worlds presume. At least some of the very properties and particulars that make up some possible worlds, including our fancifully actual Aristotelian world, cannot be disentwined from other possibilities that obtain only in other worlds. The actual cannot be as easily separated from the merely possible as the standard conception of distinct possible worlds suggests. The actual and the merely possible are entwined, and in a very local way, particular by particular and property by property. Here's a case that even a physicalist may accept: This electron here in reality has a charge of a nature that will necessitate various effects under various merely possible background circumstances. The charge is itself at least partly

---

[54] John Punch held such a view. See Coombs (1993). This is not Meinongian, despite Russell's claims, because Meinong held that while abstracta subsist, mere possibilities do not. See Meinong (1960) and Russell (1905). See also Parsons (1980).

modal in nature, not just because it involves necessitation but because what it is cannot be isolated within a single possible world, the actual one, that is wholly distinct from other possible circumstances; yet it is a basic element of the actual world. The charge must be a part of the actual world itself, but it is by its very nature something that involves other mere possibilities also. In an existing charge, the actual and the merely possible are essentially entwined, and in a highly localized way. This second characteristic feature of the type of modal structure found in the superworld is that the actual and the merely possible are entwined in a local way. What we might call the "truth grounds" of certain modal claims are local, yet not merely actual.

As I noted, there are more and less complicated ways to develop this general type of modal structure, with locally entwined actuality and subsistence. Some for instance involve different varieties of subsistence. But what will be most relevant for our purposes will be the relatively simple form of this type of structure found in the superworld. In such a superworld, there is only one type of mere subsistence, which underwrites what we might call "real possibility."

What is actual and what merely subsists are often deeply entwined and interdependent in this conception, and the entwining is local, linking specific parts of actuality with specific mere possibilities with which their being is enmeshed. So notice that on this conception, the actual retains a kind of priority in fixing the kinds of possibilities that are relevantly real. Those are merely subsisting things, but they must be entwined in being with parts of actuality. There are no merely subsisting things without that actual root. Because possibility and actuality are locally entwined, distinct possible worlds do not root this sort of modality. But still there might be a vast system of intertwined real possibilities and actualities, locally rooted in the actual, and including all of what is said to exist "in the actual world," but also including merely possible elements with the other mode of being, with which those actualities are locally entwined. This whole structure is the superworld. It is focused on its actual root, and includes all of what is actual, the entire actual world if I can put it that way, but it also includes not fully existing but merely subsisting parts, so to speak, surrounding or sticking out of relevant actual parts with which their being is entwined, which are real possibilities in the relevant sense.

For instance, assume that here in actuality is something with a basic property that is a phenomenal color. That actual basic property may have by its essence relations of similarity to merely possible properties that merely subsist, yet that are intertwined in being with the actual property by those similarity relations. Those various subsisting properties are parts of the superworld, locally entwined with the actual property in the ways noted. If this conception were realistic, then there would be only one superworld, because it is rooted in the actual world. The superworld does not contain all the merely possible worlds per se, even if we restrict ourselves to possible worlds containing only real possibilities of the sort with which we will be concerned. The mere possibilities the superworld contains are very local; they are alternatives bound to this particular or that property, alternatives that are entwined with the actual being of those things. How they might be recombined into whole alternative worlds is another question.

Am I really going to ask you to believe in a superworld with such extravagant metaphysics? As I said, no. Absolutely not.

I think it merely provides a good model of what the content of our experience would require if that experience were fully veridical. In fact, I think, such experience isn't fully veridical. I think the superworld is one kind of illusion of our experience. But such a superworld provides an important clue to what in fact constitutes our phenomenal consciousness. That is why this quite baroque metaphysics is worth our attention in a discussion of phenomenal consciousness.

Focus once more on the case of phenomenal color.

If we presume that very specific phenomenal colors, very determinate hues such as a specific scarlet, are among the basic properties of things present in reality, that our experience is veridical in that way, and if we also presume, as some claim, that they are actual entities fully independent in being of other actual properties, then various puzzles arise: If something has a specific scarlet shade, yet that necessitates that it be red. Red is a so-called determinable color, which can be filled out in different ways to yield a specific hue, yet it also seems a quite fundamental fact about the Aristotelian world that some things are blue and some are red, that they have certain determinable colors as well as very specific and determinate hues. That relation between properties is a part of the essence of a specific scarlet. It is by its nature an instance of red. But then those two properties cannot be fully independent entities, as fundamental properties are often conceived to be. What's more, as we noted before, that specific hue, and for that matter also that determinable color, has of necessity various relations of similarity and difference to other colors, colors that, if not instanced in actuality, are still possible. The scarlet of some ball is very similar to some other red hue that nothing in reality has, even nothing in the reality suggested by our experience, yet that relation seems a part of its nature. Further, a special similarity to unique red seems an essential feature of that color. The scarlet seems to be in some sense an independent entity, yet dependent on and entwined in being with other colors, some of which may not even be instanced in the world, and so may not exist.

I think the resolution of these puzzles, if our experience were veridical, would require that we embrace the general type of modal structure characteristic of the superworld, and give up one sort of independence of metaphysically basic properties, by recognizing a characteristic kind of modal structure. If we presume that specific colors, very determinate hues such as a specific shade of red, are among the basic properties that are instanced in reality, then there will be modal structure of this sort: If something is a specific red, that necessitates that it also be red, and if something is red, that necessitates that it be some specific red among the full range of possibilities. That relation between those instanced properties is a part of their essence, part of their characteristic modal structure, and that specific hue, and also that general color, have of necessity various relations of similarity and difference to other colors that, if not instanced in actuality, are still possible, and on the view we are developing subsist. One detailed way this might work is if the instantiation of each specific color would necessitate the instantiation of a more general color of which it is a specific form, which in turn would require at least the subsistence of all the different more specific forms which that general color might take, and hence the obtaining of the various relations of similarity and difference among them. And perhaps above the so-called determinable colors, for instance red and green, there is also the yet more general determinable of being colored. On this conception, the actuality of a specific red requires the actuality of the property red in general and even color in general, and that in turn assures the mere possibility of all the colors, colors bound by essential similarity relations to each other. At least the mere possibility of all these variously related colors, their subsistence, is a part of what an actual specific color is, but they are also all rooted in it. The actual and merely possible are in this way locally entwined in the hue of some object.

I think that modal structure like that is part of the superworld suggested by our experience whenever we see some scarlet thing. I also think that similar things might be said of other sensory properties and even forms of particularity. I think that the general structure of the superworld applies to all our sensory qualia. I also think, as you may also, that this is an implausibly rich metaphysical conception, too rich for reality. Nevertheless, it provides an important clue as to how we actually have such experience, even though that experience is not fully veridical.

How does such a richly modal understanding of the content of our experience help explain how we have such experience? Because our neurophysiology has an actual modal structure that corresponds quite well to this fanciful superworld structure that we experience.

It is very unlikely that phenomenal color, color just as we experience it, reflects only veridical and objective aspects of the world. For instance, we have noted that color as we experience it involves the unique hues and a circle of similarities that is captured by a child's color wheel, and specific determinable colors like red, and these structures do not plausibly correspond to any objective features of reality. But these structures in our experience are plausibly due to our human neurophysiology of color vision, and in particular to its modal structure.

I will argue in detail in Chapter 4 that your neurophysiology of vision is such that your actual state of experiencing a determinate red implies the real possibility of the experience of other colors and necessitates the actual experience of determinable red. For instance, when you are in state of experiencing something as red, that involves a certain specific pattern of neural activation in the context of a background structure that, if activated in some alternative way, would constitute the experience of green. No metaphysical hocus-pocus is required. There are two layers of structure in the aspect of your nervous system responsible for your color vision. One stable background layer is shared by more specific states of actually seeing green and actually seeing red. And it assures even when you are actually seeing red that there is the real possibility of your seeing green. This in turn explains why the actuality of the red you experience would be locally entwined in being with the mere subsistence of green in the manner of the superworld, if only that red were actual. More generally, your neurophysiology of vision is such that your actual state of experiencing a determinate red implies the real possibility of your experience of other colors with which it bears essential relations of similarity, including a special relation to the real possibility of your experience of unique red, and necessitates your actual experience of determinable red, and this aspect of the modal structure of your neurophysiology of vision explains the modal structure characteristic of the phenomenal colors you experience, in which the instantiation of some determinate red would necessitate that the object in question also be determinable red and in which that determinate red would have essential relations of similarity even to colors not instanced in reality.

That is the general idea behind the theory of sensory qualia propounded here, the third key aspect of my proposal about phenomenal consciousness. We will consult the apparent essences of apparent sensory properties to discover the forms of entwined subsistence and actuality they would require to be real, and then look to see how experience of things with that implausible modal structure is rooted in the plausible modal structure of our human neurophysiology, a neurophysiology which, I argue, constitutes that experience. This has also been a brief introduction to the general idea of modal structure, which will also play important roles in the other two key aspects of my proposal.

6

Now we are in position so that I can properly summarize the whole thesis of this book and sketch its plan. As noted, there are three main positive features of the general account of sensory phenomenal consciousness to be developed here.

The first involves a modal structural understanding of our physical, animal nature, which is the crucial general basis of our phenomenal consciousness. Animals are the most intuitive locus of phenomenal consciousness, and this core idea involves a certain understanding of their nature. It itself has three components. The first of these components reflects the fact that our phenomenal consciousness is crucially that of agents, who have an abstract modal structure which is the linking of a range of possible stimulations and a range of possible actions in various possible ways. Second, our phenomenal consciousness is crucially that of a living thing of the familiar sort found on earth. We are constituted out of living cells of familiar types, and which crucially involve both metabolism and replication. Metabolism and replication both involve modal structure in certain ways we will soon consider. Third, our phenomenal consciousness involves the very complex dynamics of our neurophysiology, which is what mediates between stimulation and action in us, and which has, I will argue, a relevant sort of modal structural complexity. One example that has already been noted is that it sometimes matters that actual neural firing occurs in a context in which other sorts of firing are really possible given that relatively fixed background state of our neurophysiology. This first main feature of my positive view is the topic of Chapter 2. This chapter also includes an introduction to the necessary rudiments of neurophysiology, which underwrite this model.

The second main positive feature of my proposal is a conception of the neurophysiological basis of our experience of qualia in particular, of phenomenal consciousness. Chapter 3 will develop this idea. It does not involve information or representation, but rather identifies what constitutes phenomenal consciousness with the much more complex modal structure of the neurophysiology that intervenes between sensation and action. I call this "the whole nervous system model." This requires an allied account of phenomenal consciousness as involving what I call "hyperexperience." This is richer in certain ways than traditional conceptions of that consciousness. Some examples of mere hyperexperience, which has a kind of lesser presence in our phenomenal consciousness, include our awareness of the immediate past of seen objects and of some of their spatially occluded aspects. Another important related aspect of my proposal is the claim that there are different forms and degrees of phenomenal consciousness, even beyond the complexities found in hyperexperience. But, putting such complexities for the moment aside, the basic idea of this second main feature of my view, put very bluntly, is that what constitutes our phenomenal consciousness is the entirety of our nervous system that intervenes between sensory stimulation and motor action. The apparent modal structure of our phenomenal consciousness in general is mirrored, predicted, and explained, by the actual modal structure of more or less our entire neurophysiology.

What is phenomenal consciousness, at least in a sentient human animal? It is a modal structure that links sensory stimulation to action. What is the neurophysiology of that sentient human? A complex thing that has a modal structure, perhaps at a certain level of abstraction that ignores fine details, and links sensory stimulation with action. And those two modal structures are the same thing. This is a natural but unconventional idea, and this conception faces a number of obvious objections. So much of Chapter 3 is devoted to answering pressing worries. Some are rooted in commonsense phenomena like intuitively unconscious aspects of beliefs and desires, while others are rooted in such things as the unconscious representations deployed in cognitive science and surprising psychological phenomena like blindsight.

The third main feature of my view is an account of particular sorts of sensory qualia, which depends on a match between the specific apparent modal structures of various qualia and the specific actual modal structures of various sorts of relevant sensory neurophysiology. I argue that our neurophysiology suffices in detailed ways to constitute our phenomenal consciousness in the case of sensory qualia. This component is initially developed in Chapter 4 through further attention to the relatively well-studied case of color experience. It is a way in which details of our sensory qualia reflect details of our neurophysiology, through a kind of shared modal structure. As I've indicated, the full range of possible phenomenal colors seems to be rooted in any actual instance of any phenomenal color, and such colors involve a complicated structure of essential similarities and differences. That is part of the modal structure of the color properties we experience. And I will argue in Chapter 4 that this reflects in various detailed ways the actual modal structure of our neurophysiology of color vision.

While the notion of modal structure is very general and abstract, various specific sorts of modal structure are relevant in these ways to the understanding of all three main facets of the view I propose. My proposal involves pressing on both sides of the key equation at once, pressing up from below, in its first facet, on our understanding of our neurophysiology, and in its second facet, pressing down from above on our understanding of our phenomenology, in both cases in a way attentive to modal structures of various sorts. And the details of the realization of our experience of particular sensory qualia will involve matching modal details below and above. So the general theory of the phenomenal consciousness associated with human sensory experience proposed here might be called the MOdalized Up and Down and in Details account, the MOUDD theory, if you like that sort of nomenclature.

The remaining chapters will apply that general account to a series of different cases, and in that way develop it. Chapter 5 considers some other senses. Chapter 6 distinguishes various sorts of spatial structure and particularity that are sometimes present together in our visual experience, in a kind of palimpsest. For instance, glimpsing the red of a barn outside through a window with a yellow frame, you can see it at once on the distant barn and yet as characterizing a visual field in which it is very close to the occluding yellow edge of the window. This chapter develops a modal structural understanding of the neurophysiology that roots core features of these sorts of experience. Chapter 7 applies a modal structural understanding of our neurophysiology, and the allied notion of hyperexperience, to our conscious experience of time, and to what I argue is a kind of causal experience. It does this on the way to an account of our conscious experience of the robust particularity of ordinary concrete objects present in our visual experience. One novelty of this discussion is that I argue that the proper modal structuralist understanding of our neurophysiology can help explain the complex structure of temporal consciousness explored by Husserl.

Once the MOUDD theory is on the table, Chapter 8 reviews it and considers a standing dualist objection to any physicalist account.

CHAPTER 2

Neuralia on Earth

This chapter articulates the first main feature of my proposal regarding the constitution of the phenomenal consciousness of our sensory experience, a certain modal understanding of our life as animals with unified nervous systems.

We aren't the only ones. Although this book is focused on human consciousness, still such diverse animals as squirrels, bats, octopi, and fish also intuitively involve such sensory phenomenality. At least short of dualist presumptions, it is implausible to claim that "consciousness" is a natural kind term exactly like "water," a term for something whose hidden essence we could discover by empirical means in the same way that we discovered that water is H-O-H. Still, there are paradigm actual cases, and we and these other animals are among them. So the first step of my account is a certain modalized understanding of the empirical facts regarding what these paradigm examples involve. We need not decide here whether hypothetical living things of various sorts that lack some of these features – or sophisticated robots, or bacteria, plants, or jellyfish – have phenomenal consciousness of any sort, though I am inclined to say that they may in some cases have various sorts of resemblance to the paradigm cases that make such a claim not unreasonable. Further, among the paradigm cases of phenomenal consciousness, we are most central.

In this first step of my proposal, we will focus on our nature as neuralia, which is to say animals with nervous systems that have specialized sensory receptor cells and muscle cells. Sponges need not apply. In Chapter 3, we will further restrict our focus in such a way as to make it also useful to exclude even now animals with merely loosely organized nervous systems, like jellyfish. But it is a certain modal structural understanding of the basic facts about neuralia that we will be specifically concerned to develop here. The notion of modal structure is a very general notion, but we will be isolating several specific sorts of modal structure in this chapter that are a part of my proposal about our sensory phenomenal consciousness.

https://doi.org/10.1017/9781009514774.002

1

There are three points involved in this initial feature of my proposal, the focus of this chapter.

First is a certain abstract and commonsense but modal conception of the animals that are the most intuitive locus of phenomenal consciousness. Like squirrels, octopi, and fish, there is a range of possible sensory stimulations to which we are differentially sensitive. That range of possible stimulations is one sort of very simple modal structure. And there is also a range of different possible things we may do, which has an analogous structure. But this involves modality in another way as well. What we do has effects; it is causal. The relation of causation may involve necessitation by a cause of an effect, which is modal in one sense, and even when a cause merely affects the relative probability of one effect or another, that is modal in a looser sense. So our range of possible doings involves two sorts of modal structure, and in us and those other animals, these two kinds of things – stimulations and what is done – are linked, in the sense that what we do is responsive to sensory stimulation, in a quasi-causal way. But there is also a variety of possible ways that, given a specific sensory stimulation, we may respond by doing something. There is a kind of linkage in such an abstractly conceived agent between different possible incoming stimuli and different possible outgoing causal and hence modal effects, although it can vary over time what specific linkage that may be, as learning or different internal conditions like hunger and thirst or even something fancy like choice occur.

It is among such abstractly characterized agents that we most expect to find phenomenal consciousness. But not perhaps in all of them. Robots or thermometers or computers arguably sometimes meet such an abstract modal characterization. But they are not alive. And it is among living things that we are most confident phenomenal consciousness is found. And so we come to the second crucial point.

2

Among such agents, we are most confident that phenomenal consciousness is found among the living ones. Perhaps, as some believe, certain artifacts, or even particles, are also phenomenally conscious. But we don't generally presume so. So while I will not assume that phenomenal consciousness in all its forms requires animal life or even life, I will assume that our particular form of phenomenal consciousness requires it. I will focus here on what seems the central and uncontroversial case, the phenomenal consciousness of living animals, among which ours is found.[1]

There are many complex and difficult questions about what wildly unfamiliar or merely possible situations would count as life. There are also reasonable questions whether even such a regrettably actual thing as a virus is alive. But I will evade these issues by a yet more specific focus, on the actual life of which we are certain, which I'll call "life on earth."

I have no idea how to define "life," and in fact suspect there are no hard-edged facts about that, that it has no essence of that definite sort. But certain key features of the life we know are captured in current definitional attempts that descend from Schrödinger's seminal *What is Life?*.[2] These definitions generally specify that life involves both replication and metabolism. Living things reproduce and grow. But they also take in energy from their environments, either by eating other living things and digesting them or by processes like photosynthesis, which allows them to maintain their own complex organization far from thermal equilibrium.

These definitional ideas can be developed in various ways, and with various refinements. For instance, it isn't silly to suggest that living things are reproducing systems able to carry out at least one thermodynamic work cycle, to fend off the degradation the second law of thermodynamics inflicts on a closed system.[3] It isn't silly to suggest that living systems are autopoietic, organized so that their constituent processes produce the components necessary to continue those processes.[4] It isn't silly to claim that living organisms are self-constructing, self-maintaining, energy-transducing autocatalytic entities that are capable of evolving by variation and natural selection so their forms and functions are adapted to their environment.[5] Take your pick. Such refinements won't matter to us, though we will soon return to the significance of both metabolism and replication.

Beyond anything obviously crucial in a viable definition of life,[6] like metabolism and replication, anything plausibly required of any life found in other galaxies or other possible worlds, there are two other important features of actual life here on earth that will matter to us: Living things on earth are constituted by living cells, or are cells themselves. And these living cells are built from a very limited range of chemical constituents.

---

[1] The significance of life to consciousness is a theme of E. Thompson (2007) and Godfrey-Smith (2016b, 2017).

[2] Schrödinger (1967).

[3] Kauffman (2000: 4).

[4] E. Thompson (2007: 98), following Maturana and Varela.

[5] Harold (2001: 232).

[6] Although autopoiesis is supposed to involve cells in some way.

The nature of the cells by which our phenomenal consciousness is constituted is, I claim, important to its nature.[7] So I hope you will forgive some details.

The cells of actual living things are constituted almost totally by a very few chemical elements, largely of carbon, hydrogen, oxygen, nitrogen, sulfur, phosphorus, and hydrogen. Potassium and sodium also play starring roles in what will follow. It will be useful to bear the nature of these elements in mind. You will recall that an atom has a nucleus that consists of various numbers of positively charged protons and uncharged neutrons, bound together by nuclear forces, and also involves "orbiting" electrons that are negatively charged. The electrons are bound to the nucleus by electromagnetic force. An atom is in this way somewhat analogous to a system of planets bound to a central star by gravity. But this analogy is only partially illuminating. One important difference is that the electrons of an atom, unlike the planets orbiting a sun, can have, because of the way quantum mechanics works, only a few discrete energy levels. Another is that quantum mechanics forbids any two electrons in an atom from having the same state, where the state is specified by four "quantum numbers," including the energy level. These quantum mechanical facts are what assure that there are a very specific set of elements with very specific chemical properties, including certain possibilities of recombination. If atoms are ordered by regard to their atomic number, by the number of protons they contain, then they fall into the surprising order represented by the periodic table of elements, with a striking chemical similarity of different elements repeating with a certain rough periodicity as the atomic number of elements increases. What is most crucial to the chemical properties of an atom isn't really its exact number of electrons, which must however balance its number of protons, and so reflects its atomic number; rather what is crucial is how many of those electrons inhabit the outer or valence shell, which is the highest of the discrete energy levels allowed by quantum mechanics in that atom, as compared with how many possible spaces for electrons in that outer shell are allowed by quantum mechanics. An element is most stable and least chemically reactive when it has an outer shell that is fully filled, with no possible spaces unfilled. This itself is one kind of modal structure, but the main immediate point to notice is that if the valence shell is not filled, the atom will tend to bond to other atoms in certain possible combinations, to form molecules. That is itself a second type of modal structure. There are various forms of bonding.

---

[7] For this trend in philosophy of science, see Chirimuuta (2014) and Glennan (2017).

For instance, in an “ionic” bond, one atom lends one or more of its valence electrons to another, so as to fill its valence shell, and the two “ions” which result from the exchange, which because of the exchange have become overall electrically charged in opposite ways, are pulled together by that difference in charge. This happens, for instance, when sodium and chlorine form a salt molecule. In a “covalent” bond, atoms share an electron in their valence shells. This happens for instance when hydrogen and oxygen form water.

No doubt there are forms of merely possible life with a basis that is not chemical, and maybe life could be formed on another chemical basis than the familiar one. But the chemical properties of carbon, hydrogen, nitrogen, sulfur, phosphorus, hydrogen, oxygen, potassium, and sodium, among others, are very important to life on earth. And they are due to the subatomic structure we’ve noted. A carbon atom, for instance, has four electrons in its valence shell, as opposed to the eight allowed by quantum mechanics. It is hence capable of bonding to four other atoms at once. It has a considerable tendency to form covalent bonds, for instance with other carbon atoms, even covalent bonds where several electrons are shared by two atoms. So it can bond into sheets or rings or long strings with other carbon atoms and atoms of other sorts. Organic molecules, which constitute living things, are often chains or rings largely of carbon atoms, with a few other atoms, often nitrogen, oxygen, phosphorus, or hydrogen, added in. In life on earth, organic macromolecules of four types are especially crucial. These are proteins, nucleic acids, lipids, and polysaccharides. They have characteristic chemical features that are crucial to life as we know it, due to the features of the atoms that constitute them and the way those atoms are bonded together.

The cells that constitute all life on earth are built largely from this limited set of molecular constituents. We have briefly noted two forms of salient modal structure at this chemical level, and this indeed is a third. Our understanding of the subatomic constituents of molecules like protons and electrons is largely modal, which is yet a fourth; physically basic properties like charge and mass are largely understood by regard to their causal role in fundamental physical laws. This four-fold modal type of its chemical and physical basis is a modal structural aspect of life on earth, which is I believe important in detail to the nature of our qualia. But I will concentrate here on other ways in which forms of modal structure are important to life on earth, relevant to the other and more definitional aspects of that life: metabolism and replication.

To properly understand life on earth, especially those forms with nervous systems, we need to consider in some detail the nature of the cells built from these molecules, which crucially incorporate both replication and metabolism. The nervous system, which I will argue constitutes our phenomenal consciousness, is built, like all familiar animals, from such living cells.

But since human cells are complex, and humans are made of trillions of cells, it is most illuminating to focus initially on a relatively simpler model case with which our own cells can be usefully compared. It is a case of special significance to human life, since our guts contain more of these cells than there are human cells in our bodies, and we could not live without them. In that sense, each of us is a symbiosis between a vast colony of such cells and human cells.[8] Let's dive inside ourselves and these usefully simple cells and have a look.

3

We begin with *E. coli*.[9] It is a single cell creature, and it is a prokaryote, which means its cell has a relatively simple internal structure. But there are certain chemical structures common to the cells of virtually all living things on earth, to creatures as different as humans and *E. coli* bacteria, and which are uniformly important. So this is a good place for us to begin our minuscule explorations. That will also allow us to consider some relatively simple forms of modal structure relevant to animal life in a less hostile environment than the complexities of human cells provide.

One simple form of modal structure plays a significant role in current biology and neurophysiology, as well, we have seen, in many contemporary accounts of consciousness. This simple structure is implicated by the notions of information, representation, and coding, at least when they are conceived so as to involve a modally robust correlation between, on one hand, a range of possible represented things and, on the other hand, a range of possible representers, where each particular represented thing is reliably linked with a particular representer. A representer, a sign, may hence be said to carry the information that the represented thing exists or obtains. Think of how the number of rings in a tree bears information about its age. Talk of information of at least roughly this sort is ubiquitous in current biology, neurophysiology, and cognitive science.

---

[8] Even human cells may be surprisingly uncooperative. See Ehrenreich (2018).

[9] One introduction consists of Goodsell (2010) and Harold (2001: 63–97). For current details, see Albert et al. (2015).

Further, there are three ways that information is plausibly present in *E. coli*, two relevant to arguably essential features of life – replication and metabolism – and one relevant to a close analogue of sensation. I will now introduce these general features of *E. coli* and develop the ways in which information and that sort of modal structure comes in, before we turn to the somewhat contrasting nature of human cells. But one caution: We love our machines, whether cars or microscopes. And it is almost irresistible and aids comprehension to say that certain molecules and molecule complexes function throughout almost all living cells as tiny machines or mechanisms that perform certain tasks. Their various chemical properties and structures differ in a way that allows and explains this, in a roughly uniform way across virtually all life. But we should beware. To speak of machines is teleological, since machines like other tools have functions, and there is a reasonable question about the legitimacy of that talk in some natural contexts. What's more, the common and revealing descriptions of the working of these little mechanisms that I will give are very idealized. Many relevant structures do not really pass through a definite series of specific states like a familiar machine, but are rather stochastic. They proceed through a basically random series of stages driven largely by external forces, but with a relatively slight statistical bias in a certain productive direction. It is unlikely that they will perform their characteristic job in exactly the same way twice in a row.[10] These are important facts about the dynamics of life that we should remember, though I will still take the easy and appealing expository road, and talk about molecular machines regardless.

An *E. coli* bacterium is a very simple cell, but it is still fantastically complex in relevant detail. It includes about 2,400,000 protein molecules of thousands of different kinds. It consists overall of 300 million molecules, not counting the 40 billion water molecules that constitute nine tenths of its mass.[11] Inside its three-layer cell envelope are some molecular machines found throughout life. Life crucially involves both replication and metabolism. Consider first some mechanisms relevant to replication and growth.

The bacterium includes a single, huge, circular molecule of DNA. DNA, which is found throughout life on earth, is a nucleic acid. It consists of two chains wound around in a double helix. Each chain consists of a series of four types of nucleotides. Each type of nucleotide involves a similar sugar and phosphate structure, which join to form the

---

[10] P. B. Moore (2012); Skillings (2015).

[11] Harold (2001: 66–68).

backbone of the chain, but also one of four distinctive molecular structures: adenine, guanine, cytosine, and thymine. Nucleotides pair together to form the rungs of the twisty double-helix ladder. The two chains of the double helix are complementary in the sense that where one has thymine, the other has adenine, and where one has guanine, the other has cytosine.

RNA is another nucleic acid that plays important roles throughout life. But it is a single helix rather than a double one, in which DNA's thymine is replaced by uracil, which is smaller by a carbon atom and a few hydrogen atoms. DNA is for these reasons more stable.

In a sense that we will shortly examine, DNA is said to be the genetic code that differentiates the cells of one living animal from another. The sequence of paired nucleotides in a given DNA molecule seems to carry information, like the sequence of words in a text. In *E. coli*, and virtually all life on earth, DNA is apparently at the root of and the blueprint for a chain of protein production. It works roughly like this: DNA splits like a zipper. Then what is called "messenger" RNA, matching the DNA sequence, is formed on its now exposed template. This messenger RNA travels to molecular structures called "ribosomes," where proteins are synthesized.

There is, for the most part shared across life, a standard three-nucleotide code, according to which a certain sequence of three rungs of a DNA ladder specifies the construction of one particular amino acid out of twenty different possibilities. When these are synthesized by a ribosome, in some linear combination specified by the sequence of messenger RNA and hence ultimately DNA, they will go on to fold in a complicated way into a given protein. Fans of modal structure will see another instance of it here.

Different proteins function like different little machines in a cell to perform various duties, or combine to do so in another modal structural phenomenon. Their various chemical properties and structure differ in a way that allows and explains this. And these mechanisms of various sorts are also shared across different forms of life.

Consider for instance the ribosomes of *E. coli*. They contain about sixty different proteins, as well as RNA molecules, organized into two submachines that work together somewhat like gears. What is called "transfer RNA" leads amino acids to the ribosome. As messenger RNA is fed into the ribosome, the two protein submachines move back and forth and attach the amino acids together in a matching sequence, to form a protein.[12] Ribosomes are little protein factories, and virtually identical in all living cells.

---

[12] Falkowski (2015).

But to form the bonds between amino acids requires energy. Besides the replication involved in life, there is crucially also metabolism. Talk of DNA "coding" naturally invokes information and so a simple form of modal structure in ways we will need to examine. But metabolism is also said to involve information and so such modal structure in a way we will need to probe. So first consider some of its details.

Adenosine triphosphate, or ATP, is a nucleic acid molecule that contains a sugar and three phosphate groups linked one after the other. When ATP is cleaved into adenosine diphosphate, or ADP, on one hand, and a separate phosphate on the other, that produces chemical energy. This mechanism is used throughout living cells to provide energy for the synthesis of proteins, for motility, and to pump ions across membranes.

ATP is created by cells in various ways, which all depend on energy being taken from somewhere else. Some bacteria, which live in very extreme climates, take energy from inorganic processes, such as the oxidation of hydrogen sulfide or reduction of carbon dioxide to methane. But we will focus on two more familiar cases. Some algae, plants, and bacteria take energy directly from the sun's light through photosynthesis. Some living things, including animals and *E. coli*, take energy from their environments through respiration and the breaking down of organic materials by oxidation, organic materials that are themselves dependent ultimately on photosynthesis.[13] These two processes – oxidation rooted in respiration and photosynthesis rooted in light – depend on the same sort of mechanism, the operation of catalytic proteins in the membranes of chloroplasts or mitochondria in complex cells and in the cellular membranes of relevant bacteria.[14] ATP is formed when protons move across these membranes. This is accomplished by little machines called coupling factors, which are like tiny motors that penetrate the membranes.

> They contain a shaft that is a set of proteins spanning the membrane and physically inserted into a set of larger proteins (a *head group*) that sits at one end of the shaft .... Protons on one side of the membrane bind to and move along the shaft to cross the membrane. In the process, their flow physically turns the shaft counterclockwise .... As the shaft physically turns, it mechanically moves the larger proteins ..., which bind ADP and phosphate.... The motor can also operate in reverse.[15]

This molecular machine is common throughout life on earth. But it requires that there be an electrical gradient, with a surplus of protons on

---

[13] Harold (2001: 40–44). [14] Ibid., 42–44. [15] Falkowski (2015: 56–59).

one side of the membrane and a shortage on the other. In photosynthesis, this is due to the absorption of light by molecules called pigments, which leads to a complex chemical reaction that creates the gradient. We will eventually see that processes like this are also crucial to human vision. In respiring bacteria like *E. coli* and the mitochondria of complex cells, electrons pass to oxygen from reduced organic material through a series of electron carriers, the respiratory chain, more or less as in a bucket brigade. And as they do so, protons are transferred across the membrane, creating the necessary electrical potential, and also a difference in pH. The membrane is not generally permeable to protons, but the coupling factors are embedded in the membrane, and allow the protons to flow back down the electrochemical gradient, creating ATP in the process.[16]

There are other more idiosyncratic molecular nanomachines in *E. coli* cells, which, as we will see, invite talk of information in yet a third way. These cells have flagella turned by tiny molecular motors, which are controlled by sensory proteins that monitor the level of nutrients.[17] When an *E. coli* bacterium is swimming in some direction in which the nutrients are not increasing, the protein signals the motor to reverse direction, which causes the cell to tumble and change direction. This is not unlike sensation in animals.

Real mechanical and electrical machines have a determinate structure of parts reliably performing determinate causal roles in the overall system of the machine in question, parts whose various states are counterfactually dependent on the states of other parts in a determinate way and which make determinate contributions to the determinate job of the whole. This is a kind of relatively determinate modal structure of the whole, which can be well captured by functionalist analysis. But remember that the machines we have been discussing are really just statistical biases in complex underlying physical structures. In reality, there is far less determinacy of the kind of machine-like modal structure present in these cases than my simplifying talk may suggest. Still, one relatively simple form of modal structure does play a starring role in all three of the aspects of *E. coli* that we have considered.

Begin with replication. There is a clear sense in which information is present in *E. coli* DNA. This is above all information regarding the proteins that constitute such a cell, and hence ultimately about the little machines present in its structure, and indeed to a great degree about its overall form. About 88 percent of the DNA in *E. coli* codes for proteins.[18]

---

[16] Harold (2001: 82–87).

[17] Goodsell (2010: 66).

[18] Harold (2001: 71).

There is also a sense in which that information is carried by messenger RNA from the DNA to the ribosomes for the construction of the relevant proteins, which will then go on and fold in characteristic ways, and form relevantly functional parts of the cell.

We've already noted a link between one simple modal structure and information. Any sort of covariation or correlation of possible states of *s* with possible states of *r* allows *s* to carry information about *r* in at least a weak sense. The tree rings in a tree carry information in this sense about its age, but so too the age about its rings. A very primitive kind of modal structure and dynamics is present in such cases, correlating one set of alternative possibilities with another. Covariation can occur with various degrees of probability, but in the cases in front of us, even high probabilities are involved. But covariation begs out for explanation. And due to the relevant explanations, there are hence slightly more specific senses in which "information" can be said to be present in these cases. The relevant covariation involved in these stronger sorts of information involves a kind of asymmetry, perhaps causal or temporal, but there are two ways in which the asymmetry can be relevant. The DNA sequence causes the protein it codes for, not the protein the DNA, and so it is plausible in one sense to say that the DNA carries information about and is a kind of blueprint for the protein, not the other way around. On the other hand, we also often say that information is carried about a source by a signal or message, as a tree's rings record its age, which involves the other direction of dependence. Either of these directionalities is a kind of modal phenomena, involving necessitation of at least some high probability in one direction or the other. Both sorts of this slightly richer type of information are plausibly present in *E. coli*. The messenger RNA carries information about the necessitating states of the DNA, and the DNA carries information about what proteins will be necessitated. Beyond this, it is sometimes suggested that there is a yet richer sense of "information" relevant to biological cases, which requires that things have the *function* to covary in one of these ways, where that function depends on evolutionary history. Such functions might more fully underwrite my teleological talk of molecular machines as well. But without delving into the complex controversies regarding the propriety of talk of evolutionary functions and what exactly they require,[19] I think we can see that the high probabilities of directed covariation of various sorts in *E. coli* suffice for the presence of recognizable information, even if the entire universe just popped randomly

---

[19] For an introduction, see Lewens (2007) and Sober (1984, 1994: 25–118).

into existence the day you were born, or indeed, to fit the case of fast-living *E. coli*, yesterday before lunch, and so nothing evolved or otherwise had a function determined by its history.

At least when deploying such a non-teleological notion of information, we need to be wary of the baggage of the word. "Information" sounds semantic, and it may be objected that there is no meaning present in these probabilistically causal processes. Nevertheless, it is true that the notion of information used in communication theory, developed by Claude Shannon, is applicable to such cases.[20] Something in this sense carries information about its source when its state is correlated with that of its source, and when it carries more information about its source, the better predictor it is of the state of its source. This allows a quantification of the amount of information a signal can carry. For instance, if there are 8 equiprobable possible states of the source, then a signal must carry 3 bits of information (since $8 = 2^3$) to fully specify the state of the source. Messenger RNA plausibly carries Shannon-information about DNA. On the other hand, that notion is indifferent to the actual alternative states of the source, to any intuitively semantic information carried by the signal. The signal carries the same amount of information in this sense whatever it is about in detail. So Shannon-information is not really a fully semantic notion. Signaling in any very intuitive sense requires more. There is nothing in Shannon's notion that requires that the information in the "signal" be capable of being understood by anyone or anything. But when we speak of a signal carrying information of a fully intuitive sort, we presume that there is not only a source and a means of transmission of the signal whose possible states covary, but also that there is a reader or receiver that is appropriately sensitive to those differences in the signal, which can read the signal and act in a way appropriate to the specific semantic information about features of the source being carried by the signal. Still, even that structure is to at least a degree present in *E. coli*. This is because the ribosomes construct the proteins on the basis of specific information carried by the messenger RNA from the DNA.[21]

There is a complexity we should note. The DNA of *E. coli* does not always produce the same proteins. For instance, when lactose is present, the bacterium ramps up the production of an enzyme that can break that down in a metabolically useful way. It does this because there is a repressor

---

[20] Shannon (1948).

[21] If true signaling requires also some arbitrariness in the code used by the signal, still that is arguably present in DNA and RNA.

protein ordinarily bound to a short bit of DNA that is near the DNA that encodes for that enzyme, which prevents that DNA from being transcribed into RNA, and which falls off when lactose is present. This is a kind of genetic switch, which turns the production of the enzyme on and off.[22] Still, this complexity probably should not undermine our intuitive sense that the DNA carries information about that enzyme, since a book may bear information even when it is unread, or temporarily locked in a bookcase.

The lover of possibility may see a tempting opportunity at this point. The way in which DNA acts as a kind of blueprint may suggest that all possible life of one sort is specified by different possible sequences of DNA. And I've already suggested that there is one sort of combinatorial modal structure here. But I think any temptation to identify this with possible forms of earthly life should be avoided, for several reasons. Even in the case of *E. coli*, there is a complex molecular environment, including for instance ribosomes, in which DNA functions so as to specify proteins, and in which proteins fold in characteristic ways. That is a stable enough environment for *E. coli* so that we may intuitively speak of information in these cases. Also, the proteins constituting the ribosomes and much of the rest of the chemical environment in which the DNA functions to specify proteins are in fact encoded in the DNA. Still, it certainly does not follow from these facts that *E. coli* DNA on its own would suffice to generate such a bacterium. *E. coli* reproduces asexually. It is a kind of rod that grows and then splits in the middle. The activity of *E. coli* DNA that we have been discussing takes place in that spatial and chemical environment. So it may well be that a part of what an *E. coli* bacterium inherits and that governs its form is the shape of its ancestor cells so that what is called "epigenetic inheritance," not carried by DNA, occurs even in such a simple case.[23] There are also other reasons to resist this opportunity. The dynamics of cell development and reproduction is somewhat complex even in the case of *E. coli*, though I believe it is still in that case simple and stable enough for talk of DNA information to be appropriate, even if it isn't *all* the information that is so relevant to the replication and growth of such cells. But the more general notion of DNA as providing blueprints for possible living things is probably not appropriate. For one thing, some living things deploy RNA in DNA's usual role. And there are some sequences of DNA that could not support properly the machinery required for their own replication through growth and reproduction, and so do not characterize possible living things.

---

[22] Carroll (2005: 54–61).

[23] Harold (2001: 99–115).

Also, we will see that in more complex cells, the relation between DNA and proteins is considerably more complex than in *E. coli*, in a way that magnifies the importance of differences in epigenetic inheritance. Still, on the other hand, it is striking that there is a more restrictive range of DNA that might be thought to characterize life on earth, in such a way as to delimit the significance of such epigenetic factors. For instance, much of the DNA that codes for the key components of the molecular machines we have been discussing is quite stable across at least respiring forms of life, and is a surprisingly small component of DNA in that life.[24]

Whatever the fate of the tempting opportunity, it is reasonably clear that there is information relevant to replication and growth in *E. coli*. But there is also a second way in which information is crucial in the life of *E. coli*, reflecting a feature not of replication and growth, but of metabolism. But the way modal structure is involved in this case is somewhat different.

Energy is conserved. That is the first law of thermodynamics. But usable energy is not generally conserved. For instance, heat flows from a hotter to a colder body, so over time we can expect the different elements of a closed system that start out with different temperatures to come to thermal equilibrium, which will make it eventually impossible for us to use the temperature difference in some productive way, say to drive a steam engine. The second law of thermodynamics captures this fact by saying that in a closed system, entropy will not decrease but may increase. Entropy could stay the same at the lucky limit, if the process in question is frictionless, driven merely by an infinitesimal imbalance of forces, and so reversible. But otherwise and realistically and generally, entropy increases in any closed system.[25] And entropy is a measure of disorder. To greater entropy, there corresponds greater disorder. As a closed system moves to a uniform thermodynamic equilibrium, it exhibits less and less order.

Entropy has a lot to do with possibility. Consider a bunch of particles. They have various positions and momenta.[26] If there are n particles, what is called a "phase space" for that system is an abstract space of $6n$ dimensions, where each point in the space represents some possible configuration of all the particles, specifying all their positions and momenta. Over time, the system will move from point to point in that space, governed by dynamical laws. Now divide the phase space into subregions, where each subregion constitutes possible configurations of

---

[24] Falkowski (2015: 95–96).

[25] Van Ness (1969).

[26] There are great complexities if we relax my simplifying assumption that we are dealing with classical particles.

the particles that are macroscopically indistinguishable. This is called "coarse-graining." Some subregions, characterized by a given macroscopic description, will be much larger than others, because for instance there are many more ways for a bunch of atoms to be at thermal equilibrium than to be in some particular highly ordered state involving specific temperature differences. Boltzman defined the entropy of the state of such a system as $k\log V$, where k is Boltzman's constant, and V is the volume of the subregion that contains it. It might seem subjective which subregions count as macroscopically indistinguishable, but it turns out that fine details like that don't matter very much given the huge differences in entropy that play a role in underwriting the second law. It doesn't much matter what exact coarse-graining is performed; the result is much the same. The basic idea behind the second law is that if the system starts out in a small-volume subregion, it is overwhelmingly likely that it will end up in a high-volume subregion.[27] The basic reason for this is that there are many more micropossibilities that correspond to that second macropossibility. That is the kind of modal structure involved.

If entropy is to decrease in some system, it must not be a closed system, but rather one that steals some negative entropy from the environment. Since the life of a bacterium is quite ordered, with an order that must be maintained against any natural increase in entropy and indeed is created by growth, and since the bacterium does thermodynamic work in swimming around in pursuit of nutrients, work that requires energy, the bacterium must steal energy and order from its environment. You cannot unscramble an egg without expenditure of energy that comes from somewhere else and so introduces more disorder into the world at large than you remove from the egg by the unscramble. The egg will not unscramble on its own.

There are many puzzles about how the statistical mechanics of large systems of atoms and molecules can deliver the truth of the second law, since the underlying dynamics of physical particles is temporally symmetric.[28] Somehow the irreversible is constituted by the reversible. In fact, the second law is most likely only a probabilistic law, which, properly put, says only that it is extremely unlikely that the entropy of a closed system will decrease, that it is extremely unlikely that an egg will unscramble itself. There was only a very great likelihood that our bunch of particles would end up in a higher volume subregion of phase space. Nevertheless, whatever the technicalities, we know that in reality, living things must steal usably ordered energy from the environment or die, and the metabolic

---

[27] Penrose (2005: 686–702).

[28] Nash (1974); Sklar (1993).

machines already described are a crucial part of the way that *E. coli* and living things in general do this.

What does all this have to do with information, as I promised it would? The basic idea of the link is that order and disorder are measured by information. An ordered structure can carry information which when disordered it cannot; scramble the letters of this paragraph into a pile to see this. Indeed, the form of the second law corresponds closely to the form of Shannon's law regarding the amount of information a signal can carry, which he was even encouraged for this reason by von Neumann to call a law involving entropy. Scanting important complications, a more concrete form of the coarse-graining idea involved in Boltzman's definition may serve to get the key point across. Imagine a gas in a box. There are a huge number of different exact configurations that the molecules of the gas in the box could assume, which they might indeed actually assume as they bounce around randomly over time. Presume that the different exact configurations are all equally probable. Some of those states are intuitively highly ordered. Perhaps the gas molecules spontaneously coalesce into, or begin as, dense little cubic boxes an inch on a side within the larger box, separated from each other by an inch of empty space. There are lots of different ways the molecules could be in detail and still possess, overall, that ordered boxy arrangement. But there are many more detailed configurations of the gas molecules that are much less intuitively ordered. For instance, there is a relatively huge number in which the gas is spread out more or less uniformly across the whole box, though which differ in all sorts of tiny unimportant details. It is overwhelmingly likely that if the gas molecules start out in an intuitively highly ordered state like the first possibility, they will bounce around and end up in an intuitively lowly ordered state like the second possibility. The characterizations of those two possibilities at the intuitive macro level are such that the first is much more highly ordered than the second. But taking all the detailed possibilities into account, the first is much less likely than the second. The second is much less likely to change over time, since it will be preserved across many more different changes in the detailed states of the bouncing gas molecules than the first. And one way to put the link with information is that it will take much more information to convey that the molecules are in the first, boxy macro-configuration than in the spread-out second one. Another is that it will take much more information to specify precisely the exact molecular configuration that underlies in any given case the second spread-out state, given that we already know that the gas is spread out, than it will take to specify the exact molecular configuration that underlies in any give case the first boxy state, given that we know of the boxy state. The boxy macro state is in this sense closer to the details. We might send this information through a signal, hence there is a link with Shannon-information. But observe that the way information is relevant here is fundamentally rooted in another sort of modal phenomenon than the covariation present in signaling. As we noted, there are more micropossibilities that underlie one macropossibility than the other.[29]

So now we have identified ways in which both replication and metabolism in *E. coli* plausibly involve information and so relatively simple forms of modal structure, in the case of metabolism two forms of it. There is a third way in which information is intuitively present in *E. coli*, which more directly involves signaling than the entropy relevant to metabolism.

Recall that such a bacterium has a flagellum whose action is modulated by sensory proteins, so that it swims towards increasing nutrients and flails randomly when nutrients are not increasing. In such a homogenous cell, where all molecules have more or less equal access to each other, there can be a simple three-step signaling pathway, carrying information about the environment. The first component interacts with the external environment and produces a particular molecular messenger that covaries with some environmental condition. That messenger then migrates through the cell and then it generates a relevant chemical effect at some specific site, so for instance the motor of the flagellum turns.

Such systems can harbor complexity, not only because there is a multiplicity of different messengers carrying different messages to different targets, but because different signaling molecules taking the same route have different kinetic properties.[30] Nevertheless, such phenomena involve information and signaling in a relatively intuitive sense, involving a similarly simple modal and dynamic structure as that involved in the specification of proteins by DNA.

So we have found three ways in which information is relevant, involving two types of relatively simple modal structure, in the relatively simple life of *E. coli*. Now we are in a position to tackle the case that most concerns us.

4

Human cells are more complex than *E. coli*. They are "eukaryotic" cells, which means that they are compartmentalized in complicated ways and

---

[29] Loewenstein (1999) claims that the role of information in both replication and metabolism suggests that information is an essence of life.

[30] Weng, Bhalla, and Iyengar (1999).

include separate organs like mitochondria that even have their own DNA. Human cells are, of course, parts of highly complex and sexually reproducing multicellular animals, consisting of trillions of cells performing different roles. While these complexities do not threaten the continuing relevance of information to metabolism, that cannot be said of its other two roles in *E. coli*. We will need to look at some of the details before confronting the roles of yet more complex forms of modal structure in human life.

The various cells of a complex multicellular animal share the same DNA. But they are highly specialized and differentiated. In relatively simple creatures like *E. coli*, it is reasonable to say that DNA carries information about the construction of proteins and cells in a fairly straightforward way.[31] But how do different sorts of cells in a human, when they are developing or when they are mature, produce just the right sorts of proteins in the right sorts of proportions suitable to their differential nature? Does that undercut the relevance to the human case of the way in which DNA carries information about proteins and cellular structure in *E. coli*?

Focus for the moment on the first question. In multicellular animals with differentiated cells, there are developmental and maintenance processes that modulate which proteins are created at which times and in what contexts, by turning on and off the synthesis of proteins that could proceed on the basis of DNA that is present in all the cells.[32] For instance, when there are certain sorts of preexisting chemical gradients, along that gradient there may be differential "expression" of DNA. The DNA may create one kind of protein at one location along the gradient and another at another.[33] One thing that allows this is that there is a very complex set of mechanisms that control the activity of protein synthesis. The activities of our genes are tightly regulated by elaborate complexes of proteins that assemble on our DNA. There is an enzyme, RNA polymerase, that transcribes DNA into messenger RNA. And there are transcription factors, often consisting of many proteins, which control the rate of this. There are both activators and repressors.[34] All this modulates the creation of proteins from DNA, sometimes on the basis of chemical differences in the cell or its environment. We have already seen that genetic switches play a role even in *E. coli*. But it is a question whether the much greater complexity of these

---

[31] For relevant complexities in *E. coli*, see Carroll (2005: 58–61) and Harold (2001: 79–82).

[32] Goodsell (2010: 84).

[33] Keller (2002: 173–197).

[34] Tjian (1995).

types of processes in eukaryotic cells undercuts the presence of information about proteins in their DNA.

There are also other complex steps that intervene between our DNA and protein creation, for instance proofreading and error-repair of RNA.[35] In human DNA, at least 95 percent of the sequence does not code for proteins, which is very different from the case of *E. coli*. Those regions that do code for proteins are called "exons." Those that do not are called "introns." In most eukaryotes, the parts of RNA that correspond to introns are spliced out after "transcription," after the RNA is formed on the template of DNA, and alternative splicing, the production from the same transcript of different RNA segments, which encode different proteins, occurs. There are also other types of "editing" of RNA.[36] Also, protein-coding sequences in DNA have no clear beginning or end, so there is what is called "frameshifting," in which overlapping stretches of DNA generate different proteins. In other words, some parts of DNA code for more than one protein. Also, there can be transcription of functionally active RNA that does not code for proteins.[37]

What's more, in eukaryotes, there are yet further reasons, beyond those present in the case of *E. coli*, to believe that DNA is not the only vehicle that is important in heredity. There is some evidence that some traits are passed to offspring through RNA and not DNA.[38] And there seem to be other sorts of epigenetic inheritance due to chemicals that stick to DNA but do not alter its sequence.[39] Living things in general inherit more than their DNA, and in multicellular eukaryotes, there is even more complexity of this sort.[40] And in sexually reproducing eukaryotes, like humans and trees, the relation between a given bit of DNA and protein synthesis is yet more complicated, because there are pairs of chromosomes inherited from two parents, and hence Mendelian dominant and recessive genes, and other analogous phenomena recognized in classical genetics, play their characteristic roles.

It is not unreasonable to say that the DNA of *E. coli* carries information about its proteins, and even beyond that about its functional parts and overall structure that these proteins constitute. There are two sets of possibilities, for instance different forms of DNA and different proteins, that covary in a directed way. In a recognizable sense, a given form of DNA necessitates a given protein. There is of course a certain background that is needed for this directed covariation to occur, a cellular environment

---

[35] Keller (2000). [36] Sarkar (1996). [37] Gibbs (2003); Pearson (2006).

[38] Rassoulzadegan et al. (2006); Soloway (2006). [39] Gibbs (2003). [40] Griffiths (2001).

with ribosomes of certain sorts. But when such a background is relatively fixed and commonly present, the simple modal and dynamical structure that constitutes one sort of information is plausibly present. Even simple genetic switches of the type regulating enzymes that break down lactose in *E. coli* do not seem to undermine the existence of this information. It is unrealistic to expect deterministic links between a possibility of one set and a possibility of another in a living thing. The tiny mechanisms in question are really just probabilistic biases towards certain characteristic outcomes in largely random processes. Nevertheless, a sufficiently high probability of a certain possibility from a second set given a certain possibility from a first set does intuitively suffice for information.

However, the relationships between human DNA on one hand and proteins and especially functional parts of cells on the other are considerably more complex. Some of the complexities do not *obviously* disrupt the presence of information, but may give us reasonable pause about that.

For instance, if protein-coding DNA is simply turned on and off in certain contexts, as in *E. coli*, it plausibly carries the same information about proteins throughout. But the cells of a multicellular animal share DNA but differ in important and significant ways, from neurons to liver cells. There are highly complex variants of this sort of switching that occur in multicellular animals. Development and growth in such creatures is systematically governed by biologically standard genetically encoded mechanisms for turning protein synthesis by DNA on and off in useful ways, often on the basis of spatial cues, which are the focus of evolutionary developmental biology.[41] There are relatively specific forms of DNA found throughout multicellular life, with little variation, that serve throughout to shape overall body development, to organize appendage development, and to guide the development of things as functionally specific as eyes and hearts despite great variation across species in the detailed nature of eyes and hearts.[42] In development, these genes are turned on and off by genetic switches considerably more complex than those controlling lactose metabolism in *E. coli*, that integrate the action of all sorts of chemical activators and suppressors in complicated ways, which themselves often reflect complex spatial features of the growing animal or its environment.[43] This elaborately controlled systematicity is a two-edged sword. It may deepen our intuitive sense that DNA carries information about proteins through its very systematicity, but it involves a much more complex dynamical and modal structure than that present in the genetic switches

---

[41] Carroll (2005: 53–131). [42] Ibid., 61–79. [43] Ibid., 109–131.

of *E. coli*. And it may undermine our sense that certain protein-specifying DNA codes also for quite specific structures performing quite specific functions in an animal, since the causal link is more tenuous and conditional in that case than in the case of *E. coli* and lactose metabolism.

We've been focusing on just the turning on and off or analogous modulations of the synthesis of specific proteins. But there are other complexities like RNA editing and frameshifting in eukaryotes, not to mention significant epigenetic inheritance and the transcription of functional RNA that does not code for proteins. The mechanism that intervenes between DNA and proteins in eukaryotes does not just function like an intelligent spigot. Still, one kind of complexity that would not fatally undermine the presence of correlation-based information, despite dynamically relevant shifts in a background required for specific forms of correlation, would be a situation in which the background simply shifts between two relatively stable states, with quite different correlations assured by each, rather like the action of a shift key on a typewriter. It may still make sense to talk of different information present given each of the shifted backgrounds, or conditional information contingent on the background.

However, on the other hand, if there are hundreds of dynamically relevant background states, or perhaps hundreds of different types of thousands of variant background states, that are all dynamically relevant, all supporting different merely probabilistic correlations, then the relatively simple modal structure characteristic of information is plausibly lost. It seems to me that this often occurs in the chemical dynamics of complex multicellular eukaryotes, especially when we are focused on the production of functional cell parts, but even sometimes when we are focused just on the production of proteins. And it may well be that this means that even our DNA is not information-bearing in the simple manner of *E. coli*. Of course, it doesn't really matter what we call what is going on in this case, as long as we don't call it something that misleads us about what is really going on, about what we already know is really going on. We know that the relevant correlational dynamics for human DNA is very complex, and it is reasonable to worry that talk of information is misleading in this context.

There is a great abundance of crosscutting but relevant modal structure present in such cases. There is a wide variety of crosscutting facts about what would be true at any given time in such a biological structure if certain facts were held fixed and others varied in certain ways rather than others. And there is a great variety of facts about how the structure would vary over time given various alternative inputs if various modal alternatives of the first sort were in place. To tag this fact, I will call this a case of "complex modal dynamics." I won't attempt a suitably general characterization of complex modal dynamics, or even of the somewhat specific sorts found in human replication. But certain specific forms of complex modal dynamics will begin to matter to us during the third step we will need to take in this chapter.

5

Here's where we are so far, after the second step in our three-step development, in this chapter, of the first of the three major features of my proposal about our phenomenal consciousness. First, our phenomenal consciousness is crucially that of agents, who have an abstract modal structure. There is a range of possible sensory stimulations to which we are differentially sensitive. There is also a range of different possible causal things we may do. These are linked, but there is a variety of possible ways that, given a specific sensory stimulation, we may respond by doing something.

Second, our phenomenal consciousness is crucially that of a living thing on earth. We are constituted out of living cells of familiar types, and which crucially involve both metabolism and replication. Metabolism involves the two sorts of modal structure we have noted, in both bacteria and us. One constitutes the presence of information. But in humans, though not bacteria, replication involves a more complex form of modal structure, a complex modal dynamics that I haven't attempted to delineate.

There are questions about how these points fit together. One natural link between our crucial agency and our crucial metabolism is evident. Animals only act in the world through the expenditure of energy. But the link between agency and replication is less obvious, and indeed it may seem that life's metabolism and replication may easily come apart, as they do in viruses. But I think there is such a link.

We have noted and dismissed attempts to deliver qualia by causal-informational relations refined by appeal to evolutionary history. But I think that replication is crucial to our phenomenal consciousness and our agency in another way. We are complex organisms, which harbor a great diversity of cells but are naturally unified agents. And it is the way that cell replication figures not so much in our reproduction and early growth, but in the maintenance of our integrated unity despite the diversity of our cells, that allows us to be so. Even if some human somehow miraculously arose randomly out of a swamp two days ago, this integration occurs right now.

There is also a second relevant consideration. Given the way we have traced that the detailed mechanisms of our life work, it seems sometimes arbitrary to class what happens in us as specifically a matter of either metabolism or replication. The little machines of replication expend energy, and the mechanisms by which metabolism operates are genetically modulated. We will also see that genetic switches even play a role in the working of our neurons, as we turn now to the third step in this chapter's progress, in which we consider the appropriate analogue in eukaryotes like us of the simple signaling that may help constitute recognizable agency in *E. coli*.

6

The third crucial point about our phenomenal consciousness articulated in this chapter is that it involves the modal structure of our neurons.

In complex cells, there is a more complex sort of "signaling" that occurs even within a cell. In *E. coli*, a certain chemical signal will migrate through the cell, signaling the presence of a certain environmental feature, and generate a certain result. But in complex eukaryotic cells, there is often relevantly complex interaction between different chemical-signaling pathways, and the compartmentalization of the cell can cause the same type of chemical pathway to carry different signals, as in different wires, and there are also physical scaffolds that direct and modulate signals.[44] But the crucial point for now is that multicellular eukaryotes also involve complex signaling and coordination *between* cells, which involve analogous complexities, some of which are especially important to animal life on earth, and especially to their phenomenal consciousness.

Animals on earth do not merely possess highly differentiated eukaryotic cells due to replication and growth, but repeat in an especially striking way at the level of the whole organism some of the most striking features of simple creatures like *E. coli*. Animals move in a directed way through their environments, doing the thermodynamic work that requires. And in neuralia like us, animals unlike sponges who have neurons, and also who, unlike jellyfish, have unified nervous systems, this involves their nervous system in a way which is crucial, I claim, to their form of phenomenal consciousness.

Multicellular creatures as well as some clusters of single-cell creatures exchange what might be thought to be information between cells in a

---

[44] Scott and Pawson (2000).

number of relevant ways. Cytokines are local chemical messengers between cells, and could be said sometimes to warn of danger or inform in other ways when relevant activity is needed.[45] Also, hormones in some sense carry signals through the blood.[46] But in animals on earth, the most striking sorts of sensitivity to environmental conditions directing action of the whole organism involve nervous systems that link sensory organs focused on environmental objects to motor organs like legs or fins that allow active and directed motion, and to other motor organs like hands or biting mouths that allow for the manipulation of such objects.[47]

Since all animals with such nervous systems, not just humans, are the most intuitive locus for phenomenal consciousness, we should attend to some features of how they all work, and the modal structure that involves, in the third and final step of this chapter. The very rough story of how they all work, but focused for expository efficiency on the human case, goes like this.[48] In humans, the nervous system involves approximately 100 billion cells. The cells of the nervous system vary, but the most paradigmatic cells involved are neurons. These are cells with an electrical gradient between their inside and outside, a difference in electrical potential that is discharged in the "firing" of the neurons. Neurons have a typical shape. There is a bumpy structure of dendrites on one end of the neuron, and a long branch-like axon leading outwards on the other. Some neurons have many axons, and some have none, but one is the typical case. Different neurons connect at gaps called "synapses." The axon of a given neuron in humans typically synapses on the dendrites of about a thousand other neurons.

Neurons fire, on their own, in a relatively steady pacemaker way. And when they fire, that modulates the release of various chemical neurotransmitters from their axon into the synapses, neurotransmitters that migrate across to the dendrites of other neurons, where they bind to the dendrites in chemically relevant ways. The firing rate of a neuron is modulated by the neurotransmitters it receives from other neurons, and also sometimes by direct electrical connections with other neurons. And its firing in turn modulates the firing of other neurons.

Certain specialized neurons are sensory receptor cells, whose firing is modulated directly by environmental factors of specific kinds, and certain neurons are motor neurons. They synapse on muscles, which they control and contract by their firing. So when certain sensory receptors are

---

[45] Goodsell (2010: 87).

[46] Ibid., 100.

[47] Trestman (2013) calls these "complex active bodies."

[48] Bear, Connors, and Paradiso (2016); Kandel et al. (2013).

stimulated in certain ways, their firing is affected in a way that affects the firing of other neurons, and so on all the way to the motor neurons whose firing leads muscles to contract in certain ways.

Now consider a slightly more specific level of detail, with once again a special focus on human neurons. During the regular, pacemaker firing of neurons, an electrical spike, called the "action potential," flows out from the cell body of the neuron down its axon. This is how that works: There are ion pumps, and also ion gates that can be opened and closed, in the cell wall of the neuron, which create differences in the concentration of sodium ions and potassium ions between the outside and inside of the cell, generating voltage differences much as in a battery. There is a negative base potential of about 70 millivolts, which is called the "resting potential." When there is a slight positive push in the resting potential, that opens sodium gates, and positive sodium ions rush in down their concentration gradient, spiking the membrane potential in a positive direction, to about 50 millivolts positive. That is the action potential, which in turn causes potassium gates to open and positive potassium ions to flood out of the neuron down their concentration gradient, reducing the membrane at that point to even slightly below the resting potential.[49] And this spike moves on down the axon. That is what it is for a neuron to fire. It might be thought that firing of neurons registers sensory events, and carries information about them. But there are many difficulties for this idea. There is pacemaker firing in a neuron. So it isn't at all plausible to say that each firing of each neuron registers some environmental event. Rather it must be at best some inflection of the pacemaker-firing pattern that can be said to "represent," and of course it might not be a single neuron firing in some way that matters. Perhaps rather a certain pattern of firing across some cluster of neurons registers something environmental. But there are many more relevant complexities than just this. It has been gradually learned that other parts of the brain than neurons are also active in crucial ways. Synapses seem to be strengthened by learning, and indeed sometimes synapses grow a second synapse.[50] So the firing of neurons does not occur in a fixed neuronal environment. And what happens at synapses can be itself dynamically very complex. For instance, it may involve more than one neurotransmitter. And other cells besides neurons matter. "Glial cells" influence where synapses are formed.[51] And a certain sort of glial cell, called an astrocyte, wraps around a synapse and apparently alters what happens there.[52] Even myelin, the fatty insulation around the axons of

---

[49] Groh (2014: 56–61).

[50] Barinaga (1999).

[51] Fields (2004).

[52] Sasaki et al. (2012).

neurons that allows them to carry electrical signals quickly, develops in a way that affects functioning and coordination of neurons.[53] Some genes are turned on and off in the nervous system in a way that is dynamically relevant,[54] and there are even jumping genes that are specially active in the brain, that modulate the form of DNA present in some brain cells. Of course we shouldn't forget the action of hormones, which can affect what happens at synapses. There are at least a hundred different neuropeptides released in very large numbers in pulses throughout the brain.[55] Also, there are neuromodulators secreted as well as neurotransmitters, which function like local and slow-acting hormones.[56] The dendrites themselves harbor significant dynamical complexity. Firing is the fastest thing the brain does, but, for instance, glial cells act quite quickly. So these other sorts of dynamical features of the brain cannot be dismissed as mere fixed background. And it may not even be that the firing of neurons is the central locus of information in the brain. Another general possibility is that activity in the synapses is what represents.[57] Perhaps this is especially plausible if glial cell activity is crucial. To find some specific crucial locus of representation or registration of, of information about, environmental events in such a structure would be no mean task; it is far from clear that there is one that is privileged over all others. But this complexity of neural structure, of indeed modal structure, is, we will see, instead useful to the constitution of phenomenal consciousness in another way.

Before we turn to that issue, it is worth noting that other human neurophysiological details, which we will eventually see matter to our qualia, involve the idiosyncrasies of different sensory receptor cells.

In human vision, light-sensitive photoreceptor cells, which are analogous in some ways to cells involved in photosynthesis, are crucial. We will discuss photoreceptor cells at length in Chapter 4, when we consider color qualia. But there are also considerably different sorts of sensory receptor cells that play an analogous role in different senses. For instance, in the "somatosenses," in other words the various senses that might be distinguished as aspects of touch, there are many different sorts of sensory receptor cells. There are muscle spindle receptors that respond to the length of muscle fibers, and Golgi tendon organs that respond to the force supplied by a muscle to tendons. These are set up in a complicated oppositional structure matching the mechanical oppositional structure of joints and limbs. And their action potentials are affected by the opening

---

[53] Fields (2008). [54] Loebrich and Nedivi (2009). [55] Leng (2018).
[56] Nusbaum et al. (2017). [57] Cao (2014).

and closing of ion channels through stretching.[58] Also there are skin senses, which involve a variety of different specialized sensory receptor cells that respond to different sorts of stimuli, for instance vibration or pressure.[59] In hearing, things work in yet a different way. Compression waves in the air vibrate the eardrum, which in turn, through the mediation of tiny bones, vibrates the basilar membrane in the ear. That has differences in thickness and rigidity along its length, which means that it vibrates at different places more for certain frequencies of sound waves rather than others. And attached to it along its length are various hair cells, which are sensory cells tuned to fire at different sound frequencies by their position on the basilar membrane. There is a kind of fixed roof, the tectorial membrane, against which they are pushed by the vibrations in the basilar membrane, and this causes the hair cells to fire by mechanical means.[60] The organs of balance in the ear involve other hair cells.[61] They include three fluid-filled semicircular canals oriented in various ways, which respond to motion of the head because the fluid drags behind the motion of the canals, hence engaging hair cells. And there are also the utricle and saccule, filled with little stones that function in a similar way to detect linear motion.

Of course, motor neurons are attached to certain muscles attached in turn to a skeleton in an oppositional way at certain joints. So there is complex muscular-skeletal structure that inflects what is done on the basis of neural firing as well.

The immediate point to note about all this is that there is considerable complexity of detail relevant to the way the nervous system bridges environmental stimuli and action in humans. Things are much more complicated than in the simple case of chemical signaling of environmental conditions in *E. coli*. It seems appropriate to speak of informational signaling in *E. coli* when, in response to the increasing or decreasing nutrients in its direction of travel through its environment, chemical signals migrate from the membrane of the bacterium so as to control the motor of its flagellum. This is a very simple form of modal structure. A chemical sign correlated with an environmental condition, and which generates an appropriate response, is either present or absent. It is a kind of binary signal. And it plausibly carries information about the environmental condition in question. Humans, like other animals, have a nervous system that mediates between stimulations of sensory receptor cells and motor action controlled by motor neurons. And this is somewhat analogous to

---

[58] Groh (2014: 52–54).

[59] Ibid., 64–65.

[60] Ibid., 110, 139–141.

[61] Ibid., 180–184.

what happens in the bacterium. But the dynamical structure of the nervous system is fantastically more complex than the simple chemical signals in *E. coli*. It is complex in some ways analogous to those that may undermine our certainty that human DNA carries information about functional structures and even proteins. In fact, I think it may be more relevantly complex even than what happens in the growth and development of multicellular animals, that it is a more dramatically complex form of complex dynamics. In the case of DNA and protein production, even in humans, we know roughly on what sorts of plausibly information-bearing entities to focus. We know how to identify analogs of words that bear information. We know where to seek correlations. There is for instance DNA, and while it has overlapping stretches and there are questions about what constitutes a functionally active gene,[62] we have a sense of how protein-coding uniformly works across life. RNA is relatively closely tied to DNA, even in eukaryotes, and there is the action of ribosomes that is closely modulated by the form of RNA in certain ways, and constrains our sense of what sorts of correlations and so information are relevant. But there are many different things in the human nervous system that could be plausibly said to probabilistically covary with one or another feature in sensory stimulation or the environment, if a certain background is fixed, and which have some effect on motor action given that background, and also many different things that could be considered such a relevant background. Many such backgrounds would be temporally unstable, but in the brain this should be expected in any case where correlation could deliver intuitive information about the environment, since the brain is a very dynamic structure, and any of these backgrounds could vary in complex ways. There are many more things in the brain that could be plausibly considered dynamically relevant differences or carriers of information in this latitudinarian way than in the case of DNA and protein, and which overlap in complex ways. For instance, there is the firing of a neuron, and the firing pattern of that neuron, and the inflection of the firing pattern of that neuron, and a pattern of firing involving many neurons, and a firing pattern involving some but not all of those neurons and perhaps some others besides. And each could be dynamically relevant to motor action in some way as long as one sort or another of background is held fixed. Indeed, absences of certain sorts of firing might be relevant in similar ways. And there are also various sorts of features of

---

[62] Griffiths and Stolz (2007).

the environment or stimulation with which different such things would covary if various sorts of backgrounds were held fixed. And there are also quite different types of things that might be considered carriers of information or dynamically relevant differences in the brain. For instance, even blood flow in the brain varies in a way that seems to be directed and affects the activity of neurons.[63] We have already noted that glial cells, which were traditionally thought to merely provide structural and metabolic support for neurons, modulate activity at synapses.[64] A central locus for correlation might not be neural firing but what happens between neurons, at synapses,[65] which may be modulated by hormones. There are lots of specific aspects of what happens at a synapse that could be dynamically relevant or covary with environmental conditions if some sort of background is held fixed. Of course, what happens with synapses can affect neural firing, and neural firing can affect what happens at synapses. There are many types and levels of modal and dynamical grain in the nervous system, which overlap and interact in complex and probabilistic ways. It seems to me that environmental information of so many different types could be attributed to something in such a structure, if different arbitrarily specified backgrounds were held constant, that it may make no sense to attribute any specific semantically or psychologically relevant information at all.[66] But in any case, we aren't interested in information unless it delivers qualia, and we have already seen that it is unlikely to do that in at least the case of color experience, because there aren't the needed features in the environment to represent. Once we relax the constraint that relevant dynamical structure must be correlated with environmental features, even more dynamical complexity is apparent. It is not like the case of thermodynamics either, where there are relatively determinate ensembles of substates that correspond to each macrostate. There is rather an overlapping dynamical mess. And of course the human nervous system is also more complex than what happens in *E. coli* in another way, simply because there are hundreds of billions of complexly interacting neurons involved, and many trillions of synapses. And the firing pattern of the average neuron is modulated by the integrated effects of the firing patterns of a thousand other neurons.

---

[63] C. I. Moore and Cao (2008).

[64] Cao (2014: 895).

[65] Ibid.

[66] Cao (2012) argues that many different sorts of information are present in the brain, in a somewhat richer sense of "information" than I presume, but that only at the level of the whole organism is there suitable correlation of that richer sort with the world.

All of that very complex structure mediates between sensory stimulation and motor action in humans. It is certainly a case of very complex dynamics. The nervous system involves a much more complicated modal structure, relevant to its dynamic evolution over time, than the simple directed covariation between two sets of possibilities in a fixed linkage that underwrites information. The modal structure of the human nervous system is a kind of wild generalization of that simple form. It is true that the neurophysiological structures that intervene in humans between sensory stimulation and motor action are linked by causal action and exhibit counterfactual dependencies, such that states closer to motor action often depend on those closer to stimulation. But the links are often merely probabilistic tendencies. And just as there are different sciences of ordinary material objects that capture interesting causal regularities at different levels of scale for certain kinds of phenomena, for instance geology and chemistry and physics in application to volcanoes, we should perhaps expect various causal regularities at different levels of grain to be relevant in different regions of the highly complex and variegated human nervous system. What's more, the modal structure internal to occurrences at some specific level of grain may be dynamically relevant even when it is not causally relevant to what happens at the specific level of grain under consideration. It would not be surprising if there is not a simple hierarchy of interesting regularities, as familiar models of the sciences suggest, but rather spatially overlapping and crosscutting causally relevant structures. And what counterfactual dependencies there are within even a given level of causal structure will depend on what other potentially variable features at that level are presumed as fixed background, even under the assumption that other levels of structure are held fixed. There are relevant ramified counterfactual dependencies expressed by conditionals within conditionals, and counterfactual dependencies of various sorts that cross levels of grain. And real constraints on possible alternatives to things that occur or might occur at any point in our neurophysiology may also be quite relevant to its modal structure, and exhibit analogous counterfactual dependencies. This is a much more complex situation than that of the simple dynamical structures that intuitively carry information in *E. coli* "sensation."

But from our perspective, this is just as well, because we seek a neurophysiologically based account of our qualia, to which we have already seen reason to think information is not very relevant. Still, what useful things can be said about this complex neurophysiology? What particular kinds of modal structure in the nervous system should we seek to isolate as relevant to our concerns?

7

The apparent modal structure of color qualia, and other forms of phenomenal consciousness that we will consider, will help us to answer these questions. But let me first note some contrasts to the answers I will be developing.

Since your nervous system is a physical entity, it is of course true that a full characterization of its current microphysical state (including that of any causally relevant environment) together with fundamental physical laws (which, however, may be fundamentally probabilistic) will suffice to explain its dynamical trajectory over time. Moreover, its microphysical basis at a moment plausibly constitutes whatever it is at that moment. But the psychologically relevant complexity of that structure at a moment and over time, the sort relevant to qualia, may well involve some sort of abstraction from that very exact level.

There is a familiar model of how this might have worked: functionalism. There might be some abstract level of relevant psychological states, with each state characterized by its causal relations to other such states and possibly to environment causes or behavioral effects. And then, beneath that abstract level, there might be a variety of alternative neurophysiologically detailed states that might constitute or realize each such abstract state. Through horizontal relations among abstract states, this conception partakes in something like the modal structure characteristic of informational correlation, and through vertical relations between abstract states and their detailed realizers it partakes in something like the modal structure characteristic of thermodynamic phenomena.

But this standard conception fundamentally misrepresents the abstract structure of our neurophysiology relevant to phenomenal consciousness, which does not plausibly exhibit this tidy regimentation into discrete layers. Rather, as I've noted, there are causal and other determination relations that scramble up any such discrete layers. One bit of evidence for this scrambling throughout biology is the fact that attempts to model biological phenomena in realistic ways frequently need to fit together, in a way that is quite difficult, different abstract models at different levels of scale that on the surface involve contradictions. For instance, sometimes boundary conditions held fixed at one level of scale need to involve conditions capable of variation at the other level if comprehensible and tractable models are to be developed.[67] But let me focus more specifically

---

[67] S. Green and Batterman (2017); Noble (2008, 2012).

on the central case already in play, that of color qualia. We noted in Chapter 1 the reasons to expect that, in the case of color qualia, a different sort of relation between layers is of significance than that which functionalism develops, an expectation that will be vindicated in Chapter 4. When you are in the state of experiencing something as red, that plausibly involves a certain specific pattern of neural activation in the context of a background structure that, if activated in some alternative way, would constitute the experience of green. There are two layers of structure in the aspect of your nervous system responsible for your color vision. One background and relatively fixed layer is shared by more specific states of actually seeing green and actually seeing red. And it assures even when you are actually seeing green that there is the real possibility of your actually seeing red.[68] This explains why the actuality of the green you experience would be locally entwined in being with the mere subsistence of red in the manner of the superworld if only that green were actual. This is not the sort of relation between layers that functionalism deploys. And notice also that it is reasonable to expect that the relatively fixed context of some firing in your neurophysiology might itself change over time at a slower pace, and that given the way in which neural firing affects synaptic connections we might expect cross-level dynamics of some sort.

This is one relatively simple form of what I have called complex modal dynamics, and will be an illuminating model to bear in mind as we proceed through the neural labyrinth in what follows. We need to look beyond functionalism. But consider some other contrasts with familiar ideas.

When we look at the neural details that explain our qualia, it will be in a way analogous to the mechanical tradition in scientific explanation.[69] I can be taken to be working in this tradition, to be attempting a kind of mechanical explanation of phenomenal consciousness. But we should remember that the micromachines in question are more in the vein of statistical biases rather than hard mechanical connections of a familiar sort. So the analogy is limited.

One last contrast: My proposal is also reminiscent of the dynamic tradition in cognitive science, in which the trajectory of a neural or psychological system through a state space is what is central.[70] But in the cases we face, even the nature of the relevant state space may also have a trajectory over time.

---

[68] For an analogous discussion of neural predispositions and prerequisites, see Northoff and Heiss (2015).

[69] Chirimuuta (2014); Craver (2007); Glennan (2017).

[70] Cosmelli, Lachaux, and Thompson (2007); Spivey (2007); Van Gelder (1995, 1998). For skepticism, see Eliasmith (1996).

8

While the modal structure relevant to color vision will be an important model for us as we proceed, still the kinds of modal structure of our neurophysiology relevant to phenomenal consciousness are quite diverse. We will be guided in what sorts of neurophysiological structure to seek by the modal complexities of phenomenal consciousness itself, which we will begin to explore in Chapter 3. Still, it may be useful to attempt now a more general characterization of the kinds of neural modal structure that are plausibly available to us as possible resources. So let me make such an attempt.

There are three steps required to explain the general idea of this, for which the various complexities we have traced in this chapter have prepared the ground. Here is the first. As we have noted, in the case of color experience, real possibilities, which are alternatives to some actual state, are an important aspect of our neurophysiology of color vision. The basic idea is that having an experience of red involves being really capable of an experience of green. There is some relatively stable neurophysiological background which the capacity to see green and to see red share, and then more detailed activity in that context constitutes only one of those experiences as actual. The halo of real possibilities that surround a given fully actual state of experience are crucial to its nature. Over time there may be a sequence of such states in which what is fully actualized migrates among stable real possibilities, as is the case with color vision. That is one kind of relatively simple modal dynamics, and the one which is perhaps most important to us. But it is also true that the nature of the space of real alternatives may vary over time, within a higher-order space of alternatives of possible first-order alternatives. And in addition it may be, as for instance neural firing patterns over time modify forms of neural connectivity, that dynamics at the more specific level can affect dynamics in the containing level, in a cross-level effect. But remain for the moment focused on a single moment in the history of some human's neurophysiology. At any given moment, that nervous system, in its relevant context, is in a given state. But that momentary state harbors significant relevant modal complexities right in the moment of the sort just noted, involving real alternatives to various fully actualized states, and at various levels. That is the first step in a characterization of the sorts of complex modal dynamics we will be tracing in this book. Here is the second. Right in the moment, it is true of the neurophysiology in question that if certain of its relevant details were altered in some specific way, then certain of its other relevant details would also be altered in some specific way. That sort of counterfactual conditional claim, rooted in its current state, might capture relevant modal structure. But it is important that those details often have a modal character themselves. We have just noted examples of this, and there are various crosscutting relationships of that counterfactual sort among such details, as well as more complex dependencies that are naturally expressed by embedded counterfactuals or counterfactuals with extensive conjunctive antecedents.

Here's one relatively simple case of what I mean, putting together the first and second steps: A neuron is firing in a certain way. But there are also real possibilities for alternative forms of firing of that neuron, rooted in its actual state, holding all else at the moment steady and indeed restricting ourselves only to forms of alternative firing that are in some sense natural for that neuron. But its firing is also part of a pattern of firing in a larger cluster of neurons. And the really possible alternative forms of that larger firing pattern might be inflected by the alternative forms of really possible firing that our original neuron might counterfactually assume. It is not just that if the original neuron is in fact firing in one particular way that the whole pattern is constrained. It also might follow from some pattern of firing in the original neuron becoming really possible that there is an alteration in what is really possible for the firing of the complex of neurons in question. There is in the moment in that case a kind of dynamical relationship between real possibilities. Of course, there aren't just neurons and their firing. Our original neuron has many synapses; focus on just one. There are real alternative possibilities for what might be happening at that particular synapse at the moment in question, whose actuality or even real possibility might depend counterfactually, not merely on the actuality of really possible alternative states of, but also alternative sets of real possibilities for, various things: the hormones or neuromodulators in local circulation, activity by local glial cells, the firing and other states of various synapsing neurons, activity and states of relevant dendrites. And, of course, we should not forget the huge complexities introduced by the huge numbers of neurons and synapses in a human brain.

There is a vast complexity of modal dependency relations among all sorts of current aspects of your neurophysiology, which include dependencies between real possibilities for various aspects of that neurophysiology. Those dependency relations among alternative possibilities and actualities constitute all in a moment a kind of modal dynamics; they relate different current possibilities of different aspects of your neurophysiology. Hence, in some sense, they set relevant possibilities for the whole as it currently is. This is very unlike the simply regimented relations between levels characteristic of the modal structure relevant to thermodynamics and functionalism.

Even if all this modal structure, reflecting these first two aspects of complex modal dynamics, is there in our neurophysiology, why think it is psychologically relevant? Because, as we will see, though in some cases more clearly than others, the explanation of our phenomenal consciousness requires something like this complexity. To solve the hard problem of consciousness, we are going to need to push fairly hard on both ends of the equation, on our understanding of our neurophysiology and on our understanding of our qualia. This sort of complex modal dynamics will help with the first part of that task, and help deliver what we will see we need when we properly pursue its second part.

But this structure at a moment is not all there is to the kind of complex modal dynamics at issue. Now let's put this monster in motion, and take our third step in its characterization. We've already noted that there may be temporal evolution not just in some fully actualized state among various real possibilities, but also temporal evolution in the nature of those real possibilities. And the facts of what those forms of temporal evolution will be may be, at least to some extent, rolled up in the current state of your neurophysiology. But there is, second, also a vast temporal dynamical complexity rooted in your current neurophysiological state that even a behaviorist can capture. There is a range of alternative complete stimuli you might experience, each which would be linked, with various probabilities, with various behaviors, given your current neural state. And third, there are also various sorts of temporal trajectories your intervening neurophysiology itself would undergo, so to speak on its own, under various hypothetical internal conditions, hypothetical conditions that include the actualization of various real possibilities that now obtain given your current neural state, but that also include the obtaining of alternative ranges of real possibilities for certain aspects of your neurophysiology, as we recently discussed. There might be temporal dynamical links not just between fully actualized neurophysiological states but also between the obtaining of various real possibilities for a neurophysiology at different times. Of course, we must mutually embed these various sorts of temporal complexities, the internal developments assured by your current state, the externally linked possible developments involving stimuli and behavior, and the internal possible ones linking alternative neurophysiological states over the similarly short term, if we are to understand psychologically relevant temporal dynamics over even such a short term. In addition, there may be interesting facts about the temporal evolution of someone's psychology and neurophysiology over longer time scales, which are not so closely bound up with the current state of that neurophysiology in the way I have focused on.

So here we are: When Ramón y Cajal discovered neurons, he also discovered that there was a directionality of neural conduction, of neural currents. Other causal influences often have a similar directionality in the brain and nervous system. But this no more automatically legitimates speech of information or signaling in such a case than in the case of the current from a switch to a light. And the "wiring" of the human nervous system is not much like that of electrical systems or familiar machines anyway. In some animals, there may be a direct neural link from some sensory receptor stimulation to a specific avoidance behavior, say from touch felt by the siphon of the giant sea slug to withdrawing its gill, which is not unlike what happens in bacteria.[71] But our neurophysiologies are not very much like that. Rather, what intervenes in us between stimulation and action is very complex modal dynamics, of the sort we have been considering. This is not to say that there aren't important uses of the notion of information in cognitive science and psychology as in biology, but there is much more to the modal structure of our nervous system than that simple form. But we have already encountered reason to think that information will not be much help in explaining our qualia anyway. So it will be useful that all this modal complexity of the nervous system is available to us.

9

Here is a summary of the three steps of this chapter. This first main feature of my account of the phenomenal consciousness involved in our sensory experience involves these three points. First, our phenomenal consciousness is crucially that of agents, who have an abstract modal structure we traced, linking possible stimulations and possible actions in various possible ways. Second, our phenomenal consciousness is crucially that of a living thing on earth. We are constituted out of living cells of familiar types, which crucially involve both metabolism and replication. Metabolism involves two sorts of modal structure as we have noted. One of these involves information. But in humans, though not bacteria, replication involves a more complex form of modal structure. Third, our phenomenal consciousness involves a complex modal dynamics in our neurophysiology of the sort sketched, which is what mediates between

---

[71] O'Shea (2005: 92–97).

stimulation and action in us. In particular, we have seen reason to believe that it sometimes matters that actual neural firing occurs in a context in which other sorts of firing are really possible given a relatively fixed state of our neurophysiology.

So far we have been pushing up from below, developing a certain modalized understanding of the neural and living resources that plausibly constitute our own phenomenal consciousness and that of other animals with unified nervous systems here on earth. In Chapter 3, which develops the second main feature of my proposal about our phenomenal consciousness, we push down from above, with an exploration of the modal structure of that phenomenal consciousness, which will also allow us to identify the specific neural locus of our phenomenal consciousness, and relevant forms of modal structure in that neural basis. Eventually, in the third key feature of the proposal, we will attempt to put these two things together, in a detailed explanation of the specifics of our sensory qualia by the specifics of our sensory neurophysiology, by appeal to matching modal structure.

CHAPTER 3

The Whole Nervous System Model

Chapter 2 developed the first feature of my account of human sensory qualia and phenomenal consciousness, a modal structural understanding of animal life on earth, and in particular animals with unified nervous systems. This chapter sketches the second feature of that proposal, a general account of where, in the unified nervous systems of animals, phenomenal consciousness is located. I call it "the whole nervous system model."

1

Say we've shrunk, and now fly about the stupendous cellular and material complexity of some human animal. Our gigantic friend has sensory receptor cells of various sorts at the input side of a vast and complicated neurophysiology, which eventually terminates at the output side in motor neurons, which control the contraction of muscles arranged in certain ways, so the human acts one way or another in the world. That neurophysiology and those muscles are composed of living human cells, with the kinds of chemical and modal structure discussed in Chapter 2.

What in that massive neurophysiology and musculature constitutes phenomenal consciousness in that animal, the details of which are relevant to details of their sensory qualia?

Put bluntly and crudely, all of it. All of that constitutes the phenomenal consciousness of the human. That is what I mean by the whole nervous system model.

Who is wrong if I am right, beyond dualists and those who accept forms of physicalism we have already dismissed? Those who articulate a more specific "NCC," a specific neural correlate of that consciousness;[1] those who articulate neural conditions for phenomenal consciousness that not all

---

[1] Chalmers (2010: 59–100).

https://doi.org/10.1017/9781009514774.003

animals with unified nervous systems meet. This and the first feature of my proposal together imply that all living things on earth with unified nervous systems mediating between stimulation and action have phenomenal consciousness of some sort, although this is not meant to imply that no other entities can. There are two types of popular candidates for the NCC. Some propose that the activity of quite specific neurons constitutes conscious experience, for instance that the NCC is activity in coalitions of pyramidal neurons, which link different regions of the cortex.[2] On another type of view, a certain sort of neural activity, and not certain specific neurons, is what is crucial. For instance, the transient synchronized firing of millions of neurons distributed in different regions of the brain might be the NCC.[3] I, to the contrary, claim, baldly put, that the whole nervous system constitutes our phenomenal consciousness.[4]

Well, perhaps that's a little too blunt and crude. There are complications and qualifications.

Our neurophysiology includes the autonomic nervous system, which regulates the heart and the glands, and that seems obviously irrelevant to phenomenal consciousness.[5] But notice that we are focused on just the part of the nervous system that connects sensory receptor cells and motor action, so that's already been ruled out. Unless, as indeed is plausibly true, glandular hormones are relevant to the dynamics of the nervous system that connects sensory receptors and motor neurons, and so they matter after all in much the same way that nerves connecting sensory receptor cells and motor action matter. A second complication is that our focus is specifically on the phenomenal consciousness involved in sensation, not all of phenomenal consciousness. But we will see in this chapter these distinctions become somewhat blurred. And it is worth understanding the broader context of that specifically sensory consciousness in any case.

A third complication is that relevant arrangements of the muscles implicitly involve the skeleton, while the skeleton and also the muscles that I've explicitly deployed in the proposal are not intuitively part of the nervous system or what constitutes our sentience. But, I reply, what we do

---

[2] Crick and Koch (2003); Koch (2004).

[3] Greenfield and Collins (2005), developing work of Singer and Llinas.

[4] Northoff (2023) and Northoff and Zilio (2022) develop another account according to which whole brain activity is crucial to consciousness. On this alternative view, not just the nervous system but also the environment is crucial, and relations among spatial and temporal patterns of firing are fundamental. These do not exhaust relevant modal structures, but, on the other hand, go beyond them in certain ways.

[5] Bear, Connors, and Paradiso (2016: 521–549); Gazzaniga, Ivry, and Mangun (2002: 93–94).

out in the world, which depends on our skeleton and muscles, is relevant to aspects of our sensory consciousness, but only in somewhat limited ways, and not mostly in the cases to which we will centrally attend. So you can largely ignore this detail of what I'm claiming if that makes you more comfortable. It is also true that some things we do are internal, for instance talk to ourselves in the mind's ear. But as we will see hints of, this is not unrelated to motor action. The spatial arrangement of the sensory receptor cells also matters in ways analogous to the arrangement of the motor neurons and muscles, which my rough characterization doesn't explicitly mention. But I will presume this to be an aspect of what the sensory receptor cells are, and so implicit in that characterization as well.

A fourth complication is that it is reasonable to complain that there must be tiny details of the nervous system or muscles, for instance how a particular calcium atom spins on a particular axis, which are not relevant at all in the structure that intervenes between sensory stimulation and motor output and to phenomenal consciousness in particular.

I agree, though we have seen that dynamics in the nervous system is complex enough that we cannot be certain before we begin what sorts of details are relevant. Still, there are certainly some plausible constraints on what level of detail is in fact highly relevant to phenomenal consciousness. So let me elaborate these, and do so by appeal to modal structure.

There is some range of alternative sorts of relevantly possible sensory stimulations that a human might undergo, which does not presumably include what would happen to them on the face of the sun, or how they would explode in the vacuum of deep space. And there is plausibly a maximum level of the detail in that range that matters, that could make any difference under any circumstances to what they would do. Likewise, there is some range of alternative possible things a human might do out in the world, which does not include flying around the earth by flapping their arms, and there is also plausibly a maximum level of detail of that which is relevantly under the control of the nervous system, at least in any intuitive sense. I don't expect anything below that rough level of detail to matter very much to phenomenal consciousness, although it could conceivably be relevant to some aspects of qualia that aren't dynamically relevant to the explanation of intuitive action on the basis of stimulation.

Still, we should consider explicitly the level of relevant detail of what intervenes between stimulation and action. Even *E. coli* in some sense reacts to sensory stimulation with motor action. If it is moving down a nutrient gradient, a chemical signal will cause its flagellum to flail and it will tumble randomly. But that is very different from the human case, and from the case of animals in general, and not just because no neurons are involved. There is a direct and fixed link between a given stimulation and a given response in *E. coli*. But a human can be in a state such that any arbitrarily chosen discriminable sensory stimulation, from within the range of relevant sensory stimulations, is coupled to any arbitrarily chosen suitably discriminable motor action, from among the range of possible actions.[6] If the red light comes on, hop on your left foot. If the green, sing the Jetson's theme song. Animals in general are flexible in analogous ways, although perhaps not so broadly, instantly, and at will.

But it isn't just a particular function from all relevant global sensory stimulations to all relevant motor actions that characterizes the sensory consciousness of a human at a time, even a function incorporating such idiosyncrasies as being primed to hop around if the red light comes on. That would be a kind of behaviorism of an unintuitive and implausible kind. Rather, a human at a time is in a very complex neurophysiological state that would underlie such a function, perhaps a function incorporating only probabilities of various outcomes given various inputs. And that internal state includes all the complex dynamical structure we noted in Chapter 2. It is important that different complex neurophysiological states underlying the same such function, even such a function incorporating the same probabilities, may introduce differences in phenomenal consciousness. The modal structure of that intervening neurophysiology matters, which includes much more than what assures a particular such behavioral function. Regarding the relevant level of detail of that intervening structure, the first thing to say is that it is constrained somewhat by the constraints on relevant input and output already mentioned. But there is more to say about it, even here at the beginning. So consider that structure more closely.

There are at any given time a variety of different ways the neurophysiology of a given person might be, which are not wildly hypothetical but rather in some sense real possibilities for that person with such a neurophysiology, which are rooted in their complete actual state by local variation in something like the way in which the subsisting parts of the superworld are locally entwined with the actual. But there are differences from the simple case of the superworld. One difference is that the local entwining of the possible and the actual in our neurophysiology involves the entwined and overlapping structures characteristic of complex modal dynamics. That, we have already seen, is quite complex, almost certainly

---

[6] There may be exceptions, say involving torture.

more complicated than the relatively simple structure of the superworld. But another difference, we have seen, is that this does not require, as it does in the superworld, metaphysical oddities like subsistence.

There are great complexities of such modal structure of our neurophysiology, even synchronically. But it is plausible that anything relevant to sensory phenomenal consciousness in any very significant way, in the relatively gross detail we will be discussing here, would have to be a kind of modal structure whose availability is relatively stable over time and even across different "normal" human individuals, and could intervene between sensory stimulation and motor action across that range of people at the level of detail we noted as relevant earlier. But there are still no doubt many complex neurophysiological modal structures of live humans that meet that constraint, and they may underlie individual variations in details of phenomenal consciousness that ordinary talk does not capture easily and that we will be largely ignoring here. These are all real possibilities that do not involve subsistence, but which are rather alternative states available to some neurophysiology in something like that way that a fixed structure of neurons could involve different patterns of firing. The actual state of the neurophysiology can be preserved, at some relevant level of abstraction, and yet be different in certain ways. And those differences are the real possibilities in question. I am not talking about wildly hypothetical metaphysical possibilities regarding alternative forms of neurophysiology a given person might have and yet remain that person, say with a different sort of nervous system perhaps more like that of an octopus than any real human. I am not talking about alternative neurophysiologies within the standard human range that that given person might have and yet remain that person. I'm not even talking about the form their nervous system will actually take in the future or the past. I'm talking about what are quite real possibilities for that person given the actual neurophysiology that intervenes between their sensory receptors and motor action right now. For instance, holding their neurons and synapses fixed, it would be really possible in the relevant sense for different sorts of sensory receptor firing than they actually enjoy to occur, which would reverberate through to motor action with various probabilities. I'm talking about real possibilities of that general sort, which are closely rooted in the actual state of the person's neurophysiology at the time. When we are dealing with the strange metaphysics of the superworld, there are metaphysical resources available that help fix modal structure in a determinate way that are not available in the actual states of someone's current neurophysiology. Yet we will see that some of the apparent modal structure of the superworld is mirrored in the less metaphysically robust but actual nature of the neurophysiology with at least substantial determinacy.[7] That is how that neurophysiology will deliver experience as of such a superworld, which I have said is relevant to our experience of sensory qualia.

Now let me summarize my response to this worry about levels of detail and add some final refinements. There are three critical components of the modal structure of someone's current neurophysiology, involving real possibilities that I hold are relevant to their phenomenal consciousness: First, in a given person at a time, there is a range of relevant ways that they might receive sensory stimulations. It presumably does not include abnormal conditions that would incinerate the person in question, or electronic fields that would manipulate their brain. It no doubt helps determinacy here that the relevant range is roughly the same for all humans, at least humans of the same size and shape. Second, there is also a range of relevant ways that the person might act by moving in the world, given their musculature and skeletal structure. This is subject to greater difficulties of determinate specification than the range of relevant stimulations. It doesn't matter what they would or wouldn't do if hemmed in by rigid and self-moving body armor, or in the strong gravitational field of Jupiter. Normal conditions matter. But these are no doubt hard to specify. In addition, at any given time the body of a person will be in a particular stance, which will constrain their movements, and which is ordinarily accompanied by a particular reduced set of possible sensory stimulations, indicating that their body is thus and so. And they will be in a certain environment that may inflect their action even if not so distortingly as gravity on Jupiter. We need a more hypothetical range of movements than one fully constrained by someone's actual current stance and external state, although on the other hand, in some cases, some aspects of that stance and state may be relevant constraints. While we expect some help here with determinacy by appeal to normal human motor capacities, "normal" people differ in motor capacities in significant ways.

But the biggest difficulties in determinately specifying the relevant modal structure of someone's actual neurophysiology have to do with a third thing, with what intervenes in that person at that time between their sensory stimulations and motor action, and which would hypothetically

---

[7] I will often try to capture relevant modal structure by appeal to counterfactual conditional claims, but I am not confident that they are adequate instruments for our task, and in any case they often involve treacherous indeterminacies of truth. See Mendola (2014: chapter 2). But one thing that helps us here is that many of the relevant conditionals focus on relatively localized circumstances.

link something from the range of admissibly relevant sensory stimulations, with perhaps some probability, to something from the relevant range of motor activity. No doubt that intervening structure would be disrupted in its dynamics by conditions outside of normal, for instance by the presence of strong electromagnetic fields or unusual temperatures, and of course it isn't that it would take the same form even in all normal conditions. I think somewhat differently when I'm warm and when I'm cold. And once again it would no doubt be difficult to specify relevant normal conditions. But we need yet more modal determinacy than they would provide. On my view, the modal structure of the internal details of all that actual intervening neurophysiology can at least potentially matter. So it is potentially relevant for instance what could normally happen in a particular neuron or synapse, or what would happen in that neuron or synapse under various counterfactual suppositions that also must be plausibly constrained to normal alternative conditions. And of course there are many different sorts of overlapping modal structure in the complex modal dynamics of the brain that involve similar difficulties in specification. But one thing that will help us with this complexity is that the kinds of structure recognizably relevant to prominent features of sensory qualia in humans are, as I've said, relatively stable over time and over different normal individuals. Fortunately, we will see that there is enough stability and uniformity of this sort of neural modal structure in most humans that there are relevant things that can be said about the human experience of sensory qualia at the level of detail we will pursue here.

All subtleties aside, here's one way to put the basic idea not only of this chapter but also of this book: What is phenomenal consciousness, at least in a human animal? It is a modal structure that links sensory stimulation to motor action. What is the neurophysiology of that sentient human? A complex thing that has a modal structure, probably at a certain level of abstraction that ignores fine details, and links sensory stimulation to motor action. And those two modal structures are the same thing.[8] Another way to put this idea, in light of the current interest in mechanistic explanations in philosophy of science,[9] is that it is a mechanistic explanation of phenomenal consciousness, but in terms of specifically modal mechanisms.[10] What phenomenal consciousness is in a sentient human is

---

[8] And they are realized by the same thing, if some further level of detail exists in reality.

[9] Craver (2007); Glennan (1996); Machamer, Darden, and Craver (2000).

[10] Might there be some physical details of our neurophysiology that are relevant to the nature of our qualia but are not constituted by the modal structure of that neurophysiology? I don't think so,

a kind of stuff, neurophysiological stuff, though understood at a somewhat abstract level, perhaps abstract enough to make dualists somewhat happy, though not so abstract as to make physicalists unhappy.

Why do I propose this model? First, because it is general enough to capture a natural division between those entities – living animals whose sensation and action is mediated by a unified nervous system – that are most intuitively consciousness, and those that are not. Second, because it provides an arena sufficient to deliver all the resources that we will see are necessary to constitute our sensory phenomenal consciousness. Third, because it underwrites, in a somewhat modified form, a quite traditional view of what provides unity to the mental, namely consciousness, a view favored for instance by Descartes and Locke, and one that deserves at least more elaboration and defense than it currently enjoys. Fourth, because, as we will see, a proper exploration of the structure of our phenomenal consciousness reveals complexities that undercut the familiar reasons why that traditional view is no longer popular.

2

Chapter 1 noted that there is a sense in which this second main feature of my proposal about phenomenal consciousness involves pressing down from below, developing a certain modalized conception not of the physical root of our phenomenal consciousness but of the nature of that consciousness itself. This may not be evident yet, but it is important to the main action of this chapter. It crucially supports the whole nervous system model. This modal structural analysis of phenomenal consciousness itself involves two crucial, linked notions, which I will call "hyperexperience" and "different forms and degrees" of phenomenal consciousness. We will eventually see that these notions are needed to answer natural objections to the whole nervous system model.

Begin with hyperexperience. The superworld of our naive Aristotelian experience has an actual core that is locally entwined with various mere possibilities. The being of any actual phenomenal orange of a hat would be entwined with the being of an at least subsisting unique yellow and unique red. That is part of its essence. Such a structure is present in the apparent content of such an experience, or is at least very close to being present there, which is indeed the root of the phenomenologically based claim

---

because I believe there are no physical details that cannot be captured by modal structure. But that is a long story for another day.

about the essence of phenomenal orange in question. I call modal structure like that, within the phenomenally conscious content of experience, a form of hyperexperience. But we will shortly see that hyperexperience is more complex than the superworld in certain ways, and involves other sorts of cases, and it also suffers certain limitations not present in the superworld. Still, begin with experience of the familiar case of phenomenal color, whose place in the superworld we have already considered.

There is the fully present stuff you experience quite vividly and immediately, including the phenomenal orange of a chair you are looking at in good light, whether or not phenomenal color is really out there in physical reality. That's in the core of your hyperexperience of a moment. But, unlike the actual core of the superworld, the core of your immediate experience does not include everything there is, or even everything that you naively take there to be. It doesn't include the unseen backside of the chair. But, like the superworld, its core is locally entwined with mere possibilities, or rather appears to be, right there in phenomenal consciousness. For example, as noted, the orange that you seem to experience is entwined with the mere possibility of unique red in that very experience, in much the way it would be in the superworld. I think you have mere hyperexperience of that unique red, or at the very least of determinable red, when you experience an orange chair. I think you can introspect that possible red of the chair in a way somewhat analogous to the way you can introspect your fully present experience of the orange of the chair, in a way that is relevant to what it is like to experience orange. Your experience of that orange is the experience of something as having an essence that involves such merely subsisting red, or at least, on the alternative metaphysical model we considered, possible but perhaps uninstantiated red.

It may be hard to believe in the wacky structure of existence and entwined subsistence of the superworld. But I have promised and we will see that the neurophysiological basis of your experience of orange plausibly involves something like a certain specific kind of neural activation, say certain firing patterns, in a relatively fixed neural context that would, should other firings occur in it, constitute the experience of red. The real possibility of the experience of such a red for you is assured by your actual experience of orange, and not in a way that involves wacky subsistence. And it inflects what intervenes in you between sensation and action, though perhaps in a more subtle way than fully actualized experience. So there is also a third party here. The very neural structures already in question assure that there is less than fully vivid hyperexperience, mere hyperexperience that is not fully present experience, of even that merely possible red when you see an orange chair, and which is a part, though not a very dominant or central part, of your phenomenal consciousness, just because it so inflects. There is a kind of match among three things, the wacky modal structure of the superworld suggested by your experience, related forms of hyperexperience, and the fully mundane modal structure of your neurophysiology of color vision that is relevant to what you do.

What is mere hyperexperience, that part of hyperexperience that is not its fully present core? It is less than fully present, but I believe it is somewhat present, in our phenomenal consciousness. It is present in a lesser or different way. It is more than merely what you *could* introspect in some very wide and hypothetical sense of "could." But it is somewhat close to being introspectible in the same immediate way as fully present phenomenal experience that is not actually being introspected at the moment, though it has a somewhat lesser presence than even that.

In fact, in the case of red latent in orange that we are considering, such mere hyperexperience has an especially lesser degree or form of presence. So this case with which we begin, which we do so because it is related to what has come before, is in some ways the most doubtful of all the forms of mere hyperexperience I will be suggesting in this chapter. And so now we come to the third main idea of this chapter, that there are different forms and degrees of phenomenal consciousness.

This is a somewhat familiar position in a somewhat familiar debate. There are those who think that phenomenal consciousness is either/or, and those who think it can be a matter of degree.[11] And there is relatively recent literature exposing some of the confusions and overlooking of distinctions to which this debate has been prone.[12] So let me more closely articulate the particular view in this rough space that I will be developing here.

Subsistence in the superworld is a yes or no matter. But the kind of presence within phenomenal consciousness I have already noted can be a matter of degree. Hyperexperience is already in that way more complicated than the superworld. But this does not imply that the very existence of phenomenal consciousness in us, or anyone, is not a yes or no matter. I admit that I am somewhat sympathetic to the thought that entities outside of the class of unified neuralia we are discussing here, perhaps even all living things, have something analogous to the paradigm sort of phenomenal consciousness we are discussing, and maybe different degrees

---

[11] For the second idea, see Lycan (1996: 39) and Van Gulick (2007).

[12] Bayne, Hohwy, and Owen (2016); A. Y. Lee (2023).

of resemblance to that paradigm case, although this book develops no claims of that sort.[13] But the degrees of presence within phenomenal consciousness I am discussing here are not that sort of degreed *existence* of phenomenal consciousness. There are properties like mass that come in degrees while yet there is a yes or no fact whether something has mass, and I think that the existence of at least hyperexperience is like that.[14] And in addition, I think that the structure of hyperexperience is more complex than merely one axis of degreed presence can capture. That is why I have mentioned different forms as well as degrees of phenomenal consciousness. I think there are, even in we humans, different sorts of phenomenal consciousness that cannot be organized along a single scale from least to greatest.[15] This makes the structure of hyperexperience quite complex. I also think that there is reason to believe that there are forms of phenomenal consciousness in us that are even outside the complex structure of hyperexperience, that aren't rooted in fully present and paradigm phenomenal consciousness in the way that hyperexperience is, in which hyperexperience preserves one structural analogy with the superworld. Still, in neuralia like us, they are parts of our phenomenal consciousness as a yes or no matter.

The existence of degrees of consciousness is sometimes motivated as a prediction of theories of phenomenal consciousness.[16] And it is true that this third, "form and degrees" aspect of my proposal in this chapter is suggested by the whole nervous system model. But my central argument for it will involve introspective exploration of the modal structure of our phenomenal consciousness, buttressed by replies to various objections that the whole nervous system model will naturally occasion. This positive argument will slowly unfold throughout the chapter. But let me begin with a case I take from contemporary enactivists, who claim that representational states play a much smaller role in cognition than has been generally assumed.

According to Alva Noë, perception is constituted by sensorimotor knowledge, which is a kind of know-how,[17] analogous to Heideggerean skillful coping.[18] Whatever the general accuracy of this proposal, which,

---

[13] So phenomenal consciousness may have evolved gradually, contrary to Tye (2021).

[14] This is a theme of A. Y. Lee (2023).

[15] This is a theme of Bayne, Hohwy, and Owen (2016).

[16] For discussion, see Tye (2021: chapter 1).

[17] Noë (2006); O'Regan (2011) is a different version. For a treatment of what is for these views a difficult distinction between neural dominance and deference in determining qualia, see Hurley (2007).

[18] Noë (2006: 66).

I think, is dubious in some respects, he has some suggestive cases. One is your experience of a whole tomato when you only see its face. Noë holds that this is the possession of sensorimotor knowledge, which is know-how. "Our perceptual sense of the tomato's wholeness – of its volume and backside, and so forth – consists in our implicit understanding (our expectation) that movements of our body to the left or right, say, will bring further bits of the tomato into view."[19] One might object that the room next door is not perceptually present despite our possession of such an implicit understanding relevant to it, so that case must be different from this one. But Noë says that movements "or changes in the room next door will not provoke (visual) sensory change," and your relation to the room is less movement-dependent, since blinking "affects your relation to the tomato in front of you, but not to the room next door."[20] So that is, he thinks, the relevant difference. We will shortly return to some details of this. But let me now just put the case in the idiom I favor.

It is plausible that we experience the tomato as having a backside and that this involves an expectation that it would look a certain way if we walked around and looked at it from the other side. But I think we sometimes, in such a case, have mere hyperexperience of the backside of the tomato. Still, there are significant differences between such a case and hyperexperience of the merely possible color of the chair that we considered earlier.

That color is fully hypothetical, but our expectation of the way the tomato looks on the other side is an expectation about reality. Its looking that way to us is, I think, presented in experience as closer to actuality than a merely possible color some chair might have. In phenomenal consciousness, it has a greater degree of presence, or perhaps another form of it, than some merely possible color. The case of the tomato is also different because it involves a linkage between a hypothetical action on our part and a hypothetical look that the tomato would consequently have for us, and that is a difference which is arguably present to us in hyperexperience. But not in the superworld apparently presented to us through that experience. In the superworld apparently present through experience, there is no involvement of our walking around in the very essence of the color of the back of a chair. Despite the fact that our experience of such a thing plausibly involves a novel sort of hyperexperience involving our action, in the superworld of our naive experience, the backside of the tomato does not crucially involve *us*. There were some philosophers in the sense data

---

[19] Ibid., 63. [20] Ibid., 65.

tradition who held that an objective tomato was constructed out of all the possible sensations of the tomato any creature might have. But if you literally think that objects are permanent possibilities of *sensation*, as Mill did, you have a different conception of the modal structure of reality than that presumed in the commonsense Aristotelian conception of the world that, properly understood, is the superworld. Rather, in the superworld, there is merely an objective tomato with a certain actual phenomenal color on all its sides. So the relation between our hypothetical experience of its backside upon our moving, which is reflected in our hyperexperience, and the nature of the world it suggests, is complicated and indirect.

Because of these differences between merely possible colors and backsides, while there is a way in which the tomato presents itself in our phenomenal consciousness as having a backside, that involves a more complicated modal structure in hyperexperience than a mere binary sort analogous to that between the actual and the merely subsisting. I won't attempt an exhaustive analysis of such complex structures here, though in our progress we will look at various cases that require various complications. Fortunately, our central focus is sensory qualia, which mostly do not involve these extra complications.

The three key notions of this chapter – hyperexperience, forms and degrees of phenomenal consciousness, and the identification of phenomenal consciousness with at least most of the nervous system – may well seem questionable. Indeed, they may seem absurd or crazy. But I hope to convince you by the end of this chapter that this conception isn't obviously false, that you don't *know* it is false, by answering natural objections, by adding refinements that those objections require, and by explaining more fully what I mean. One central positive argument that you should give this view a hearing will be our discussion of the most intuitively plausible cases for hyperexperience, which are other than the cases reflecting the modal structure of basic sensory properties like phenomenal color. But our discussion will also serve to plausibly account for those parts of the nervous system that you may initially think are not at all relevant to phenomenal consciousness, and so answer natural objections to the whole nervous system model and the hoary view about the ubiquity of phenomenal consciousness in mental life that it reflects. I will be attempting to suitably inflate the content of phenomenal consciousness to answer these objections. Chapters 4–8 will mostly be concerned to develop the aspects of hyperexperience that are closely tied to basic features of the superworld, since both those kinds of things reflect our sensory neurophysiology. And so some of the more ambitious things I say in this chapter are not strictly necessary for what follows. But I hope by the end of the book that you will not only get used to this overall conception but also see how various detailed considerations we will later engage do fit well with it. I will be asking you to think here in what is now a somewhat unfamiliar way, but we should expect that a solution to the hard problem of consciousness will require that.

3

Before we get into the difficulties and details, it may be useful to note a few relatively close comparisons.

This view is not, of course, representationalism. It does not deliver qualia by appeal to the informational content of mental states, and causally based information plays no role in my proposal, not merely at the first-order but at any higher-order. That may encourage worry that it cannot explain our ability to introspect our qualia, which may seem to require higher-order thought or perception or some global workspace, and so some form of higher-order representation. But the general view of introspection I favor can be adopted from first-order representationalism, and skirts this concern: We, no doubt falsely, experience color qualia as out there on the objects. This is at base because various forms of sensory stimulation causally determine what colors we see and what color qualia we experience. While what is out in the world does not help constitute our sensory consciousness, on my view, still the state of our sensory receptor cells does, and the states of sensory receptor cells are closely controlled by forms of sensory stimulation that we receive from our environment. We talk about and act on our sensory stimulations, and also the environmental objects which in turn cause them, largely in the guise of the phenomenal consciousness which all that causes in us. On such a conception, all we need to do to introspect our qualia and sensory consciousness in the relevant sense is to learn a new way to talk about what we naively see as out in the world. We merely repress corrections we are otherwise inclined to make when reporting and acting on what we see, corrections that lead us to ignore the experiential fact that a tree looks smaller as it recedes. We merely learn to talk about the uncorrected looks of things, just as they appear to us in experience. How can mere possibilities, the kinds relevant to mere hyperexperience, be introspected? Because, as we have already noted and will see in greater detail in what follows, even alternative real possibilities for some actual neural state can be dynamically relevant to action and also help constitute the phenomenal consciousness that lies behind action. How can things in our neurophysiology that are other than sensory stimulations, which are also relevant to qualia on my proposal, be introspected? Partly because they are relevant to what actions are ultimately taken on the basis of sensory stimulations, for instance to what is said. Still, there is no doubt that we can recognize, react to, and act upon similarities and differences in our sensory qualia in various ways, so it is likely that whatever constitutes at least some sensory qualia does not reach all the way to particular actions without some intervening comparisons of qualia first occurring. This is a complexity to which we must be sensitive as we proceed.

While it is borrowed from representationalism, in our argumentative context this general view of introspection suggests an inversion of orthodox representationalism. It suggests that while sensory representations do not constitute qualia, perhaps qualia constitute sensory representations. That, in turn, suggests another close comparison.

Phenomenal intentionality theory, or PIT, is the view that phenomenal consciousness is the root of intentionality or aboutness, that all intentionality arises from phenomenal consciousness.[21] One big difference with my proposal is that I am not here developing any theory of intentionality at all. But still, while I am not proposing an account of sensory content per se, I do think that the account of sensory qualia developed here could plausibly be deployed in that role. And some forms of phenomenal intentionality theory develop, as I do, an extended notion of the content of phenomenal consciousness. There are various hard cases for PIT, for instance standing beliefs and abstract thoughts, which are also hard cases for me. What are called "inflationary" forms of PIT deploy controversially rich forms of phenomenal consciousness to help with such cases, while in "derivationist" forms the hard cases are supposed to derive, perhaps through inference or dispositions, from the easy ones.[22] One form of PIT might both inflate and derive. In this chapter, I will be developing an account of hard cases that can provide resources to both derivationist and inflationist PIT. So while I am not proposing any theory of intentionality here, I am a fellow traveler, and one way to think of what I am doing is as developing an account of how the physical can deliver the resources required by fans of phenomenal intentionality without what seems their customary, and I think unfortunate, recourse to dualism.

---

[21] Bourget (2010); Farkas (2008); Horgan and Tienson (2002: 520–533); Mendelovici (2018); Pitt (2004).

[22] Mendelovici and Bourget (2020).

Since the qualia I am attempting to deliver may conceivably play such representational roles without themselves being constituted by causally based representation, my view is also in some ways a close contrast with enactivism, which attempts to root many intuitive mental states not in any fundamental sort of representation but rather in know-how.

We just considered in a friendly way Noë's enactivist account of the presence of the backside of a tomato in perception.[23] I already described that case in a somewhat different way, but to bring out a greater difference between my view and enactivism, contrast the back of the tomato with another case. There are tomatoes in the refrigerator in the kitchen, and I have the sensorimotor know-how to engage with them, but they are not plausibly present in my perception.[24] What's the difference with the back of the tomato? Noë suggests that we have a relation to the back of a visible tomato that involves both movement-dependence, in the sense that the tomato's look will change if we ourselves move, and also object-dependence, in the sense that its look will change if it moves around, and that both sorts of dependence are reflected in relevant sensorimotor know-how.[25] My relation to the tomatoes in the refrigerator is not object-dependent in the same way; I wouldn't see analogous movement. So, he claims, perceptual presence requires both these sorts of know-how. But now consider a third case, due to Amy Kind.[26] There is some soda inside an open can you see on your desk, and you know that. The backside of the can is present in your experience in the sense we are discussing, but the soda is not, although you possess both sorts of know-how regarding that liquid. Noë can't deliver that difference. So what is it? Kind suggests that there is imagination of the back-side of the can. What do I think? My proposal about hyperexperience of the back-side of our tomato is in some sense midway between Noë's and Kind's. We will return shortly to various differences between imagination and the flavor of hyperexperience that is relevant here, including phenomenological ones, but one salient difference between them is that an imagination of the back of the can will be from some particular rough angle and distance, whereas hyperexperience encompasses all the possibilities. I think Kind has proposed a form of presence of the back-side that it is too robust, and would be unable to plausibly encompass the full range of alternative views of the can required.

---

[23] Noë (2005).

[24] We will return, skeptically, to accounts of hyperexperience suggesting a view not unlike Noë's.

[25] Noë (2006: 64–65). His further proposal that the difference in question is a matter of degree is not implausible, and would complicate what I say later about belief.

[26] Kind (2018).

But, on the other hand, she seems to me right that a condition more robust than that Noë proposes is required, and delivers a more robust sort of phenomenology than he countenances. I think the perceptual presence of the back-side requires neural conditions with a kind of modal complexity beyond those his enactivist proposal demands, and underwrites a kind of presence in our qualia, in our phenomenal consciousness, which he implicitly denies. The modal complexity of phenomenal consciousness is, I think, greater than that which enactivism can deliver, at least in its current forms.

4

We turn now to pressing objections to and some consequent needed refinements of my initial proposal. Our remaining discussion in this chapter will be organized around problems for the view that neurophysiology in general constitutes phenomenal consciousness. But intuitive support for and exploration of the modal structural complexities of hyperexperience will be a crucial component of my replies to some of these objections, while other replies will involve forms of phenomenal consciousness that are beyond hyperexperience and in that way help to delimit its range. So our negative discussion will also have a positive upshot. We will begin with objections rooted in commonsense psychology, and later consider others rooted in scientific psychology and cognitive science, and even in pseudoscientific psychology. Here's the first. I claim that even the peripheral nervous system is part of what constitutes our phenomenal consciousness.[27] This is because sensory receptor cells and motor activity are important resources that help deliver intuitive differences and similarity among some sensory qualia, in ways we will begin to trace in Chapters 4 and 5. But fans of the significance of the brain in particular may object by appeal to certain commonsense cases.[28]

In dreams, we apparently have phenomenal experience without the stimulation of the peripheral nervous system.[29] For instance, we can dream of a vivid red rose without stimulation of the sensory receptors involved in seeing red roses, and even apparently without activity in the primary visual cortex. What's more, the ordinary routes from sensory stimulation to action are broken in dreamers, in fact actively suppressed, even though on my view such routes are part of what constitutes phenomenal consciousness.

There are those who would question whether dreams should be considered phenomenally conscious. But they intuitively involve qualia. And

---

[27] See also Aranyosi (2013).

[28] Elpidorou (2015); Hill (2015).

[29] Hobson (1988).

dreams are clearly rooted in the brain, and so if they aren't phenomenally conscious would still constitute an objection to my proposal. So I cannot rely on this point in reply. Also, dreams are not wholly disanalogous to other phenomena we will treat later in our discussion of objections based in psychology and cognitive science, as involving distinctive forms of phenomenal consciousness that are not parts of hyperexperience. But my main reply to this worry about dreams is this: There is no very strong reason to believe that the qualia of dreams are exactly like the qualia of waking perception, that they are not dimmer or more abstract. So activity in some more limited portion of our neurophysiology than that which is active in analogous perception may well suffice to constitute the relevant dream qualia in accord with my proposal. In addition, my proposal allows neural structure that is intuitively inactive, which constitutes merely a kind of fixed modal background for intuitive activity, to be relevant to the nature of qualia. Sensory receptors may play that role even when they are not stimulated. What's more, dreams may have some effect on future action, or at the very least on possible action that would occur should the suppression of action cease, and so might be constituted by neurophysiological modal structure of the sort that would under various alternative circumstances mediate between sensory stimulations and motor actions in relevant ways, in hypothetical ways that yet may be relevant to actual current hyperexperience. Finally, dreams may sometimes be plausibly held to involve lesser degrees of mere hyperexperience than those at its core. It is expectable on my proposal that dreams do not involve exactly the paradigmatic forms of phenomenal consciousness, the core of hyperexperience in its ordinary qualitative form with a full degree of presence. But such complexities imply nothing that commonsense introspection forbids.

Any remotely adequate discussion of anesthesia involves considerable complexities,[30] to some of which we will return during our discussion of objections rooted in cognitive science, but it requires treatment in some ways similar to dreams. Another somewhat analogous problem case is general paralysis, which is especially difficult for me in instances when motor output is part of what constitutes details of phenomenal consciousness.

But, I reply, even in that case, there may be inhibited or blocked dispositions to motor action that play some consciousness-constituting role. And aspects of sensory qualia due to motor action may be affected without eliminating such qualia altogether. And there are internal effects

[30] Cole-Adams (2018).

such as hormone release beyond motor activity that are not undercut by even general paralysis. All this probably also happens during dreaming.

Another case that might be thought to show that peripheral nervous system activity cannot be part of what constitutes qualia is the case of phantom limb pain. Those who have lost a limb may feel pain as if it were in such a limb, even when the limb and its pain receptors and nerves no longer exist. And some experience pain in a limb, say a leg, when it is not stimulated but there is spinal cord damage, indeed even when the spinal cord is severed.[31] But, I reply, in fact the most significant aspect of what constitutes pain qualia is I think not peripheral in the specific way under pressure from such objections, but involves motor and emotional tendencies. Pain qualia are an unusual case in that respect, to which we will return.

The phenomenological aspects of the more or less commonsense cases we have just considered are difficult to understand in detail on any view. But it may well seem that common sense harbors much more certain counterexamples to my proposal, to which we now turn.

5

The principal source of commonsense resistance to my proposal is that there seems to be much more to our psychologies than our phenomenal consciousness, and those additional things seem to be constituted by our neurophysiology. Consider for instance beliefs. One key objection is this: Not all our neurophysiology constitutes phenomenal consciousness because some of it constitutes our beliefs, and many beliefs do not affect phenomenal consciousness.

My response to this objection will be a brief sketch of a conception that of course demands more elaboration than I can provide here. It has several components.

While our phenomenal consciousness is not due to information-based representation, still it does present us – in a perhaps misleading way – with certain perceived environmental objects at a time. We may have experience as of a red tomato on a green chair. And one part of my reply to this objection is to note that our phenomenal consciousness in such cases, the phenomenology of perceived objects at a moment, our hyperexperience, is richer than is often claimed. This has been a theme of a considerable amount of recent work, much of it stemming from Susanna Siegel.[32] And

---

[31] Berger and Gerstenbrand (1981).

[32] Siegel (2006, 2010).

there are famous historical antecedents for this conception in the phenomenological tradition, of which I developed a modal structural version in *Experience and Possibility*. It is at a minimum a live view, which I adopt here in such a modal structural guise. There are, I claim, certain sorts of beliefs, and hence their neural bases, which are reflected in phenomenal consciousness, contrary to what might first be supposed. This often involves mere hyperexperience, but of a sort that is especially close to full presence. Such a treatment also extends to certain sorts of know-how.

Remember what we said earlier about tomatoes. And when you see a chair as a chair, for instance, that is also plausibly reflected in various ways in your phenomenal consciousness. You are aware that it has a backside that you don't presently see, but which would look a certain way to you if you or the chair moved in certain ways;[33] the chair plausibly presents itself to you in experience as just such a complex three-dimensional thing. But what's more, the chair presents itself as capable of a certain sort of utilization, for instance sitting on.[34] If you are a driver, your car, indeed any car, presents itself as ready to be driven by a certain sort of engagement with the wheel and the accelerator. Your knowledge of how to drive and your beliefs about cars are no doubt rooted in your neurophysiology, but also plausibly reflected in such relatively subtle features of your phenomenal consciousness. A hammer presents itself in this way as ready to hand, especially if you know well how to use it.[35] And, if you know how to speak and read a language, words in that language present themselves in a richer way to you than if you don't, perhaps so rich as meaning what they do, or at least as bearing a component of that meaning. It is virtually impossible to see and experience the words of a language you know as merely sounds or marks.[36]

Since tools seem ready to hand in certain ways, and chairs to have unseen backsides, it is quite plausible to think that the neural basis of our possession of related skills and beliefs is a part of what constitutes us to have such experience, such phenomenal consciousness. In these cases, differences in beliefs and skills seem to directly introduce differences in that consciousness, particularly in mere hyperexperience.

This requires that hyperexperience have a more complex form than the simple binary structure of subsistence and existence of the superworld. We considered a few of the complexities recently. And some plausibly

---

[33] This may involve a certain degree of vagueness and a certain amount of error.

[34] Gibson (1979).

[35] Heidegger (1962).

[36] For discussion, see Azzouni (2013).

involve a kind of modal breadth in hyperexperience.[37] We often seem in some such sense aware of a variety of things we might bodily do, and a variety of things that would be experienced consequent to each action, perhaps with various degrees of uncertainty, which would in turn perhaps provide further opportunities for action. We are so aware that we *could* walk over and grab the tomato and take a bite, or flee it, and if it tastes one way or another, act variously. Although there is a hypothetical aspect of such things, they are closer to actuality than any fully hypothetical red of the green chair, and characteristically have a greater presence, or, perhaps better put, a different form of presence, in phenomenal consciousness.

But while there are these reflections of know-how and belief within phenomenal consciousness, such beliefs and know-how need not be identified with such phenomenal consciousness. If beliefs are part of the brain, they need not be spatial parts, but merely modal parts, capable of in some sense independent variation. And their neural basis might overlap, and in fact given the way they introduce differences into phenomenal experience plausibly does overlap, with the neural basis of that consciousness, without those two things being identical.

Still, that is only one set of cases. There is a second set of beliefs that plausibly require a different sort of response. I think there are three types of responses arguably available in these new cases. One of them assimilates these cases to the first set, and the other two don't. I leave it up to you to pick your favorite response of these three, or mix them in the way I prefer and will endorse, although I do think the first response is usually implausible. Here are the new cases. It is quite unintuitive to claim that you are presently experiencing many things about which you surely have beliefs. For instance, you believe there is an oval office in a big white house not far from a large obelisk. You believe there is a very large museum entered through a glass pyramid. These beliefs can't obviously be treated in the same way as those that inflect your experience of things right in front of you now like the tomato on the chair, the cases we've already considered. Still, it may be that some of these new type cases, or conceivably even all of them, can in fact be properly assimilated to those we already considered.[38] That's the first possible reply. If you glimpse a corner of your neighbor's couch through the open door, and have some familiarity with their living room, that may engage in some way phenomenal consciousness of the whole room, or even perhaps living room furniture in general, in

---

[37] Trestman (2014).

[38] Greater vagueness and error would be involved in this application of our former strategy.

something like the way your sense that a seen chair has an unseen back does. And maybe you even so experience that room as part of a house in a town on a planet. There are many routes out from where you are that would take you through the world, perhaps organized in your expectation by your approaching and overtaking and then leaving behind various objects you might see on your progress through that world. And just possibly that involves an especially ethereal form of mere hyperexperience even of the far side of the moon. You also have beliefs about the future and the past. But maybe routes to where you are now or forward from here suggest there is hyperexperience of all that as well, right here and now. We'll come back to cases like that.

But it is certainly reasonable to object that this first strategy cannot plausibly be pushed all the way to a proper treatment of all cases of belief and know-how of the sorts now under consideration. Indeed, we were earlier skeptical about such an account of even tomatoes in the refrigerator. Our current experience of a coffee cup doesn't plausibly take us all the way to the moons of Jupiter. And so it seems necessary to refine my initial crude proposal about the neural basis of phenomenal consciousness to handle these cases of beliefs about things beyond current experience. That's the second possible response. And that's what I think. Refining revision of the whole nervous system model is required.

On this second conception, which I prefer for many of the cases now in question, these beliefs and the relevant know-how would only affect your phenomenal consciousness *if* you were presented with other environmental objects than you are currently presented with. On this conception, it is not that all of the neurophysiological basis of all your beliefs and know-how relevant to things you think actual, all the modal structure that would intervene between any relevantly possible stimulation of sensory receptor cells and motor action for you at a time, constitutes your phenomenal consciousness at that time. Rather, that neural basis *would* help constitute other sorts of phenomenal consciousness *if* you were otherwise stimulated. It potentiates such alternative phenomenal consciousness in that sense; it puts you in a particular state of readiness for it if only you receive the requisite stimulation. This is a kind of modal halo or penumbra for your current phenomenal consciousness. It isn't part of your hyperexperience, on this conception, though it is still modally related to it.

There is also a third possible response to this set of cases, which significantly varies the second idea. Perhaps the neural basis of your beliefs about things beyond experience constitutes not what would make it true that you would have a certain sort of conscious experience under conditions of hypothetical stimulation, as the second response suggests, but rather constitutes your own sense of that, how you take it that your phenomenal consciousness would be under various hypothetical conditions of stimulation. And maybe it is also important that you believe that one set of possible stimulations rather than another is available here in reality. The second response would instead have to fold that somehow into different reactions to possible stimulations, say presenting them in phenomenal consciousness as fanciful or realistic.

This third response may seem to require a more significant refinement of my claims about the neural basis of phenomenal consciousness, since the neural basis of these beliefs would represent our consciousness rather than be entwined with it. But I don't think this is so. I think that the two sorts of beliefs deployed in the third response can be delivered by the resources already made available by the second response. To explain this point reasonably well will first require consideration of other problem cases. But we will get there. We can begin with this point.

If how you take it that your experience would go if you were presented with certain objects is very much awry from how it would go, it isn't at all clear that you have the beliefs about the world that you might hence take yourself to have, rather than fantasies about what your beliefs might be. If you claim, even to yourself, to believe that cars are not shadowed by evil monsters, but whenever actually presented with a car, you expect such a monster to lurk on the other side, it isn't obvious that you have the beliefs you claim. And so the third response seems to have at most limited application. It may refine our beliefs about the concrete world, by virtue of our possession of, for instance, relevant linguistically mediated beliefs about that, things we say to ourselves, but it seems implausible that it could generally trump the resources deployed in the first or second responses.We will also shortly see that linguistically mediated beliefs are in fact reflected in phenomenal consciousness, and similar things can be said about other ways in which you might possess the beliefs required by the third response. Linguistically mediated beliefs, and their analogues, are other problematic cases we need to consider in their own right, and we will also need to do that before we can properly see how the resources of the second response in front of us can serve to deliver the third, to the degree that is required. But they can.

It might be that some cases of belief regarding things outside of immediate experience require one of these three treatments and others require another, and that some require mixtures. That is in fact what I think and will presume in what follows, though I lean quite heavily on the second response for the cases we have been discussing. But I also think it does deliver the resources required for the third response when that is necessary, as I will shortly better explain.

6

Before we turn to the important case of linguistically mediated belief, there is a difficulty to confront related to another set of problem cases. One worry about the proposal on the table is that it doesn't allow a proper difference between things believed and merely imagined. The faint experience of the far side of a chair and of a merely imaginary flying submarine may seem quite the same phenomenologically. It may also be important to consider imagination as an important resource to deliver certain sorts of belief, not just as a significant contrast.

Since there is a hypothetical aspect of the cases we have been considering already, and since I have claimed that there is even mere hyperexperience of the fully hypothetical red of an orange chair, it is natural to wonder if hyperexperience might be even further bloated than is required by its deployment for all cases in the first possible response to the problem of beliefs about unexperienced actual objects and aspects of such objects. We can imagine many things that we do not believe exist, even beyond those directly entwined in being with those we believe to exist. So perhaps there is a kind of delicate phenomenal consciousness of such mere possibilities when you actively imagine them, or even a very delicate hyperexperience of all of them even when you don't. Maybe even all of one's possible sensory experiences of a moment, or all of one's possible actions, are present in mere hyperexperience in some way.[39]

If we take the phenomenally expansive first strategy for these cases of imagination as well as sometimes for beliefs about what isn't immediately present in experience, if we consider even your imagination of a merely hypothetical apple to be in some sense present in your mere hyperexperience along with the unseen far side of the moon, it is natural to think that the internal structure of hyperexperience will need to be complex in ways that capture this difference.[40] But we are probably already committed to that anyway. And even if we adopt the second or third strategies for beliefs,

---

[39] Another possibility is to expand hyperexperience to include fantasy objects, but not generally take the first route for handling belief about objects outside of experience.

[40] This isn't obvious, since it might be that mere imaginings are reflected in phenomenal consciousness, while their difference from belief-relevant imagery is not.

it is plausible that imaginings are at least sometimes present in phenomenal consciousness and indeed hyperexperience, and present in a somewhat distinctive way.

The similarity of structure of hyperexperience and the superworld is limited on any of these models we are exploring. In any case, the structure of hyperexperience is more complex than that of the superworld. The superworld involves but one sort of subsistence that is not existence. But the way in which a merely possible red thing is present in phenomenal consciousness of an actual orange thing is plausibly different from the way in which the look of the hidden backside of a chair is present in the experience of a chair, even on the most restricted conception of hyperexperience we've been considering. If we extend the model of hyperexperience to encompass belief in the whole world, or things imagined, then many more complications ensue. And there are also further complications involving the phenomenological experience of absences, like Pierre's not being in the café when he is expected. The structure of hyperexperience is certainly complex, and may be very complex indeed. But there remains an important structural analogy between the superworld and hyperexperience, whichever detailed model of hyperexperience we adopt. A merely subsisting red is the same type of thing as the actually existing orange with which its being is entwined in the superworld; it is a concrete color. And it seems plausible that the hidden look of the back of the chair, or the imagined look of a merely hypothetical chair, is also in that sense the same type of thing as the present look of the front of the chair. At least much of our mere hyperexperience that is not fully present experience seems to be of the same general sort as fully present experience, in the sense that its content is a similar kind of content, even if the form or degree of presence of that content is of a different sort. This suggests that the way in which the chair presents itself in mere hyperexperience is related in some way to alternative modes of fully present sensory consciousness of the chair, which you don't but might have. It isn't something nonsensory or very abstract instead, although it is clear that it can be somewhat vague. And this feature of hyperexperience might be generally related to another capacity that we obviously have. It is clear that most humans sometimes have a kind of quasi-sensory imagery of fantasies, which we might call "primary imagination." You can, at will, hear in your mind's ear what you wished you'd said, or see in your mind's eye what your childhood bedroom would have looked like with that paint job you never got.[41] That is not unlike

---

[41] There is variation in individual capacities for primary imagination.

sensory experience, but in most cases, to put it crudely, fainter. The neural basis of this imaginative faculty is plausibly not completely spatially divorced from the neural basis of sensory consciousness. There is neurophysiological evidence that, for instance, visual imagination involves activity in some of the very same neural regions that are active during visual perception. And of course on the type of view I'm proposing, even the neural basis for the possibility of activity constituting visual perception, for instance inactive sensory receptor cells in the visual system, may be relevant to the constitution of just such a state of imagining.

The existence of the primary imagination, when conjoined with the fact that aspects of our phenomenal consciousness relevant to belief that are not fully present are yet not different in kind, suggests that the two phenomena are linked. That seems right, but I do not think they are the same, not at least in all cases. There are cases of mere hyperexperience, those which we have mostly been discussing, that do not usually involve the primary imagination, I think. And a part of my reply to the worry now under consideration is that some of the cases from the two classes are phenomenologically distinct in some way, allowing a difference between belief and primary imagining, although I do grant that it is likely that they involve somewhat similar neural resources.

However we handle in detail these complexities about the imagination, it is clear that the hidden backside of the chair is presented in phenomenal experience as actual in a way that a mere imagining is not. This fact is relevant to understanding all the ways we have considered for dealing with beliefs about objects outside of current experience. And it is a crucial part of the distinction between imagination and belief we must also deliver.

Fortunately, some further illumination about this phenomenon can come from a consideration of one of our remaining problematic cases of belief. We have quite complicated spatial and temporal conceptions of reality. Even if we can deliver beliefs about objects outside of experience in one of the ways we previously considered, our complex spatial and temporal beliefs weave together these relatively atomic beliefs in complicated ways we should more explicitly consider. There is a question how my proposal can deliver all that.

We've been focused mostly on the experience of a moment, and we've considered three models for delivering the atoms of complex spatial and temporal conceptions, some of the objects required to constitute, say, a city over time. But when we come to a discussion of time consciousness in Chapter 7, and treat it in accord with the general proposal about phenomenal consciousness I have made, it will deliver a crucial resource for unity of these atoms. It is reasonably clear that you can have a kind of fixed intention for a motor trajectory through the world that it will take a while to execute, and a kind of expectation of the experience you will in fact have along that route. And this plausibly inflects even your phenomenal consciousness at this time. You feel where things are probably going, and so the look of the backside of the chair may be very salient in this way. You also have some understanding of how experience would unfold if you took another possible trajectory that you don't intend to take. Hyperexperience of the backside of the chair is knitted together in these ways with fully present current experience, while a mere imagining generally is not. That is one key difference between those cases. And indeed this is one grounds someone might cite to support the claim that the conception of hyperexperience as including the entire world according to one's conception is not completely implausible. Our conception of what we would experience if we traveled throughout the rest of the entire world in various ways is knitted together with present experience in ways closely analogous to the way our hyperexperience of the back of the chair is knitted in. But even if we don't take that radical step, even if we don't consider phenomenal consciousness so bloated as to encompass all belief, it is clear that there are constraints of consistency between our phenomenal consciousness of a moment and of earlier and later moments. There are expectations that may be met or violated, reflected in the hyperexperience of a moment. We don't simply see the front sides of objects, or even consider their unseen sides in isolation. We have a conception of a complicated spatial world of objects that would look various ways if we moved and they moved in this way and that, and at least a bit of that structure is often present in hyperexperience at a time, stretching out beyond what immediately confronts us. This can underwrite surprise if things don't turn out that way.

It is a quite characteristic feature of human cognition that we consider fanciful possibilities of action or of circumstance that we do not undergo or undertake. What is most interesting about a person is sometimes what they consider doing but don't do, or what they imagine, or what they don't even consider or imagine, so that not their history but these penumbral possibilities seem most explanatory and interesting about them. But that is a different case than hyperexperience of the occluded or distant. A part of the difference no doubt has to do with the will, with the way in which fantasy can range quite freely and intentionally, but there does seem also to be a phenomenological difference of another sort. Mere possibilities of actual experience are highly relevant to human cognition even of what is taken to be actual. But the complex spatial contents of human belief weave together possible and actual experiences, and the complex contents of human actions actual and possible, in a very distinctive way. As Kant hinted, our experience of a more or less objective space and time of independent objects involves a complex intermingling of what we do immediately experience and what we expect we would experience, or perhaps what we would not be surprised if we did experience, as things unfold in the world or if we were to move this way or that. We have somewhat analogous capacities to entertain how things seem to other people and so what it is to be another person. It is true that much of this capacity to conceive the past, the perceptions of others, or even the occluded backsides of present things, is closely related to primary imagination, our capacity for some sort of attenuated, freely willed experience as of things we are not in fact experiencing. But in the cases that now concern us, cases of belief, it is entwined with actual experience in important ways. And what I am suggesting is that at least some of that extra content is also often part of hyperexperience in a way distinctive from the purely imaginary. Some of our expectation about future experiences, even those hypothetical on certain merely possible actions, inflects even current phenomenal consciousness in a characteristic way, and even when we aren't actively involved in imagining such things. This is one distinction between the products of the primary imagination and the features of hyperexperience relevant to belief.

The way in which one momentary experience must lock appropriately with others, due to this structure in hyperexperience, helps distinguish imagination from belief. We have also noted some other key phenomenological differences and the relevance of the will.[42] But it remains important that what is merely imaginary doesn't fit into the run of experience in the same intimate way as the back-side of a tomato. And so this is a resource that is helpful in delivering belief about things far outside of immediate experience within the strictures of the second response. It provides some of the necessary linkage between the atoms of hyperexperience when belief is involved. But we also need a few more resources to that end as well, which are themselves also needed to handle other problematic cases, to which we now turn.

7

Another commonsense objection to my proposal rooted in belief brings us back to one resource I have already mentioned: Much of what we believe

---

[42] These factors encompass all three types of resources deployed respectively by the impoverishment view, the will-dependence view, and the nonexistence view distinguished by Kind (2020).

involves what we say, and we talk about a lot of things that are not obviously part of any concrete conception of the world that can be delivered in the ways suggested already. There are disagreements about justice and the continuum hypothesis. And so, the objection goes, the neural basis of such beliefs may well be quite distinct from that involved in constituting phenomenal consciousness.

But, I reply, the words we say, even to ourselves, and read and write and hear, are concrete things of which there can be phenomenal consciousness in ways we have already noted. So even they can be folded into this rough conception. Meaningful words in a language we know present themselves to us in a special way, even sometimes, like tools, as things that can be used in characteristic ways. And we can be potentiated to so respond to words we aren't actually experiencing, in the manner of the second response. Indeed, while much linguistically mediated belief is plausibly unspoken, we may be potentiated to respond to various linguistic and other cues present in stimulation with at least inner speech of a relevant sort that would be present in our phenomenal consciousness, so the resources of the second response have a very wide reach here. There are various indirect ways in which the meaning of even very abstract words can be rooted in the concrete sorts of conditions we have already considered. For instance, concrete scenarios can play something like the role of possible worlds in delivering familiar conceptions of the truth conditions of sentences.[43] These scenarios may enter phenomenal consciousness or its penumbra in ways we've already discussed. And it is also not completely implausible that they are sometimes locally entwined in hyperexperience with relevant experienced words and sentences, though in very complex ways. I've said something about all this elsewhere.[44] It is also the case that various sorts of linguistic and other concrete resources can be linked together in various sorts of layers, for instance with words related to vague concrete scenarios that themselves can be refined linguistically, and so on.[45] So we have a complex set of resources available to deliver linguistically mediated belief.

There is clearly much more that needs to be said about this conception than can be said here. It might be best for now just to focus on the nonlinguistic capacities of us admittedly very chatty humans, and introduce a related refinement of my proposal, which would rule the neural basis of linguistic capacities out of our consideration. Feel free to do so. But it seems to me we cannot plausibly do that very cleanly, and that

---

[43] Chalmers (2012); Lewis (1975).

[44] Mendola (1997, 2008).

[45] Something like this is suggested in the PIT tradition by Bourget and Mendelovici.

suggests that it shouldn't be done. There are a number of reasons for this. What we believe in linguistically mediated ways sometimes affects our phenomenal consciousness, as for instance we come to experience a tree in a certain specific way once we have learned to what named species it belongs.[46] And words do present themselves as meaningful in experience when we know the language, even arguably as sometimes bearing certain specific meanings. And there is the linkage between word meaning and concrete conditions noted, which at least occasionally does seem to be present in hyperexperience. And sometimes word-mediated belief about the world outside of experience does inflect the other resources we are using to deliver the third response to our initial worry about beliefs, for instance the manner of the interlocking of momentary experiences due to their temporal structure.

This isn't to say that word mediation is always crucial in founding belief about the world far outside of immediate experience. We have noted various other useful resources. And one thing that also plausibly sometimes plays such a role is a memorial capacity related to primary imagination, a kind of gross spatial map, say of some particular town, a scale model into which we can zoom at will in imagination. The primary imagination can be deployed in a useful way as long as different such imaginings can be certified as playing particular disciplined roles, as a rough map of an actual town rather than a mere fantasy. Indeed, even words we experience must be taken in the right way. This may seem yet another difficulty for my proposal. But we are about to consider a resource that helps in this regard, in the guise of a response to yet another commonsense objection.[47]

There is at least one more commonsense case closely related to belief that we should consider explicitly: memory.[48] Clearly this is also part of any adequate treatment of belief about objects outside of experience as well. Parts of our brain are plausibly the basis of memory, which is not intuitively phenomenally conscious, and this too seems an objection to my proposal.

But I reply in this way: It is customary to distinguish between declarative and nondeclarative long-term memory. Nondeclarative memory is memory to which in some sense we have no conscious access, which includes such things as knowledge of how to ride a bike and to read and

---

[46] Siegel (2010).

[47] For more on the ways thought may be present in consciousness, see Bayne (2020).

[48] For an introduction to memory, see Gazzaniga, Ivry, and Mangun (2002: 301–350). Working memory is frequently held to have a significant role in consciousness, as in the global workspace theory, and short-term sensory memory presents no evident objection to my view of qualia, and will be a focus again in our discussion of temporal consciousness.

how to perceive things accurately in familiar environments, and also learned associations between stimuli and motor responses. But these sorts of things are already in fact included in the potentiation by know-how that we have already discussed in the case of tools.

On the other hand, we are generally supposed to have conscious access to declarative memory. Tulving distinguished between episodic declarative memory, which is memory of historical episodes in our lives, so to speak from the inside, and semantic memory, which includes learned skills that are reportable, like how to tell time or count, and also reportable facts like who was Lincoln's vice president. Learned reportable skills are incorporated in my proposal in the same way as not easily reportable know-how. But does the neural basis of episodic declarative memory and semantic memory of facts constitute a telling objection? I don't think so, because such things are always capable of presence in phenomenal consciousness through presentation of possible stimuli, or quasi-sensory stimuli such as provided by the primary imagination. This is the second type of response we considered earlier regarding the case of belief. You might be asked, or you might ask yourself in your mind's ear, who was Lincoln's vice president or what your childhood bedroom looked like. Much of the memorial knowledge so queried seems to involve the capacity to recognize certain forms of primary imagination or other sorts of less-than-fully-present phenomenal consciousness, for instance as a view of your room.[49] And that plausibly sometimes inflects your phenomenal consciousness of it, in something like the manner that your knowledge of the species of a tree may do so, or at least in the guise of words you say to yourself about it.

Now notice how recognitional memory of this sort also provides a resource for the third response regarding beliefs about things outside immediate experience. And it also helps, along with the presence or absence of the will and other phenomenological differences, to distinguish fantasies and such beliefs. We can recognize some hyperexperience we might undergo as that of our childhood bedroom. Even dropped into such an experience, without regard to any connections to those experiences before and after it of the sort a plausible account of time consciousness requires, we can so distinguish between experiences relevant to the actual world beyond experience and those that aren't.

So here's another resource to help distinguish fantasy from belief. We also have the temporal interlocking between momentary experiences of what is actual available. And we have relevant word-mediated thought.

---

[49] Squire (1987).

And we have various plausible phenomenological differences and differences regarding the will. Furthermore, while there is little grounds for certainty on such complex matters so briefly treated, it would seem that all that is enough to deliver the third response regarding beliefs about what is not present within the strictures of the second, and so avoid the further refinement of the whole nervous system model that would otherwise be required.

So here's where we end up regarding belief and analogous phenomena: There is a complex neurophysiology with a complex modal structure that intervenes between possible sensory stimulations and possible motor actions. And the modal structure of that neurophysiology is the modal structure constituting, on the most expansive and implausible conception of hyperexperience suggested, your phenomenal consciousness at a time, and on less expansive conceptions of hyperexperience, which I endorse, constituting it and something additional but entwined besides, your immediate capacity for other forms of hyperexperience. On the less expansive conceptions, the relevant aspect of phenomenal consciousness is hyperexperience, but it is surrounded by the real and highly entwined possibilities of other sorts of hyperexperience, for instance those that you would not be surprised to have if you go one way or the other through the world, or those that in fact you would have whether you expect it or not. All that is complexly entwined at a minimum in the sense that it shares a complex neural base, the full neural structure that intervenes between sensory stimulation and action, and constitutes on the less expansive but more plausible conceptions not just hyperexperience but those various sorts of modal halos around it that are your contribution to the potentiation for those other experiences, though at the maximum, in the radical first response, all that complex entwining occurs within the modal structure of hyperexperience itself.

We can conclude that there may be some reason to refine my initial crude proposal about the neural basis of phenomenal consciousness, if one of the less expansive conceptions of hyperexperience under consideration are correct. And I think one key refinement is required, as I have indicated. But I claim that at least one of the less expansive forms of hyperexperience exists, though I believe not one that stretches to the far side of the moon, or even to where the Eagle landed.

Needless to say, providing adequate evidence for what I have said here about belief, memory, and language, and working out all the details properly, would be a very lengthy matter. But I hope this abbreviated discussion can serve to explain the refinements that my initial crude proposal about the neural basis of phenomenal consciousness plausibly require because of belief and related phenomena, and points to a way, for your further consideration, that certain worries about this view might be handled.

We will be focusing in what follows only on some of the most rudimentary sorts of sensory qualia in the vivid core of our phenomenal consciousness, the fully present core of hyperexperience, and what is bound up in the essences such qualia present to us. So, fortunately, these various complexities won't matter very much in most of the book. Still, we will be looking for resources to constitute basic sensory qualia throughout our neurophysiology. I hope this discussion has at least served to allay some resistance you might otherwise have had to that procedure, as well as to the rough viability of the whole nervous system model.

8

Much of the mind and the nervous system is concerned not with belief and perception but with motivation and action, with getting things done. And that too seems an objection to the whole nervous system model of phenomenal consciousness.

But this objection can be answered in ways analogous to our treatment of belief.

Consider first desire. There are at least four ways in which desire may be present in phenomenal consciousness that parallel four ways in which we have seen that cognition may also be present. If you really like apples or are hungry, that apple you see may appear in experience as especially needing to be eaten, or especially sweet and delicious upon being eaten. There are also unseen apples round the corner, which would appear sweet if seen. And there are mere fantasy apples of desire. And much desire is linguistically mediated.

But this does not exhaust the ways in which the brain plausibly roots motivation and action. Some cases plausibly involve a further extension of factors already in play.

Consider emotions. When someone is angry or fearful, objects immediately experienced or just around the corner or even in fantasy present themselves to that person in ways reflective of those emotions, in something like the way that delicious apples reflect desire. That hateful face presents itself as needing to be punched, and that dark entrance as scary and so to be shunned. There are also moods, when the whole world may, for instance, present itself as uniformly empty and worthless.[50] And of

---

[50] For an exploration of ways that emotions might be present in consciousness, see Deonna and Teroni (2020).

course there is a long tradition of thinking that desire and emotion present themselves somatically in some way, say in that longing felt in your chest or those butterflies in your stomach.

And there is also another cluster of relevant factors. There is also much of the nervous system that is devoted to the initiation, production, coordination, and control of action. For instance, the cerebellum plays a large role in this. But, as I've already hinted, action tendencies, for instance to flee or protect, are plausibly crucial to even certain sensory qualia, for instance those involved in physical pain. We have already noted that the presence to hand of a tool involves in phenomenal consciousness non-actual action in another relevant way, bound up in one's experience of the tool. And tools may be around the corner or merely imagined as well. Furthermore, we will see in Chapter 7 that some causal powers, those exhibited not only by the objects one immediately interacts with, like an incoming wave, but also by one's own actual countervailing motor activity, make a fundamental contribution to experience and hence phenomenal consciousness, with one's own powers even a constitutive element of that consciousness. And this treatment can be extended to include what is presumed to be around the corner or merely hypothetical. So this is yet another way in which motivational features of the nervous system, and in this case of the musculature that it controls, may be reflected in phenomenal consciousness.

Another aspect of action that is sometimes plausibly present in experience and hence phenomenal consciousness is planning, the orientation of things actually or hypothetically done in the world around a further goal. This goal will, when it is something experienced, be a kind of focus within phenomenal consciousness, both in itself and as a point towards which other things present in that consciousness seem oriented in something of the manner of tools. And such goals may also be available to phenomenal consciousness in the manner of things around the corner or mere fantasies.

So while there are huge complexities involved in motivation, which once again properly require much longer and more careful consideration than we can give them here, it at least initially seems that the aspects of the nervous system that are involved in motivation and action might be plausibly treated in accord with the whole nervous system model, in ways analogous to those that we have already sketched for belief.[51]

---

[51] For relevant neurophysiology, see Bear, Connors, and Paradiso (2016: 551–643).

9

Despite hoary conceptions that everything in the mind is conscious, contemporaries generally believe that there are many unconscious mental states. Such unconscious states are plausibly constituted by our neurophysiology. And so they provide reason to think that my proposal about the neural basis of phenomenal consciousness must be incorrect. There is also a somewhat hoary commonsense commitment, despite Descartes and Locke, to currently unconscious memories and beliefs and desires, and we have already considered those cases. But nowadays we expect there to be many more examples of the unconscious than that. Our currently expanded conception of unconscious mentality has two dominant roots, one in cognitive science and its forebears, and one in Freud and his ilk. We begin with the first.

There are various features of our neurophysiology that seem irrelevant to the constitution of our sensory qualia even when they are relevant to perception. Because of the dominance of information-processing conceptions, many think there are forms of sensory representation in the brain that are prior to any conscious mentation, which are closer to sensory stimulation in a hierarchy of forms of representation that culminate in conscious representation. For instance, if vision involves information-processing, it is natural to expect that representations of luminance edges present in visual stimuli, or other sorts of edge detectors, play an early and unconscious role in visual perception.[52]

But, I reply, we have left information-based stories about the constitution of conscious sensory qualia behind. Rather, we seek forms of neural modal structure that match that of phenomenal consciousness. And if there are edge-detectors in the brain that intervene between sensory stimulation and action, they can be part of the neurophysiological modal structure that constitutes visual phenomenal consciousness, regardless of whatever information about the environment they carry. If they were to correlate, through some odd accident, with waves on the sea or numbers in your checkbook, still they would provide the same relevant modal structure to the constitution of phenomenal consciousness. The fact is that we *consciously* experience edges, and there is no reason to believe that involves higher-order re-representation of lower-order representations or of their content, as information-based models presume. Rather, the differential states of the "edge-detectors" are more plausibly part of what helps

---

[52] Palmer (1999: 172–182). Canonical but superseded is Marr (1982).

constitute the overall modal structure that provides us with our conscious visual phenomenology.

A second objection: I have proposed that all of neurophysiology between stimulation and action together constitutes phenomenal consciousness. But sometimes distinct streams of sensory processing, one conscious and the other unconscious, are suggested. For instance, Milner and Goodale proposed that there is a ventral stream of visual processing relevant to identification of stimuli, while a dorsal stream is unconscious and involved in fine motor capacities.[53] But, I reply, my view is indifferent to the spatial location or segregation of causal influences in the brain that intervene between sensation and action, except insofar as they are relevant to the overall modal structure of the nervous system that does the linking. Separate streams of causal influence linking stimulation and action within one nervous system are as good as singular ones in the life of neuralia like us who, unlike jellyfish, have such unified nervous systems.

Although the notion of information-processing is somewhat less deeply imbedded in current studies of action than of sensation, it is also natural to think there are output-side analogues of the hierarchy of unconscious representations favored by current studies of perception, which specify the actions to be taken in ever greater detail as neural action proceeds outwards towards the motor periphery.[54] This may also seem to involve various forms of unconscious representation in the brain. But, I reply, this, like edge detection, is also just more grist for our mill. My proposal is that all the neurophysiology that controls motor action is also part of what constitutes the detail of our phenomenal consciousness.

But this response may seem more problematic than in the case of input-side sensory "processing." What our arms do may be dependent on various sorts of unconscious motor representations, while, unlike visual edge-detectors, those representations are not in a direct causal chain leading to what we say about our phenomenal consciousness, which does seem introspectible in that verbal way. Our mouths may not have access of a suitably direct sort to what our arms are up to. And this doesn't fit very well with the story about introspection I adopted from the representa tionalists.

The proper response to this worry requires an understanding of another set of pressing cases from cognitive science and psychology that may motivate the claim that there is unconscious perception. So let me discuss those cases before I return to this worry.

---

[53] Milner and Goodale (2008).

[54] Rosenbaum (1991: 79–118).

Psychology and cognitive science have identified various interesting and unexpected psychological phenomena for which our neurophysiology is clearly the basis, but which may seem to involve unconscious perception, contrary to my proposal. (1) There are cases of blindsight and unilateral neglect. Sometimes people with extensive neural damage do not acknowledge the presence of visual stimuli, while their behavior is still influenced by such stimuli as if they experience them.[55] (2) There are cases of inattentional blindness. Mack and Rock wrote a book-length presentation of experiments in which subjects were completely oblivious to a plainly visible stimulus when their attention was engaged elsewhere.[56] In one especially vivid study by Simons and Chabris, when subjects were asked to count how many times a basketball was passed by one of two teams in a video, 92 percent reported no awareness, when questioned afterwards, that a woman in a gorilla suit had cavorted very visibly among the players.[57] (3) There are cases of unconscious priming, in which it seems that some unreportable sensory content is interfering with some reportable content, or in other ways affecting what is done, as when the evil owner of the movie theater inserts consciously unnoticed frames of popcorn to get you to the concessions. (4) There are unattended but still attendable experiences that are psychologically salient in some ways, for instance what you saw when driving while yet focused with full attention only on your upcoming date, or the refrigerator humming that will only capture your attention when it stops. (5) There plausibly are forms of experience that do not involve the kinds of higher-order reportability or unified control through working memory characteristic of human consciousness but which persist in humans after such fancy capacities are destroyed. (6) There are relatively extreme states associated with post-comatose disorders, such as the so-called vegetative state and the minimally conscious state, and epileptic absence seizures.[58] Some of even these seem to involve somewhat degraded forms of perception, and even when they don't, they present another challenge to my proposal, which requires phenomenal consciousness in such cases when it may well seem absent. (7) And there are also the somewhat analogous cases of sleep and anesthesia on which we have already briefly touched, regarding which it is worth remembering that there are different forms of sleep such as REM sleep and different levels of sedation to which different analyses may be appropriate.

---

[55] Weiskrantz (1986).

[56] Mack and Rock (1998).

[57] Simons and Chabris (1999).

[58] Bayne, Hohwy, and Owen (2016).

I propose that we sort these cases into three classes. The first class seem to me to involve hyperexperience. For instance, while it is sometimes claimed that what is not attended is outside of phenomenal consciousness, still as you shift your attention across your visual field, especially without moving your eyes, the unattended bits seem to remain present to at least some degree, and part of hyperexperience, though not in its fully present core. Here's another case I think is like that: "In the Sperling phenomenon, subjects are shown . . . a grid of letters, say a four-by-three array. They say they see all or almost all of them but can report only three or four items. However, if subjects are cued after the stimulus is gone regarding which row to report, they can report three or four from any cued row."[59] Other cases that I think remain within hyperexperience involve not differences in presence from the core but rather unusual differences in content. For instance, a minimally conscious "patient might be aware of motion, but be unable to recognize the moving object as the kind of object it is."[60] I've also suggested that dreams can at least sometimes be treated this way.

But it is, I think, implausible to claim that all these cases reside within hyperexperience. I don't suggest that we incorporate blindsight into hyperexperience in some very faint way, even though blindsighters also have hyperexperience. Rather, I think a second class of cases, including blindsight, involve phenomenal consciousness of a second sort, or probably a second class of sorts, outside of hyperexperience. A third class of cases are those individuals in which there is some sort of phenomenal consciousness of roughly that second sort and no hyperexperience at all. Perhaps in some sleeping or sedated states even normal adults are sometimes like this. We aren't here considering phenomenal consciousness in nonhuman animals, but some also may fall within this third class. But perhaps it is worth noting for purposes of clarity that I think awake dogs and human babies have hyperexperience of at least rudimentary kinds, even though they can't report it as we can.

With that said, return now to our earlier cases. I think that unreportable motor-processing in humans is analogous to unconscious priming and blindsight. When such phenomena affect what happens between sensory stimulation and action, then they are phenomenally conscious, but often in different ways or to different degrees than the characteristic sort of human phenomenal consciousness found in hyperexperience. Otherwise, they involve potentiation for such phenomenal consciousness should other sensory stimulations occur, as allowed by our former refinement of my

---

[59] Block (2014: 445).

[60] Bayne, Hohwy, and Owen (2016).

crude proposal because of belief. This point about different degrees and forms of phenomenal consciousness is another important aspect of the whole nervous system model. Indeed, it is its third main feature.

I've given roughly phenomenological reasons to think some of the problem cases we face are within hyperexperience. But why think that blindsight and unconscious sensory priming are anything at all like such paradigm phenomenal consciousness? Why think they involve phenomenal consciousness of any sort? Because they involve the same sorts of neural structures intervening between sensory stimulation and action, and because they do so in a way that seems to require sensory qualia. The real choice is between saying that these cases involve different forms of phenomenal consciousness or saying that they involve unconscious sensory qualia, and that second thing seems like a greater terminological and conceptual strain. Notice that in, for instance, blindsight and priming, one would expect that the properties blindseen or interfering with conscious perception are of the same general type as those reportable and seen. So it seems that there are color qualia of the same general type as conscious color qualia present in cases of blindsight and priming. They affect action, but in a simpler, less determinate, more primitive, and less reportable way than the fully and paradigmatically conscious qualia which are our central concern. I think they involve a lesser or different sort of consciousness, but I think they do involve phenomenal consciousness. How could there instead be unconscious qualia? Still, if you prefer that other strange language, go ahead and put it that way. That doesn't change the reality.

We have already been dealing with what I've called degrees and forms of phenomenal consciousness, because of the complex forms of hyperexperience we discussed earlier. But the variations or degrees in phenomenal consciousness we are now considering seem often of a different sort than those. In some cases of this sort, there may be a large realm of real potential for fully realized qualia of paradigmatic sorts, which potentialities compete for full actualization and for effects on action. Or it may be that there are just lesser probabilities of effects on action. Or it may be another kind of effect on action than the ordinary sort that is involved.[61] Such forms of reduced consciousness may for instance only affect forced choices, which would be a less than dominant effect in another sense than that involved in those that lose the kind of competitions just noted.[62]

---

[61] Perhaps in humans there are such phenomenal qualia regarding things that do not enter hyperexperience, for instance pheromones.

[62] For this treatment of blindsight, see Gertler (2001); of neglect and blindsight, see Phillips (2016).

The coexistence of different forms of phenomenal consciousness within a person, outside of hyperexperience, generates natural questions about the unity of a person and of consciousness. And such questions are closely related to issues regarding other forms of unconscious mentation we should now consider in the guise of objections to my proposal.

10

Another root of the contemporary popularity of the unconscious is Freud. Whatever the questionable epistemic credentials of his various pseudo-scientific claims, some phenomena that were a focus of his concern do suggest objections to my proposal about the neural basis of phenomenal consciousness.

Dreams were one focus of his concern,[63] but we've already considered that case. He was also concerned with unconscious long-term memory.[64] And since he thought that, contrary to currently popular distinctions among forms of memory we recently considered, some memories were unconscious in a deep way, could not be brought to consciousness in any direct way, such possibilities might constitute further objections to my proposal. But he also thought that such repressed memories had effects on action. And so these cases are on his view like forms or degrees of consciousness we recently discussed, not unlike blindsight or unconscious priming.

But there is a third and perhaps deeper way in which Freud presents a challenge to my proposal. Freud believed in the importance of certain kinds of mental disunity in a person. For instance, it is almost as if the id, the ego, and the superego are supposed to be distinct agents, aware of different things, and fighting it out within a given human for control of action into the outside world.[65] The ego is phenomenally conscious, but at least not fully and directly aware of what the id is up to. There are even wilder analogous views. Pierre Janet suggested, though in a behaviorist-friendly way, that in various pathologies and even in normal humans, there are generally alternative centers of conscious agency, which are revealed as distinct only under damage or stress.[66] In other words, he suggested that there are generally multiple and distinct phenomenal consciousnesses within a given human. But my proposal implies that any phenomenal consciousness of a given human is unified simply in being constituted by

[63] Freud (1965). [64] For an introduction, see Freud (1963, 1989). [65] Freud (1962).
[66] Ellenberger (1970: 358–394); Sjovall (1967).

the single neurophysiology that intervenes between the sensory stimulation and motor output of a single human animal. So my proposal may seem wrong for reasons like these.

My reply will take a moment to state. These claims of Freud and Janet are of course at best wildly controversial. Still, my proposal does have resources to help deliver some of what they proclaimed, and also help with some obviously real cases.

The way in which we are aware of and can report and think about our own mental states, on the view I have proposed, certainly does not imply that we are infallible about our own mental states and psychologies. I can certainly be unaware of or misunderstand my tendencies to specific action under various forms of hypothetical sensory stimulation, especially because what I might pleasantly imagine I will do under certain hypothetical stimulations and what I will do under such stimulations may be quite different. In this way, we can have desires that would be reflected in phenomenal consciousness under hypothetical conditions without our being at all aware that we have them, and perhaps without being capable of understanding that we have them, except in times where they fully reveal themselves in action in some way that we cannot deny. And we are very good at denying what we are highly motivated to deny. Such resources are available to us given the refinements to the whole nervous system model we've added.

Still, Freud and Janet certainly believed in more disunity than that. They believed in competition and conflict among internal agents. But I have also already allowed that there can be different forms and degrees of phenomenal consciousness in a person, that affect action in different ways, and even some that may conflict. Blindsight and inattentional blindness and priming have effects on action which are not totally unlike the behavioral expression of an unconscious desire. In addition, there can be a conflict of behavioral tendencies and dispositions in someone. And we just noted that we can be much more aware of some of our behavioral dispositions rather than others.

But let's be more concrete, and consider some unusual but real cases. Some people, under extreme emotional pressure, apparently split into distinct personalities so that alternative characters with different names and characteristics, and with various abilities to remember and be aware of what other such characters have done and are doing, sequentially occupy a given body.[67]

---

[67] Wilkes (1988).

What I say about such cases, which is consistent with my proposal, is this: Disorders of memory are central to these cases, and our sense of personal identity and individuality over time is no doubt closely bound up with memory. So there may well be in such cases more than one person persisting over time while sharing a body and a brain. Still, that does not imply that there is also more than one phenomenal consciousness at a time. Perhaps there is at any given moment just one person-stage, with a single phenomenally unified consciousness, but the person-stages associated with one body over time may be variously collected by different shards of memory into more than one continuing person.

But maybe we think multiple personalities at least sometimes occupy one body at a single time. And there is also another set of unusual and extreme real cases that may help us to think better about these cases as well. There have been people who have had the main neural pathway between the two hemispheres of their cerebrum, the corpus callosum, severed, and under certain conditions they exhibit behavioral disunity that suggests to some that they have two distinct centers of consciousness, one in each hemisphere.[68] Still, in these cases the lower brain is intact, and shared by both cerebral hemispheres. There are forms of neural connection that persist between the two alleged centers of consciousness. Some on this basis suggest that there is but a single and fully unified phenomenal consciousness despite the commissurotomy.[69] That would be a treatment of the cases consistent with most of my proposal here. But the behavioral disunity that can be present in these cases is still in some intuitive tension with my claim that the neural sweep between all sensory stimulation and the entirety of motor action constitutes a single phenomenal consciousness. Others suggest that there is but partial unity in such cases.[70] That fits better with my proposal, in this way: Extant versions of this option suggest that each phenomenal consciousness shares some qualia with the other, so they overlap. This may be so. But recall also that there are different forms and degrees of phenomenal consciousness according to my proposal. For instance, sometimes subjects, when asked to report appearances in consciousness in three ways, by blinking, pushing a button, and talking, can do it in one way but not the other two, especially under severe time constraints. I believe that in these cases and also in split-brain cases, there may be a lesser degree or attenuated form of phenomenal consciousness and so unity that bridges the whole, which might be considered another

---

[68] Gazzaniga (1970). This is controversial. For an alternative, see Marks (1981).

[69] Bayne (2008).

[70] Lockwood (1989); Schechter (2014). For other cases of partial unity, see Hurley (1998).

form of partial unity than qualia overlap. After all, under many conditions, two split hemispheres work cooperatively and smoothly in the world, and there is no behavioral conflict and disunity between them. It might also be that in something like this way, different personalities can occupy one body at a time, unified together in some ways but not others.

Perhaps some animals have sensory organs and motor capacities that are connected but also entirely disconnected from the remainder of the sensory organs and motor capacities and neurophysiology of the animal.[71] They have a split neurophysiology, and so they would, on my proposal, lack a single unified consciousness. And it may be there are even some rare human cases like that. But ordinary humans can do more or less anything they can do on the basis of any stimulation they might undergo. They can hop on one foot when the green light comes on. And when I've been speaking of a neurophysiology that intervenes between a range of sensory stimulations and actions, I've been presuming an intertwined and single neurophysiological structure of the sort normal humans have and which allows that, because we are concentrating on the normal human case. Still, such a structure can exhibit the forms of disunity we have been considering here in the ways just sketched.

II

There are various ways in which phenomenal consciousness is unified at a time in even the most paradigmatic and straightforward cases. One's visual experience of a cup is unified in some sense with one's visual experience of a nearby pitcher. Both experiences are part of one experience, and there is a spatial relation between the two objects present in the experience. In addition, distinct sensory aspects, intuitively involving different senses, are unified in experience at a time. The cup may present in experience both a seen color and a felt texture. And one can hear something while seeing something else.

At least in the class of cases on which we are focusing in what follows, the sensory qualia at the core of hyperexperience, it is not implausible to claim that all bits of the relevant experience of a moment present spatial relations within that experience. Some tactile feels and some looks are presented as at the same place. So we will spend some time in what follows on sensory consciousness of spatial relations.

---

[71] The ventral and dorsal streams discussed earlier are not like this.

There are complexities relevant to even this basic spatial unity of consciousness. For instance, it is clear that the space present immediately in vision is more detailed and definite than that involved, for instance, in smell or even experience of sound, and it may have details that infect our spatial experience of tastes and sounds when they are identified with visual locations. And there are other relevant cross-modal phenomena. You hear the ventriloquist's voice coming out of the dummy's moving mouth, and the voice on the speakers coming out of the mouth seen on the screen. And there are other relevant complications. Remember Molyneux's question.[72] He asked Locke whether a blind person who had sight restored could recognize which felt shapes correspond to which visual shapes, and the answer to that question seems relevant to whether felt shape and seen shape are qualitatively the same. Some recent empirical evidence suggests that three-dimensional shapes are not immediately matched in this way by those whose sight is restored,[73] although the relevant capacity is quickly learned. This suggests that the answer to the question is no. However, on the other hand, there is the fact that the capacity is quickly learned, and the empirical studies in question don't involve specifically the kind of two-dimensional case on which Locke and Molyneux fastened.[74] And it is also relevant that the blind draw in much the way that the sighted do, which suggests some shared spatial structure.[75] Nevertheless, it does seem that the feel of a shape and the sight of that shape differ phenomenologically, and, as I said, different senses involve different levels of intrinsic spatial detail.

Still, despite these complexities, all that spatial structure can be fused together at least in the core of hyperexperience, in a spatially unified sensory experience of a moment. There are many issues about the unity of our consciousness, even the unity of our hyperexperience at a moment. But the most basic question is how various different experiences due to different senses constitute one experience of a moment in a given person, distinct from that of any other person or animal, but internally unified at least to the degree that we humans can compare and distinguish all the various types of elements present in our immediate experience – colors, shapes, and sounds – and also at least to the degree that we, like animals generally, can act as unified agents on the basis of their whole spatially complex experience at once.

There are two aspects or versions of this basic question. The first version is phenomenological. How does this unity appear within experience?

---

[72] Van Cleve (2014).
[73] Sinha, Wulff, and Held (2014).
[74] Van Cleve (2014).
[75] Kennedy (1993, 1997).

The second is ontological. How does this unity in experience, in fact, arise? My answer to the first version, in the cases of sensory hyperexperience that concern us, is that it appears as a kind of spatial unity. My general answer to the second is that an entire animal like us at a time has but one phenomenal consciousness, constituted by the modal structure of the single, intertwined neurophysiology that intervenes between its sensory stimulation and its motor action, which allows these various sorts of spatial fusing among different senses, although of course we have yet to see how spatial experience arises. The unity of consciousness is the unity of the animal's neurophysiology. Of course, in animals and sometimes in humans, this may involve a lesser sort of phenomenal consciousness, such as we recently considered. And we noted in our earlier discussion of introspection that relevant unity may need sometimes to be accomplished neurophysiologically before bodily action occurs.[76]

12

One last set of objections to my proposal focuses on the existence of specific brain regions thought to perform very different functions than intervention between sensory stimulation and action. The peripheral nervous system and the spinal cord are for the most part implicated in the link between sensation and action, but there are exceptions, and there are parts of the brain, especially the brainstem, which seem to have very different roles. But my response to these worries is simply to allocate the problematic regions among the various responses already in play.

There are neural structures involved in controlling visceral automatic activities, like breathing and swallowing and the beating of the heart. But as I said earlier, we have left the autonomic nervous system to the side, except as it involves hormone secretions that do intervene between sensation and action. Other brain areas are involved in the control of facial expressions, or modulate arousal and sleep, or smooth and coordinate action. But these are relevant to action and phenomenal consciousness in ways we discussed during our treatments of belief and desire. It is also plausible that many of the problematic brain structures are relevant to reduced forms of phenomenal consciousness such as we recently discussed.

---

[76] For many different questions beyond the two I have asked, see G. Lee (2014a). For general discussions of the unity of consciousness, see Bayne (2010); Dainton (2000); Masrour (2020); and Tye (2003).

13

The main proposal I have made in this chapter – (i) the whole nervous system model, conjoined with the existence of (ii) complex hyperexperience and (iii) various forms and degrees of phenomenal consciousness, but qualified and refined by appeal to (iv) the relevance of the state of our nervous system to what we would experience under merely hypothetical conditions of stimulation – is not only speculative and controversial, but also, while already complicated, still radically undeveloped. Some further hints about its general nature, further grounds for its truth, and further responses to possible objections will slowly unfold in what follows. For instance, Chapter 6 will consider potentially problematic phenomena involved in neural control of eye saccades. And some of the more elaborate modal structural complexities of the brain hinted at in Chapter 2 – for instance when they are relevant to neural systems involved in arousal, attention, and action – are plausibly reflected in some acute modal complexities of hyperexperience we have noted in this chapter. But our primary focus from here on will be narrower.

Our concern will be to elaborate this conception only for some of the most rudimentary sorts of sensory qualia, beginning with color in Chapter 4. We will see that sensory qualia require modal structural resources up and down any neural hierarchy that bridges sensory stimulation and action, so that it is unlikely that any small spatial component of the brain, or very specific kind of neural resource, is the highly localized seat of phenomenal consciousness. Also, we will see in more detail how some of the simpler modal details of hyperexperience are plausibly part of the phenomenal consciousness involved in sensory experience. While this will be in accord with the general conception sketched in this chapter, for the most part this second positive feature of my overall proposal will be merely a suggestive setting for that more focused treatment of basic sensory qualia. This is not a way of taking it back. I think such a view is required if we are to deliver phenomenal consciousness in all intuitive cases, for all neuralia with unified nervous systems. But it does suggest ways in which much of the rest of my proposal can be made independent of the central claim of this chapter, if you feel the need to make it so.

CHAPTER 4

The Modal Structure of Color Qualia

This chapter will sketch the third and final feature of my general proposal regarding human phenomenal consciousness of sensory qualia, by attention to the relatively well-understood case of color qualia. This conception will later be refined and developed by attention to other cases. It is an understanding of the detailed way in which differential features of our qualia are rooted in neurophysiological differences, because of modal structure that is shared between the content of our phenomenal consciousness and our sensory neurophysiology.

There are philosophical disputes about the nature and reality of color that would take us far afield. So I will not be concerned with the nature of color itself, but merely with color qualia, with the what-it's-like of phenomenal consciousness involving color experience, with phenomenal color just as it appears in our experience. It is largely uncontroversial among vision scientists that there are ways in which color qualia reflect the neurophysiology of color vision. But I will be developing a certain way of understanding these well-established but still philosophically surprising facts, which involves modal structure. I'll call this "modal structuralism" in accord with my usage elsewhere, despite the importance of modal structure in the other two main features of my proposal as well.

I

Let me begin by reminding you of various ways in which our color qualia do not reflect features of the objective world itself, adding some relevant details to what I've said before, and also indicating general features of our visual neurophysiology that they reflect.[1]

When you perceive the surface colors of an opaque object, light is reflected by the object, and then stimulates sensory receptor cells in your

---

[1] Standard philosophical introductions include Hardin (1993) and Hilbert (1987).

https://doi.org/10.1017/9781009514774.004

eye. Despite the role of reflected light in this, there is a certain degree of constancy in your color qualia regarding a given object even under somewhat different conditions of illumination, under different colors and intensities of light. There is an objective property of such an object that corresponds approximately well with the color it is seen to have, and with color qualia, despite changes in illumination. That is its surface spectral reflectance. This is the proportion of incident light that the object is disposed to reflect at each wavelength in the visible spectrum of electromagnetic radiation.[2] For different objects, such a disposition may be constituted by different chemical and physical features. However, objects with different physical surfaces but the same surface spectral reflectance will seem to be the same color under a full range of normal illumination.

Nevertheless, there are a number of problems with identifying surface spectral reflectances with color qualia, with phenomenal color.

First, some things that do not have opaque surfaces are seen to have colors, including lights and translucent things. These colors are not surface spectral reflectances, but yet may present to us quite similar color qualia. Second, there are metamers, which have very different sorts of surface spectral reflectances, but yet present the same color to humans who are not color-blind under normal conditions of illumination, although they can be seen to be different in odd conditions of illumination.[3] There are hence differences in spectral reflectance that are not visible to normal humans under normal conditions. Perhaps we should say that these are imperceptible differences in color, though in any case they do not normally involve different color qualia. A third but related point is that there is nothing about surface spectral reflectances in the visible range that is objectively special. Infrared and ultraviolet radiation exist, but are not seen by humans. These present no color qualia despite being reflected from objects.

Metamerism occurs because humans are, normally, trichromats. This means that any arbitrarily colored light stimulus can be matched in their experience by a mixture of three colored lights. This fact was exploited by three-color printing and early color television. Trichromacy suggests that there are three kinds of sensory receptor cells in our eyes, sensitive to different wavelengths of light, and, as we will see, this is true. That would also help explain the visible colors at least closely analogous to surface colors that translucent objects and lights are seen to have, since they

---

[2] For this view of color, see Byrne and Hilbert (2003) and Hilbert (1987).

[3] This is oversimplified. For details, see Hardin (1993) and Hilbert (1987).

involve radiation in similar frequency ranges, and it would explain the limits of humanly visible radiation.

It is possible to group together, by brute force, surface spectral reflectances and other physical conditions that are seen to have the same phenomenal color, but the principle of collection will involve the peculiarities of human trichromacy. They will be grouped together as one color, or at least as presenting one kind of color quale, because of the way they interact with our form of trichromatic vision, not because of objective physical similarities. Our color qualia seem dependent in this way on things about us and not about the objects out in the world that we see. Our peculiarities also seem to be the reason that there are no infrared or ultraviolet color qualia.

There are other objections to identifying color qualia with surface spectral reflectances. Some especially telling ones are rooted in forms of similarity and structure among color qualia that involve features of human visual neurophysiology other than trichromacy. The fourth relevant difficulty for that identity is this.

Consider the child's color wheel, rounding from red through orange to yellow to green to blue and then through blueish purple and reddish purple and back to red. Now consider it as the equator of a color globe with pure white and pure black at the poles, and hues and shades filled in over the surface of the globe and within it, so that closer points have more similar colors.[4] Consider the globe as mapping the similarity of color qualia for colors of opaque surfaces. Now focus on the circle of qualia similarity that rounds the equator. Some violet hues and some reddish hues on that equator appear quite similar, although they correspond to very different wavelengths of light, at opposite ends of the visible spectrum of electromagnetic radiation, and so to very different spectral reflectances. It seems essential to phenomenal colors, colors just as we experience them, to color qualia, that they have these relations of similarity to other phenomenal colors. But in some cases this corresponds to no objective similarity of surface spectral reflectances or light.A fifth difficulty is that there are the unique hues, for instance the red that appears neither blueish nor yellowish. There are also unique yellow, blue, and green. The structure of unique hues also seems essential to phenomenal color, to color qualia. And there is nothing in objective surface spectral reflectances or light waves that accords the unique hues any special status.[5]

---

[4] For further complexities regarding "color spaces," see Kuehni (2010).

[5] But for a suggestive proposal regarding the objective difference between blue and yellow, see Broackes (2011).

A sixth difficulty is that there are also other forms of structure among the color qualia, whereby some specific hue shades appearing to be determinates of the "determinable" property green for instance, while quite similar hue shades do not, are over the line. This structure also corresponds to nothing objective in light or surface spectral reflectances.

There are still more phenomena that reveal a difference between color qualia and objective physical properties. Color experiences change with mood, so that depression and sleep deprivation make the world look grayer.[6] There are many chromatic illusions. And the color that a surface seems to have depends on surrounding colors. But focus on color similarities and the structure of unique hues: the fourth, fifth, and sixth objections just noted.

Like metamers, these phenomena plausibly reflect peculiarities of human color vision, but of another sort. Despite the fact that we are trichromats, there also seems to be a kind of opponency in our color vision, not only of white and black but also of the unique hues on the equator of the color globe. Unique red and green are phenomenologically opposed and maximally dissimilar, as are unique yellow and unique blue, much as pure black and white. There are four unique hues on the equator, instead of the three you might expect on the basis of our trichromacy. And intervening hues on the equator partake in some degree of essential resemblance to the unique hues they are between. There are yellowish greens and greenish yellows, and indeed some yellowish greens are more yellowish than other yellowish greens. But also yellowish greens and greenish yellows remain respectively greens and yellows, while yet some yellowish greens and greenish yellows are quite similar. It seems that our vision is specially responsive to reds and greens, and also to blues and yellows, as opposing and determinable types of colors, and that there is an extreme of opposition between unique red and green on one hand and unique blue and yellow on the other.

This structure of phenomenal color suggests that there is, despite our trichomacy, a sensory system that can be tipped in its response towards green or red, and another that can be tipped towards blue or yellow. That form of oppositional structure, called opponent-processing, would explain the phenomena just noted in this way: We don't see any reddish green, because those are opposing colors of one system. We see something as unique red when the first system is tipped towards red and the second system is in a neutral position. We see something as orange when the first

---

[6] Prinz (2012: 16).

system is tipped towards red and the second towards yellow. And the reddish purples and the blueish purples both involve tipping towards red and towards blue, which explains their similarity despite the fact that they correspond to very different wavelengths and surface spectral reflectances, but there is a difference in their degree of tipping in those directions.

As we will see, this opponency in our color vision plausibly reflects the structure of our neurophysiology of color vision, just as other features of our color qualia rooted in trichromacy reflect our three kinds of chromatic sensory receptor cells. But we don't even need to turn to neurophysiology to find further evidence for this opponent organization.

Consider color blindness. There are those who are completely red-green color-blind, and unable to distinguish reds from greens. There are those who are completely blue-yellow color-blind, and unable to distinguish blues from yellows. It is natural to suppose that the first have a disorder of their red-green opponency system, while the second have a disorder of their blue-yellow opponency system.

The color-blind provide many grounds to think that differences in color qualia are not due to objective differences out in the world. The color-blind plausibly have different qualia than those with normal color vision. It is hard to believe that someone could have vivid red and green qualia and not notice the difference. The difference between normal and color-blind qualia must plausibly be constituted by internal differences between the normal and the color-blind. Put a color-blind person in exactly the same external environment as a normal percipient, and they will have a somewhat different experience of the world. So what constitutes the difference must be inside them. Even if the color-blind differ in some of their relations to their environment, they so differ because of what is inside of them. And the relevant differences seem to be not only internal but also current. People can become color-blind by a hit on the head. If current internal neural differences can deliver the difference in color qualia between the normal and the color-blind, as they apparently must, then we have some antecedent reason to think that current internal neural differences could at least in principal deliver all differences in qualia.

Because color-blind individuals are incapable of discriminations that the normal can make, it may seem that while their differences from the normal are internal, still what dominantly constitutes our qualia is not. It may seem that internal differences merely screen off the color-blind from the world, like a cataract. But that is not a plausible understanding of the differences in question. Some individuals are anomalous trichromats, who have color experience that is quite clearly different from the normal in a manner other than subtraction. And they can make discriminations among environmental stimuli that no color-normal human can make.[7] There are differences among people's color experience that are not merely intuitive omissions from the normal case.

There are other cases of this same general sort. These are individual differences in color vision and color qualia among the color normal, who aren't color-blind. For instance, there is considerable variation among color normals regarding exactly what wavelengths of light are seen as unique yellow, despite the fact that a given individual will identify the same wavelength as unique yellow over a long temporal gap. This plausibly reflects individual differences in the way color opponency systems are organized, for instance at exactly what wavelength of light the neutral point between red and green occurs.

It is quite clear that no fully objective physical similarities of external objects, similarities that could be discovered other than by appeal to how they affect humans with our peculiar color vision, can constitute the similarities of our color qualia. And it is also clear that they do reflect features of our visual neurophysiology.

2

Nevertheless, there are those who think that there are colors that exist out in reality more or less just as we humans experience them to be, but that they are not constituted by physical structures such as surface spectral reflectances.[8] They think that in addition to physical properties like surface spectral reflectances, opaque objects have irreducible phenomenal colors, which are much as Aristotle thought colors to be.

We have noted that this would require locally entwined subsistence and existence in the manner of the superworld.[9] A specific hue of red in our experience presents itself as a universal, capable of entire existence in more than one place at the same time, though as an immanent universal, which can only exist when instanced. Somehow the instantation of that determinate universal would also necessitate the instantiation of the determinable universal red. And it would also be entwined in being with merely subsisting phenomenal colors with which it bears necessary relations of similarity and difference. The actual existence of such a modally complex beast is hard to accept. Others may prefer another sort of metaphysical

---

[7] Kaiser and Boynton (1996: 444).

[8] Such "color primitivists" include J. Campbell (1997).

[9] For extended elaboration and defense of this view, see Mendola (2021: chapter 7).

model of this situation, deploying uninstantiated and hence transcendental universals rather than subsistence, but that too is not easy to swallow. And, because of the way our color vision works, with our experience of colors mediated by light and the way it affects our sensory receptors, it is hard to see what any such irreducible phenomenal color would have to do with our vision in any case. And there are even greater phenomenological differences in color experience across different species than I have noted among humans – differences that cannot plausibly be explained by different primitive Aristotelian colors of this sort. Pit vipers have infra-red skin sensors that involve no lens-like focusing or color opponent-processing.[10] Even if that case is too distant from human vision to constitute color experience of any sort, still there are obvious examples of color experience that involve different qualia than human color experience. Goldfish, bees, and many birds are tetrachromats. Their color vision involves four-different visual pigments.[11] At least in pigeons, it appears that there is a corresponding fourth opponent-processing channel, beyond the dark-light, red-green, and blue-yellow channels of humans.[12] So there are apparently more forms of difference and similarity constitutive of pigeon color space than our own. There are also differences in the ranges of wavelengths that different species see. Salmon see into the infrared range, bees into the ultraviolet range, and pigeons somewhat into both.[13] And bees are not only tetrachromats who see into the ultraviolet range, but also have a sensitivity to light polarization that humans lack. They also do not have single-lens eyes like our own.[14] Even if we are prepared to presume that only humans see the true irreducible colors of things in a godlike way, such a view is still implausible, because of the differences in color experience among color-normal and variously color-blind humans. There can be no specific irreducible colors of this irreducible sort on the surfaces of external objects that capture all the different color qualia of different humans, not even of all color-normal humans.

Still, it might seem, as some dualists contend, that we each have our own private irreducible colors of this sort, in some private space. Dualists cannot plausibly deny that there are the sorts of opponency and trichromacy we have noted, and they cannot plausibly deny that human color vision involves the interaction of reflected light with our neurophysiology of vision. But they may deny that that is sufficient to constitute our experience of color qualia.

---

[10] Matthen (2005: 166 fn), which references P. Bradley.

[11] Allen (2009); Matthen (2005: 163).

[12] Allen (2009: 200).

[13] Ibid., 199.

[14] Matthen (2005: 185).

We will return in Chapter 8 to dualist arguments of this sort. There is in fact no pressing reason to agree with the dualists about this, given the physicalist story we are in the process of developing. But there is a specific case often deployed in support of dualism that is revealing in the context of our immediate concerns.

Consider the case of your imaginary color invert: They are a neurophysiological and behavioral duplicate of you. They call the same things red and blue, and their neural processing is the same when they look at an apple or the sky as yours is. But they are hypothesized to have inverted color qualia. Along some axis across the child's color wheel, their qualia are flipped over relative to yours. The invert has your neurophysiology, but not your qualia. It may seem that such a color invert is possible, since I just described them. It may seem that what can be conceived to be distinct, like your physical structure and your qualia, is distinct. And this may seem to provide good grounds for dualism. As I said, I don't agree. But this case is still a useful foil for us. It can help explain and motivate the relatively specific class of views of color qualia among which my physicalist proposal is located. It points towards the specific kind of explanation of our color qualia by our neurophysiology that I favor.

3

The view of color qualia I favor is among a class of views pioneered among philosophers by C. L. Hardin and Austen Clark, on the basis of a more general idea proposed by Bernard Harrison, specifically, as we will shortly see, as a response to qualia inversion cases.[15] The class of views in question holds that the structure of phenomenal color experience reflects the structure of our neurophysiology of color vision in a way that forbids inversions. There are also important historical precedents for this view in psychology. For instance, Müller's doctrine of "specific nerve energies" for specific senses was extended by Young and Helmholtz into an account of color qualia differences rooted in trichromacy.[16]

In my version of this idea, the modal structure of the apparent colors that our experience seems to reveal is explained by the actual modal structure of our neurophysiology of color vision. We have noted what I take to be essential relations of similarity and difference among the color qualia of color-normal humans. I think that such colors are often mere

---

[15] Clark (1993, 2000); Hardin (1993: 113–186); Harrison (1973: 102–114, 1986).

[16] Boring (1942: 70–74).

illusions of our experience, which exist nowhere, not even on colorful sense data in the mind. I think there are merely experiences of such colors, just as there are beliefs about unicorns when there are in reality no unicorns. But if such nonveridical aspects of our experience were true, if there were such Aristotelian colors out in reality spread around on ordinary objects, then there would be out in the world colors that would have these relations essentially. They would have such a modal structure. And these nonveridical aspects of our color experience are constituted and explained, I believe, by corresponding but real modal structure of the neurophysiology that constitutes that phenomenal experience. We have seen that the modal structure of the colors we experience does not match very well the modal structure of any objective properties of objects out in the world. That's the negative part of my claim. But it does reflect the modal structure of our own neurophysiology of color vision. That's the positive part. It is because of this match that understanding the detailed metaphysical nature of the colors we seem to experience helps us to understand not merely that they aren't out there in reality, but also how we experience such properties nonetheless. In the realm of our experience, in the superworld that would obtain if our experience were fully veridical, scarlet is a kind of red, which is necessarily related to green, and this in turn requires that green is at least a subsisting color. A scarlet wagon would encode in the essence of its color the possibility of a green one. Who could believe it? But this apparent modal structure of experienced colors reflects the very real modal structure of the neurophysiology of our color vision, a neurophysiology that assures that any creature who can experience such a red can also experience green, that any creature who does experience scarlet also experiences red. The particular opponent-organization of color-normal human vision is what does the work, in a way I will shortly explain. The modal structure of the properties we seem to see is a crucial part of the explanation of how creatures with a visual neurophysiology like ours would see such properties, even though nothing in reality has them. The phenomenal colors we seem to see, with their various complex relations of similarity and difference spanning modal structures that include subsisting mere possibilities, are all rooted somewhat obliquely in the various possible states that our visual system can in reality occupy.

But that these are possible states of our visual system in the relevant sense does not require the kinds of fancy metaphysics deployed in the superworld, with subsisting mere possibilities. There are actual but somewhat general neurophysiological states of actual humans who are also potentially in a variety of more specific actual states. Your standing neurophysiology of vision is such that, depending on what neurons are active in which ways, you will see either green or red. That neurological background is actual, but grounds various real possibilities of activation. And if you are in the state of seeing red specifically, then more specific neural conditions within that background are actual, while yet under other specific neural conditions consistent with that neurological background, you would have seen green. There is a sense in which seeing red is fully actual, while seeing green has a kind of lesser degree of actuality, but still has some, reflected in your actual neurophysiology. But no metaphysically suspect lesser degrees of being or transcendental universals are required. You aren't seeing green, but that real possibility is assured by your actual background neurophysiological state, and actual differences of activation consistent with that background determine which real possibilities are actualized. And so the modal structure present in the red we experience, which involves green of a lesser degree of being, is in some sense reflected in our actual neurophysiology.

That's the general idea behind what I am proposing in this chapter. Let me stress a contrast with one popular view. There is interest in cognitive science in second-order isomorphism,[17] which is to say "second-order correspondence linking the relations between different external stimuli and the relations between their internal representations."[18] But, on my view, while color qualia similarity matters, it is not in a way that involves informational representation of real world color similarities of the sort second-order isomorphism requires. The correspondence between neurophysiology and qualia I am suggesting does not involve information-based representation.

One way to understand the details of my proposal is to explore a nearer contrast, with the much more closely related accounts proposed by Hardin, Harrison, and Clark. At the first level of detail, this will involve consideration of the color invert.

It is sometimes thought that we can have no way of ruling out the possibility that your physical twin yet has an inverted color experience,[19] and this is sometimes thought to be a good argument for dualism. Harrison's general idea in reply was to exploit asymmetries in color experience to show that such inversions can be ruled out. I think that this was a good idea, but that there are problems with Harrison's execution of the strategy. His roughly Wittgensteinian proposal was focused on color

---

[17] Edelman (1998); Shepard and Chipman (1970).

[18] Gordon and Hayward (1973).

[19] Chalmers (1996: 99–101).

terminology, and I think he presumed too close a connection between the detailed nature of our color terminology and the structure of our sensory experience and qualia. And he wasn't much concerned with the physical basis of color vision. But Hardin's and Clark's variants of Harrison's strategy *are* concerned to link our sensory experience and our neurophysiology, as I think we should.

Hardin's was developed first. He focuses on the possibility of color inverts who yet share our neurophysiology of vision. He attempts to project the structure of the normal range of possible color experiences onto the normal range of possible neural states that plausibly underlie them, on the basis of structural analogies. And he argues that we have more antecedent reason to expect, on the basis of the structure of human color experience and of our neurophysiology, that someone who has normal visual neural systems would have normal rather than inverted color experience. On such a view, unlike Harrison's, color similarities we recognize are not due to linguistic phenomena; rather our language of color naming follows the similarities present in color experience because of our relatively fixed neurophysiology. One way to think of his proposal is as providing what was, at the time it was proposed, a plausible and generally acceptable neurophysiological story about trichromacy and opponent-processing, that shows how those phenomena are rooted in neural details, and then applying that story to the philosophical case of qualia inverts.

Hardin's specific proposal was this.[20] First, there is a neurophysiological basis for trichromacy. Human color vision is mediated by three different sensory receptor cells, or "cones," which are maximally responsive to three different wavelengths of light.[21] They are called the *S*, *M*, and *L* cones, for shortwave, middlewave, and longwave. Each has a spectral absorption curve that gives the probability that a photon of a given wavelength will be absorbed by the cone and hence affect its firing, although the form of that firing is independent of what wavelength of light is in fact received. The curves in question are roughly bell curves, and while the curves of the different types of cells have different peaks, the curves also overlap somewhat, with the *M* and *L* curves overlapping especially closely. Second is a

---

[20] The general idea is due to Hering, and Jameson and Hurvich. Hardin's view has taken different detailed forms over time. Hardin (1993: 113–186) is the original formulation on which I focus, but shies away from robust claims about the constitution of the asymmetries among color primaries by physical differences. Hardin (1997) is more confident in that way, but also introduces complications about the neural basis of color experience to which we will return.

[21] For a vivid history of color science focused on trichromacy, see Mollon (2012).

more controversial, complicated, and in some ways speculative proposal about the neural basis of opponency.[22] Assume that the firing outputs of the cones can be added and subtracted when they determine the firing of other neurons on which they synapse, and that neurons have a baseline firing rate 0 that can be increased or decreased. Then L + M could be a neural code, in neurons receiving input from *L* and *M* cones, that when above the base rate codes for white and when below codes for black. It is an achromatic signal. There are also two chromatic codes on this proposal. L – M might code for red when above the base rate and green when below it, in neurons whose firing is causally dependent on that of the *L* and *M* cones. And (L + M) – S might code for yellow when above the base rate and blue when below.[23] This sort of neural code might constitute the opponent organization of our color vision, and so explain the structure of unique hues, color similarities, and determinable color borders. And it was at the time Hardin made this proposal not unreasonable to think that such a neural code might be present in the normal human nervous system, from the retina to the visual cortex.[24]

How would this help deal with the qualia invert?

First, it delivers the phenomenon of unique hues. Any plausible assignment of color qualia to our neurophysiology would need to assign qualia corresponding to unique hues to neural states in which one chromatic signal is neutral and the other is tipped, as we recently discussed. Second, any plausible assignment would assign unique red and unique green to different sides of the same chromatic signal. Likewise for unique blue and yellow. This greatly restricts plausible inversions of color experience.

But Hardin also tries to squeeze more out of this strategy. For instance, he suggests that since yellow seems phenomenally hotter than blue and red hotter than green, perhaps this reflects states of positive versus negative activity in appropriate opponent processes.[25] This underlies that detail of his own choice of the particular code just suggested.

---

[22] Hardin (1993: 35).

[23] There are controversies about the details of such schemes, and also reasons for and against various schemes, including this one in particular. See Hardin (1993: 1–58). We will return to this.

[24] Although even then Hardin noted complications.

[25] This idea is presented as speculative in Hardin (1993), and there is evidence for opponent-process cells that are active in both directions. But Hardin (1997) cites an unpublished paper by B. Katra and B. H. Wooten in support of the claim that summed average opponent-process cancellation data reflects the intuitive asymmetry between hot and cool colors.

He also points out that there are more discriminable greens than reds, and more discriminable blues than yellows.[26] So perhaps we are down to just one possible inversion, in which unique red is inverted with unique yellow and unique blue with unique green. And there are other exploitable asymmetries in color space and in our neurophysiology. Yellows are characteristically less saturated than reds. And indeed color-sensitive cells in the cortex apparently are more sensitive to relatively lighter yellows and darker reds.[27] For that matter, oranges and yellows shade into what seems the distinct color of brown, which isn't phenomenologically merely a blackened yellow or orange in quite the way that dark blues and greens are merely blackened blues and greens. And this may reflect the same neurophysiological asymmetry.

So we should expect, Hardin claims, though without any certainty, that someone with normal neurophysiology would have normal and not in any way inverted color experience. And this suggests a way in which our experience of color qualia is constituted by our neurophysiology. It rests crucially on a posteriori knowledge of the neurophysiology of color vision.

I believe that the best strategy for coming to an understanding of what neural states constitute experience of what sensory qualia is to push something like Hardin's general strategy very hard for all features of normal human qualia that are not plausibly veridical, like the structure of unique hues. Once this initial generalization of Hardin's strategy is successful, we should then take another step. We should treat the remaining qualia involved in human sensory experience in the same way.

Still, I have differences with Hardin. For one thing, I don't like his talk of coding, which seems too reminiscent of information-based accounts of representation, without the elements of the world required for that information to exist. We will also shortly see that various details of his original proposal are at least misleading, that there are highly salient complications. But my main difference with Hardin is that I think we should deploy something resembling his account across the board, to underwrite all the differential features of our sensory qualia. I think we should extensively generalize Hardin's attractive proposal.

In this, I agree with Austen Clark.[28] But I have detailed qualms about Clark's generalization of Hardin. Clark maintains that the qualia involved in a particular sense are characterized solely by structural relations within

---

[26] This was noted in Hardin (1993), but Clark was the first philosopher to emphasize its importance.

[27] Yoshioka, Dow, and Vautin (1996).

[28] For other cases, see Clark (1993: 153–160); Mendola (1997: 129–135).

the corresponding space of discriminable qualities, which the possession of neurophysiological states exhibiting isomorphic structure would fully explain, leaving no difficult gap in our understanding.[29] I think that this is too optimistic. One way to put my point is that Clark's view implies that we can develop what he calls an "objective understanding" of the phenomenal consciousness of a bat. We can do this by coming to understand the very abstract structure of its qualia space, the structural relations among the qualities it "hears" that are fixed by the discriminability and similarity of the properties it senses aurally and the location in that sensory "space" of sensory primaries (analogous to unique hues). And he thinks that this objective understanding is adequate to defeat Thomas Nagel's worry that humans will never understand the experience of a bat.[30] But I think that such an abstract structural understanding is not enough for an understanding of what it is like to be a bat sufficient to defeat Nagel's worry. Another way to put my point is this: In Clark's proposal, *very* internal mechanisms of discrimination are what is crucial to the constitution of experience. Even possession of specific internal sensory receptor cells is apparently unnecessary. This allows Clark to claim that qualia inversion of intuitive sorts can occur consistent with his proposal, for instance when these very internal mechanisms are attached to sensory receptor cells in inverted ways.[31] But it seems to me more plausible to claim that the relevant sensory receptors, say the relevant cones, play a crucial role in constituting the nature of our qualia. Some of Hardin's claims about the asymmetries that constitute experience of yellow, for instance the relevance of saturation, seem also to imply this.[32]

But despite my differences from Hardin and Clark, my proposal has important similarities to their views. It adds together Hardin's less abstract conception of relevant structure and Clark's generality. But it talks about modal structure rather than coding. And as we will see, it also differs from Hardin's in key details. But it will be useful to remain somewhat abstract for the moment so I can insist again on some central points and how I think they should be understood.

To each particular mode of sensation, say color vision, there corresponds a structure of alternative possible qualia, and to be in a state of experiencing such a property is to be in one element of a structure of

---

[29] Clark (1993).

[30] Clark (1993: 206–209) and Nagel (1974), but note his qualifications.

[31] Clark (1993: 201). He doesn't talk about variations in connections to sensory receptor cells, but rather to "stimuli."

[32] Hardin (1997: 296–297).

alternative possible states of our neurophysiology. There is a kind of modal structure shared by the qualia and the neurophysiological basis of such experience. For instance, I think the only real possibility for living animals on earth with neurophysiologies like ours to have the experience of unique red requires that it be realized in a neural structure capable of alternatively providing experience of unique green. Phenomenal green and red may seem on the surface like quite distinct components of our experience, but, if we probe their appearance, we see that their being is entwined, in the sense that the actuality of one would require the subsistence of the other. And this is because, I think, the neurophysiology of vision implies that the experience of one requires the capacity to experience the other, in the sense of "required" at issue. I don't deny that individual dualist colored sense data may be in some broad sense metaphysically possible, and might constitute the experience of some possible creature in some other way, who might experience red without being capable of experiencing green, but I do deny that is a relevant real possibility for life on earth. And I am not even claiming that it is impossible for you, through some ingenious but wildly hypothetical amalgamation of neurophysiology and dualist qualia, to have that capacity. I am talking about the real possibility of such an experience of green for you, which is assured by aspects of the same neurophysiology that in fact constitutes your experience of red. All that is required for that experience of green is a different pattern of firing in your fixed neural context, and even a pattern that would naturally occur under different conditions of normal stimulation. This is not a sense of real possibility requiring strange subsisting possibilities; only alternative states of which our actual neurophysiologies are obviously capable. In this sense, the apparent modal structure of phenomenal color is a reflection of the true modal structure of what constitutes our experience of such color. Hardin has provided one neural model for how this could work.

Like Hardin, I also stress the importance of specific sensory receptor cells more than Clark does, and this involves another element of modal structure worth noticing. Assume that the cones in particular are crucial to color experience. Then while actual external stimuli aren't crucial to the constitution of our experience of color qualia, although they do cause such experiences by inserting photons into sensory receptor cells so as to excite them, still the nature of those sensory receptors is crucial. Even their unexcited nature constitutes a range of dispositional truths about what *would* excite them that may well be relevant to the constitution of our experience of color qualia. I think it is also important that qualia are generated by stimulation of sensory receptor cells, because they are in that way a response to and presentation of the world. Color qualia seem to be features of the world because they involve, because they are in fact forms of, the mediation of sensory stimulation and action. They are rooted in the presentation of the environment and root our reaction to it.

That is the general idea of what I am proposing, although we will need to consider further details. But first it will be useful to consider some general objections and other contrasts.

4

There is a standing objection to accounts that claim that qualia are complexly constituted in such a way, and in particular that relations among qualia are constituted by relations among elements of the neurophysiological basis of our capacity for such forms of experience.[33] The objection is that qualia seem to be intrinsic and not relational properties of our experience.

I grant that it may be natural to quickly conclude on the basis of the appearance of phenomenal green, green as we experience it, that it is an intrinsic property, and even an intrinsic property whose experience cannot be constituted in the relational way I have just suggested. But it is also just as natural to quickly conclude on the basis of the appearance of violets and some reds that they involve close objective similarities of objects out in the world. Natural quick conclusions on the basis of the appearance of colors alone are not sufficient evidence on these matters. And, in any case, we have been pushing harder than is common on the nature of phenomenal colors, on the metaphysics they would require to be veridical, and we have discovered they are not as independent in being as a cursory look would suggest.

There are also considerable standing disputes about the nature of color, which may suggest objections to a Hardin-style view like mine. But remember that I make no claims here about what color is, only about what gives us experience of color qualia. Hardin himself, on the basis of his proposal, is an irrealist about color. He thinks that color must have the features that color qualia present, including the structure of unique hues, so that if there is nothing in reality that has those features, then there is no color. That is consistent with what I am claiming here. But it is not a view that is widely popular among philosophers, who are frequently concerned to preserve the truth of commonsense claims. And I am not claiming that

---

[33] Levine (2001: 96–104).

irrealism is true, though we are done with irreducible Aristotelian colors. So it is worth noting that a neurophysiologically based view in the general spirit of Hardin's proposal about color qualia is plausibly required even by alternative physicalist views of the nature of color. What I have to say about color qualia is consistent with these views. Some philosophers are relativists or relationists about color,[34] but they are so partly on the basis of the recognition that experience of color varies among individuals in the way suggested by a Hardin-style view. There are also non-relational physicalist accounts of colors, for instance those that identify colors with surface spectral reflectances. But even their stories about color similarity, at least when they are remotely workable, are in fact based in a somewhat indirect way on just such a story about color qualia. According to the hue-magnitude account,[35] which is the best developed of these proposals, to each side of an opponent chromatic process, there corresponds a magnitude of engaging, or being disposed to engage, that side of the process, which is a property external objects may have. To be unique red is to be such as to engage the red side of the red-green process maximally and not to engage the yellow-blue process in either direction. More generally, objects are said to be presented to us as having proportions of these hue magnitudes. And purple things are more similar to blue than to green because they share a higher proportion of the blue magnitude. But notice that this way of representing objects is crucially dependent on the same general type of story about color experience that Hardin tells, and that I am pursuing in a certain specific form. Objects are presented as having complex dispositional and relational properties, on these views, which depend in the end on the way they interact with our neurophysiology of color vision. It is reasonable to complain about such views that such properties as colors certainly don't appear in experience in such a complex dispositional or relational way. But perhaps color qualia themselves are supposed to be on such stories (I think implausibly) somewhat inaccurate ways of presenting such objective complex and relational properties, more or less as if qualia had qualia.[36]

So for physicalist accounts of color of any sort, something like Hardin's story about color qualia is required, and my variant would do. There are even some dualist views to which it may be closely compared. For instance,

---

[34] Cohen (2009). [35] Byrne and Hilbert (2010); Tye (2000: 145–170).

[36] Another physicalist attempt to account for unique hues depends on the fact that when objects are seen as having unique colors, there is less variation in changes in their effects on the three types of cones introduced by changes in illumination over normal ranges. See O'Regan (2011: 129–130). But this too depends on the neurophysiology of human color vision.

Adam Pautz argues, partly on the basis of ways in which sensory qualia reflect internal structural features of perceivers rather than objective physical features of physical objects, in other words partly on the basis of just the kinds of points we have been considering, that a particular form of dualism (or at least non-reductive grounding physicalism) is correct. There is alleged to be a primitive relation of being conscious of, which conceivably supervenes on but is not constituted by the physical. And in the case of color experience, it is alleged to relate us to primitive colors, colors such as Aristotle conceived colors to be. But these primitive colors to which we are related are not instantiated in reality. Still, they are present in reality. They are existing but uninstantiated properties, which "live only in the contents of our experience."[37]

I believe and have argued elsewhere that the notion of an existing but uninstantiated property is not coherent.[38] But I have claimed here that in color experience there is an appearance of something that, were it actual, would be entwined in being with something that merely subsists. And the mere appearance of a property that, were it to exist, would involve the subsistence but not existence of some other property is not wholly unlike the existence of an uninstantiated property, indeed at least two of them. Both the property that merely appears to exist and those that would subsist in relation to it, according to my story, are in different ways somewhat close to being uninstantiated properties. And if indeed Pautz' uninstantiated properties literally "live only in the contents of our experience" as he says, they wouldn't quite be like traditional uninstantiated properties in any case, which exist independently of being experienced. So there is a kind of confluence even between this form of dualism and my way of interpreting a Hardin-style proposal. Still, of course, my proposal is not a form of dualism. The relation of being conscious of such a property is, on a Hardin-style view like mine, far from being a primitive relation that is not constituted by the physical. It is instead, in reality, constituted by the physical in different very messy ways for different cases, which we are beginning to explore here.

As I've said, we will return, in a skeptical way, to the central motivations for dualism about qualia in Chapter 8. But the only reasonably realistic stories to tell about our experience of color qualia are roughly of the sort proposed by Hardin and Clark, and also favored by most color scientists. Still, I also claim, it is important to think about these stories in the right way, not only in a physicalist but also in a modal structural way.

---

[37] Pautz (2006, 2010).

[38] Mendola (2021: chapter 7).

To understand the need for this modal structural conception and some of its details, it will be useful to consider now ways in which the particular story about the neurophysiology of color vision that Hardin originally told is oversimplified, and also to sketch some additional positive aspects of my specific proposal.

5

The sorts of particularity and spatiality present in color experience are complex matters, which we will begin to consider in Chapter 6. But even in the most rudimentary sort of case, color qualia are presented to us as involving location. You see red to the right but green to the left. Whatever allows us to experience a red quale, it certainly does not forbid us from experiencing a green quale in an alternative location at the same time, even though those are opponent hues. While our experience of space is complex, it surely involves the integration of sensation and action at least to the degree that you can point with your hand towards where the green quale or the red quale seem to be located on some concrete object in front of you.

My claim is that each individual color quale is constituted by a distinct "modal filament" within the nervous system, bridging sensation and action in the experiencer. What is a modal filament? Let me begin by characterizing the simplest type of case. Such a simple modal filament is some variable aspect of someone's current neurophysiology that leads from the sensory receptor stimulation that causes that element of their experience, say of the red quale on their right, to their consequent action. It is the causal-explanatory chain that links that particular stimulation with that particular action, which explains the latter on the basis of the former in the context of the relevant relatively fixed neurophysiological background. Given the complexities of modal structure of the nervous system, there are no doubt difficulties in specifying relevant fixed backgrounds, and the notions of causation and explanation themselves harbor many difficulties, and causation may incline with a certain probability rather than necessitate, but that is the rough idea of the simplest sort of modal filament. And let me focus on what it is not. As the word "filament" suggests, it is not a single spatial location; it leads from one place to another. But contrary to the suggestion of that word, it need not be a spatial filament, like a piece of string, which involves a specific spatially localized path through the nervous system, down a single chain of neurons, for instance. It has a modal, rather than a spatial, character of that sort. What is crucial is that it can itself vary independently of other current qualia constituted by other modal filaments present within the same relevantly fixed background, independently of whatever for instance leads from sensory stimulation to our subject's experience of a green quale on the left and on to consequent action. Its individuality is modal in that way, not necessarily spatial. And a modal filament is also modal in another way. It itself involves the kind of opponent structure that is apparently constitutive of your experience of a red qualia. The opponent structure we have been considering as constitutive of color experience must characterize each modal filament constituting your experience as of a phenomenal color as at a certain location. That modal filament has modal alternatives that would be the alternative experience of a color quale of a contrasting sort, perhaps any contrasting sort. That modal environment of alternatives is crucial to what it is. This is not to deny that fixed environments other than that fundamental modal one may also be crucial. To say that such a modal filament is capable of independent variation is not at all to say that whatever visual neurophysiology constitutes it, that would also constitute such an experience independent of the rest of your neurophysiology, say if we cut out your visual system and put it in a jar. It is plausible that quite different but normal humans share a capacity for such an experience, and hence that the requisite background neural environment of this sort may assume some variety of individual forms, but that is not at all to imply that some such environment is not necessary.

There are of course many complications. There is a question about how many color quale and so modal filaments you have. But the most immediately important complication is that qualia-constituting modal filaments are certainly not all of the simple type just noted. Experience of a certain color at a certain location does not mandate specific action, as that simple case suggests. You can point with your left hand or your right when asked to signify the presence of such a green quale, or point away from it, or simply talk about it, or do nothing at all. At least specific action *tendencies* may help constitute differences between some qualia. But even action tendencies aren't relevant in all cases. For these reasons, it is worth noting immediately that the modal filaments that constitute qualia need not lead all the way to external action of particular sorts. And even when they do, they may constitute mere tendencies to particular actions. But, on the other hand, sometimes even the actions that some filament will, or probably will, lead to may constitute something of its qualitative nature. Perhaps pain qualia are like that, as we will later consider.

These details matter. But the main point to notice now is that, on this understanding, there need be no particular location in the brain that constitutes the experience of given quale. We should not presume that there is a kind of local coding pattern in certain specific kinds of neurons of the visual system that constitutes in an opponent way the experience of red and green, as Hardin's formulation of his proposal suggests, and in the way traditional representationalism presumes. Rather, the modal filament that constitutes experience of a certain color quale is something that leads from sensory stimulation, perhaps all the way to action. And it needn't even do that in the spatially localized way of a string. It could involve a firing pattern spread all over the brain, although it may originate in specific sensory receptors. And it is "surrounded" throughout its "length" by various sorts of modal structure, by various unrealized alternatives that help constitute it as having overall a certain type of modal structure relevant to its phenomenal character. The sensory receptor cells whose stimulation begins a certain filament might have been stimulated in other ways and fire in alternative ways, and that may be relevant to the filament's nature. And some motor action or neural precursor of motor action to which it leads in a characteristic way might have been caused in other ways or have various alternatives, and that may be relevant too. Perhaps most important, throughout the entire length of neural action that mediates sensory stimulation and action within the neurophysiology underwriting such a filament, there are various forms of locally alternative and hypothetical neural action that may be relevant to its modal nature as well. The modal filament may in that way gather relevant modal structure along its length.

There are various forms of similarity relevant to qualia, for instance the similarity of hues. And similarity is relevant to modal structure, even the modal structure of our neurophysiology. But it is important not to think of these similarities inside the relevant neural maze in an essentially spatial way. It is clear that the spatial arrangement of sensory receptor cells is relevant to our experience of space, and we will see later that the spatial arrangement of motor action is also crucial. And I grant that between sensation and action, our intervening neural activity, for what amounts to engineering reasons involving distance and signal speeds, probably exhibits some correlations between qualia similarity and spatial distance, or other similarities of spatial patterning. But what is crucially significant is how the intervening neurons are hooked up in that intervening maze, necessitating or inclining in the context of available alternatives, not that in being so hooked up they end up happening to be in various spatial relations. What matters is what is alternative to what and how that affects neural dynamics, not where it is located.

The modal filament that is a certain quale is analogous in these various ways to the actual core of the superworld or to the fully present core of hyperexperience. Its complex modal environment is crucial to what it is. In fact, the fully present core of hyperexperience is a collection of such modal filaments on my proposal. Perhaps some of its less present parts are constituted by aspects of the relevant modal environment of those modal filaments in our neurophysiology, or by modal filaments of another sort.

With this rough model in mind, which we will continue to refine in Chapters 5–7, we turn now to a review of important details of the neurophysiology of color vision. In the end, they require modification of Hardin's initial proposal. And they fit better with the modal structural way of understanding the relevant neurophysiological facts involving modal filaments, instead of Hardin's more traditional way that focuses on spatially localized coding.

6

There are three detailed aspects of our neurophysiology that are conceivably relevant to the constitution of our color qualia.

Begin with the sensory receptor cells involved.

As we've noted, human vision is mediated by the stimulation of photoreceptor cells. These are living neurons sensitive to visible light, to electromagnetic radiation. This radiation consists of photons, but such "particles" are also wavelike. Light visible to humans has a limited frequency and wavelength range. The relevant range of wavelengths is approximately 390 to 700 nanometers.

The human visual receptor cells are of two basic types, called "rods" and "cones." The rods register differences in illumination, and are important in achromatic vision. But the cones are central to chromatic vision. Here's how these cells work.

Opsins are a class of proteins. They originally evolved in ancient bacteria, in which they were used for photosynthesis. In some bacteria, opsins help guide the creatures in question towards light or energy sources, so they have a sensory-like function. Different opsins found in different creatures differ often by single amino acid, and mutate very easily, though within a limited range.[39] Rhodopsin is the photopigment in rods. In chromatic vision, things are somewhat more complicated. Opsins are present in cone cells that are themselves relatively fixed in structure, in the

---

[39] Matthen (2005: 178).

sense that through evolutionary pressure, it is relatively easy for the particular opsin involved in a cone to evolve. And there are differences among the opsins in different types of human cones, covalently bound to other "chromophore" molecules.[40] These photopigment molecules are made up of a chain of carbon atoms with some additional carbon, hydrogen, and oxygen atoms off to the side. Some of the bonds between neighboring carbon atoms in the chain are double bonds, which like parallel pins through two grapes prevent mutual rotation. But some are single, and so allow mutual rotation as on a single spindle. Light, even a single photon, may break a double bond, allowing for rotation and a change in the internal configuration of the photopigment molecule.[41] This change of configuration of the photopigment molecule causes it to sever its previous connection to another molecule, an opsin, which is a catalyst and a protein, built of amino acids. When that connection is severed, it opens up an active site on the opsin, which initiates a complicated sequence of catalyzed reactions.[42]

In a cone, there is a sodium gate that is regulated by a messenger molecule, which, when present, causes the gate to remain open. But at the end of the chain of reactions cascading down from receipt of a photon of the relevant frequency range is a molecule that shuts down the messenger molecule in question, hence closing the gate and making the neuron more negative, which hence modulates the natural firing rate of that neuron when a photon comes in.[43]All we see comes through such a complex chemical bottleneck. I think the general nature of these sensory receptor cells, not only their focus on a limited range of radiation but also their detailed chemical modal structure, may well be relevant to the nature of our visual qualia in general. But our focus will be on differences among such qualia.

There are different types of photoreceptor cells in humans, involving somewhat different photopigment molecules and opsins. Despite their crucial role in achromatic vision, even rods are differentially sensitive to light of different frequencies.[44] But focus on the role of cones in chromatic vision. Different types of cones have different patterns of sensitivity to wavelengths of light within the visible range. The so-called *L* cones are maximally sensitive to yellow, *M* cones to green, and *S* cones to blue light. As we noted earlier, each type of cone has a characteristic spectral absorption curve that gives the probability that a photon of a given wavelength

---

[40] Solomon and Lennie (2007: 276–277).

[41] Groh (2014: 14–16).

[42] Ibid., 16–18.

[43] Ibid., 17–20.

[44] Bowmaker and Dartnall (1980).

will be absorbed by the cone and hence affect its firing, although the form of that firing is independent of what wavelength of light is in fact received. And the curves of the different cones overlap, those of the *L* and *M* cones especially closely. Because of that overlap, a given photon may have a fairly high probability of affecting the firing of an *L* and an *M* cone. But on the modal structural understanding, it may also be relevant to the nature of a quale that it results from a cone that would have been engaged, or likely engaged, by photons of other specific wavelengths. That element of modal structure surrounds and characterizes the cones, and hence the early part of the modal filament that constitutes the quale in question.

There are complexities that attend even this trichromatic aspect of human color vision. For instance, some women carry a gene for a fourth cone pigment, which is also found in men who are anomalous trichromats and hence in which it replaces a standard pigment.[45] But the oversimplified story about receptor cells I have told should be close enough to truth for our current purposes.

Now turn to a second type of element of neurophysiology that may be relevant to the nature of color qualia. There are some ways in which it is conceivable that characteristic actions and not just sensory input are part of what constitutes even a color quale. And on the modal structural understanding, it is possible for different relevant aspects of a modal filament, the modal structure that surrounds it and characterizes it at various points over its length, to occupy different locations along such a "string." However, unlike such cases as pain qualia, there doesn't seem to be any very clear connection between particular color qualia and particular actions. But still, there are some candidates.

There is some reason to believe that primates evolved color vision like ours, which differentiates not merely yellow and blue but also red and green, because the red-green contrast allows fruit to pop out from foliage, which is useful in seeing small fruit at a distance among a dappled background.[46] Maybe something like that functionality is present in us and relevant to the nature of our color qualia through dispositions to certain sorts of actions, for instance eating certain things or avoiding others, which are engaged by certain color experiences rather than others. And we react strongly to the red color of blood, and are calmed by blues and greens. So perhaps there is some tendency to action that red gives us that blue and green do not, and can explain part of the difference in the relevant qualia.

[45] MacLeod (2010: 154–159).

[46] Matthen (2005: 180–183).

Matthen has suggested that epistemic actions – inference and learning and memory – are also relevant to the nature of color qualia.[47] This would plausibly involve, as we noted in Chapter 3 in discussing introspection, tendencies and structure upstream from specific bodily action. Still, it might seem that any characteristic epistemic actions involving red and green would be perfectly symmetrical, so that is hard to see how epistemic action could constitute their difference. But it is possible that uniform such action helps explain similarity in qualia despite differences in sensory stimulation. We treat all yellow stimuli somewhat the same in this sense. It is possible that different such action helps explain the difference between color qualia and those associated with different senses. Matthen speculates that some of our experience of sound qualities reflects the way we can, as instrumental musicians and speakers and singers, produce sounds, while we do not in similar direct ways produce colors. He also suggests that it is relevant that we deploy color qualia in inference in a particular way, that color is the last differentia. We do not make general inferences based on color alone but on color within a type.[48] Perhaps this is so. If it is, the modal structural understanding can incorporate it as relevant to the phenomenal character of color qualia.

Nevertheless, despite such possibilities, much of the neural structure relevant to differences among our color qualia – in particular to the structure of unique hues and color similarity – lies between our trichromatic sensory receptor cells and neural structure directly relevant to action. A third type of element of our neurophysiology is crucial. While Hardin was clearly on to something important about this, the details are much more complex than the simple model that he originally proposed. And to the degree they are known, they fit better with my proposal than with Hardin-like localized coding.

It has long been clear that there is some kind of opponent-processing in the retina, and in the lateral geniculate nucleus (or LGN), to which specific neurons of the retina project in a specific way. And it originally seemed to take a form in both the retina and LGN much as Hardin's proposal suggests.[49] Let me remind you that that proposal was that there is a red-green channel that subtracts *M* from *L*, and is hence positive when tipped towards red and negative when tipped towards green, which can tell us how strong a signal is at the ends as opposed to the middle of the visual spectrum. There is also a yellow-blue channel, which subtracts *S* from the sum of *L* and *M*, is positive when tipped yellow and negative when tipped

---

[47] Ibid., 14, 97–98, 229–232, 237–239.

[48] Ibid., 265–268.

[49] Ibid., 178.

blue, and tells us how much of a signal's strength is in the middle and long wavebands relative to the short. There is also opponent-processing of dark and light. And it is still clear that there are, in reality, distinct groups of neurons that do fit much of this pattern. The midget ganglion cells in the retina oppose the activity of L- and M-cones and project to the parvocellular layers of the LGN. Another group of neurons in the retina are strongly responsive to the firing of S-cones opposed to some combination of the firing of the L- and M-cones.[50]

But the identification of these opponent channels in the periphery of color vision with the unique hues has become much less plausible than it once seemed.

It was always clear that there must be other effects on color qualia by opponent-processing of different sorts. There is simultaneous color contrast. Oranges look brown if surrounded by lighter objects. Red objects seem less red if surrounded by very red objects. This allows some constancy of appearance over changes in illumination. But there isn't perfect color constancy. Rather, whatever is in red light looks somewhat different, but we are somehow able to separate out our sense of what is due to the color of the object and the illuminant. So opponent activity beyond the eye clearly affects our color qualia.

But the crucial problem with the original Hardin-style model is that the early hope that the structure of unique hues and the axes of chromatic similarity map straightforwardly on to the structure of opponent-processing in the retina and LGN has not been vindicated by later investigation. There are other explanations possible for the existence of opponent-processing in the eye. It is an efficient form of processing, since it removes redundant information, and the forms of opponency observed are special with respect to efficiency of chromatic detection, discrimination, and adaptation.[51] And it has become reasonably clear that the satisfyingly simple model proposed by Hardin must be wrong in detail, because of the way rod cells are involved in chromatic perception.[52] But the main difficulty is that retinal and LGN cells do not correspond well enough with the phenomenal axes their names suggest.[53] For instance, the red-green cells respond similarly to green and blue.[54] And a reddish spectral violet is what most excites the green channel.[55]

---

[50] Solomon and Lennie (2007: 279).

[51] Valberg (2001).

[52] Joesch and Meister (2016).

[53] Derrington, Krauskopf, and Lennie (1984); MacLeod (2010: 159–161); Solomon and Lennie (2007: 278–281).

[54] Abramov and Gordon (1994); MacLeod (2010: 160).

[55] MacLeod (2010: 161); Stockman, et al. (1993).

Perhaps, then, the relevant opponency can be located at some "higher" level of neural processing, perhaps in the V1 area of the cortex that is crucial in visual "processing" and to which the LGN projects.[56] Within V1, there are functionally distinct types of neurons that are sensitive to color.[57] And there is at least some evidence that there is a small population of LGN cells that do have the right opponency structure.[58]

But any way in which color is "represented" in V1 is at best very untidy and complex.[59] There is better evidence for spatial coding in V1 of intuitively intermediate colors arranged by similarity.[60] There seems to generally be rather broadband tuning of neurons to stimuli in the cortex, with relevant neurons arranged spatially, which may help explain why there is a unified experience of the values across the red-green, yellow-blue, and white-black axes, as well as why we experience one color space rather than mere individual colors, but which is less suited to delivering opponent structure.[61]

The bottom line is that a claim that there is some localized opponent coding of the unique hues in particular anywhere in the brain is currently entirely speculative.

But on the modal structural understanding we are developing, that is not required. Rather, a modal filament may gather along its length the necessary modal structure. And it seems that all these various detailed discoveries and complications suggest just such a conception. There is some opponent organization in the retina and LGN, but also an inflection and complication of that modal structure at "higher" levels, farther along the modal filaments in question, so that the filaments do have the relevant structure. The fact that there is not explicit and straightforward representation of the intuitive unique hues in the eye or LGN does not imply that what happens in the eye or LGN is not part of what constitutes some relevant filament and hence quale to involve the unique hue that it does. One way to put it is that just as trichromacy and any action-inclination constitutive of red qualia may help compose relevant modal structure without being at the same location in the brain, so too the opponency of color vision and whatever specifies the exact frequency of unique hues may help compose relevant modal structure without being rooted in the same

---

[56] Hardin (2014: 380); Stoughton and Conway (2008). But see also Mollon (2009) and the immediately following reply by Stoughton and Conway.

[57] Solomon and Lennie (2007: 283).

[58] Hardin (2014: 379–380); Tailby, Solomon, and Lennie (2008).

[59] MacLeod (2010: 161).

[60] Parkes et al. (2009).

[61] See Matthen (2005: 133–141) for this notion, due to Gallistel.

location.[62] The phenomenon of simultaneous color contrast should lead us to expect such a structure in any case.

Another objection to Hardin's original model is that it might be thought that opponent-processing cannot constitute the unique hues of color qualia, because there are unusual situations in which objects seem to be reddish green.[63] There are translucent things that look orangey green. Other such phenomena involve the filling in of colors, much as in what covers up for us the blindspots of our eyes.[64] This suggests that at some level red and green are registered in other than opponent channels. There are other color illusions due to competition between lower-level perceptions to which this phenomena is perhaps analogous. But none of this is a problem for modal structuralism. It is just another complication in the nature of the relevant modal filaments. For instance, the relevant modal environment may undergo a fission along the relevant string.

What is crucial is not spatial complexity but modal complexity. Things can vary while other things stay the same, and they are related in various ways to things that do not occur but could, and that modal structure of our neurophysiology delivers the relevant structure of the phenomenal. Red such as we experience in the world itself would require different degrees of being. And that is something in which it is pretty hard to believe. But the apparent modal ontology of such a field of color reflects the way we are in fact capable of such an experience, which goes some way towards explaining what would otherwise be a very mysterious phenomenon. Any neurophysiological structure capable of a fully actualized state of experiencing phenomenal red is also capable in a real but less actualized sense of experiencing phenomenal green. And this goes some way towards explaining the strange but illusory modal metaphysics of color as we experience it.

There is an objection to my way of attempting to avoid Hardin's difficulty. Since the structure of unique hues, for instance, can be revealed in behavior, it may seem there must be some location in the neurophysiological progress linking stimulus and action where the modal structure of the hues matches that of the mechanics of neural activity. But, I reply, the relevant mechanism is modal. This does not require that there be some

---

[62] Does this imply that trichromacy is in some way introspectible, for instance that we can, by deep inspection of our color qualia, foretell the possibility of color blindness of various sorts? Not necessarily. But it is also relevant that there were historical attempts to deliver color qualia structure solely from trichromacy, which may suggest this. Mollon (2012) sketches trichromatic theories, including contributions by such luminaries as Helmholtz, Grassman, and Young.

[63] Billock and Tsou (2010); Crane and Piantanida (1983); Hardin (2014: 380–382).

[64] Piantanida (2010).

specific location in the brain where that mechanical structure of opposition is located. There may of course be such a location involved in the instigation of one particular behavior over another that helps reveal the phenomenological structure in question, but still that location needn't always be part of the modal filaments constituting the color qualia involved, and in any case what feeds causally into that action, and would feed into alternatives, may be complexly distributed both spatially and temporally. Perhaps there are two different spatial streams of causal activity that affect what is done, where the second is relatively fixed and perturbs the action to which the first would otherwise lead, and such that the second generates corrections in what would otherwise be seen as unique hues. Similar things might be true regarding the causal control of upstream neural structures from which a variety of bodily actions might originate, in cases where relevant modal filaments do not extend all the way to specific action.

I have presumed throughout our discussion that a certain sort of primitive spatial particularity is present in our experience of color qualia. What it is to be a minimal bit of color in the visual field is to be seen at one point or another. What it is to be a particular bit is hence dependent on being at a certain point in the field, which point must be filled with one color or another, and is somehow fixed in its particularity by relations to other bits of the field. We will need to come back to a proper modal treatment of our experience of that particularity. But the general idea is that the modal structure of our relevant neurophysiology is reflected in the modal structure apparently present in this sort of experience of particularity as well, and not just in the apparent modal nature of the color.

7

That completes my proposal about the color aspect of color qualia.[65] There are detailed aspects of this proposal that are of course controversial.

For instance, I suppose that the neurophysiology of the rods and cones is crucial to the nature of color qualia. But there are now artificial retinas, and if they turn out to preserve color qualia, it would imply that this is wrong.

---

[65] How do the differences between different sense modalities enter into this modal structure? For instance, is difference from auditory modality essential to the experience of color? There cannot be much structure imposed by the differences between modalities, because the deaf can see the colors we see. But it would seem that the ability to differentiate color and sound is relevant.

Still, current artificial retinas allow only significantly abnormal visual experience.[66] And the general modal structural conception does not demand the relevance of the sensory receptor cells, only allows it.

My proposal also predicts controversial things about color blindness. Any account of the general type of Hardin would predict that disability of the red-green system would generate red-green color blindness and of the yellow-blue system would generate rarer yellow-blue color blindness. And color blindness does correspond to deficiencies in the production of different optically sensitive pigments in a way that roughly corresponds to this prediction. But there are detailed differences in the predictions of different accounts of this general type.

For instance, Nida-Rümelin has proposed that opponency may even allow for the existence of real qualia inverts. Some red-green color-blind people lack one standard pigment and some another, because of genetic defects. But both defective genes may be active simultaneously in one individual, hence switching the pigments between cells.[67] Perhaps these are real color inverts. My own view, however, is that this switch would simply switch which cones are relevantly red and green.[68] So my particular detailed proposal disagrees with hers at this point, although again that is not required of any modal structuralist account.

Another difference regards the unity of color structure. Does the fact that someone is red-green color-blind affect their experience of blue and yellow? Hardin's own proposal suggests not. But I think it may, and if so the color normal may have no way to fully understand the color experience of the color-blind, just as the color-blind cannot fully understand color-normal experience. In fact, on my view, any difference in modal structure of the features that centrally constitute color qualia would probably introduce some limited failure of this sort, though that failure must plausibly be capable of coming in various degrees. It is worth noting that we generally presume that the dark-light opponency is unaffected when considering such cases. If it were affected, there would plausibly be further effects on other color qualia. But I admit that it does seem implausible to claim that experience of pure black and white is significantly affected by color blindness, which is a limitation on my scheme.[69]

Perhaps these claims about color blindness suggest an empirical test for my specific proposal relative to others of its general type. For instance, there is some evidence that the experience even of yellow things by the red-

---

[66] Rojahn (2013). [67] Nida-Rümelin (1993). [68] Ibid., 152–154.

[69] Thanks to Reina Hayaki for this point.

green color-blind is phenomenally different than normal experience in a relevant way. Some rare individuals possess one color-normal and one color-blind eye, and can introspect differences, and at least one such subject perceives the yellows seen by the color-blind eye as orange.[70] That is the kind of phenomena my proposal would predict, contrary to Hardin's. However, on the other hand, notice that this case also suggests that the tipping of at least a normal red-green system isn't necessary for experience of reddish hues. At least red-green opponency at a very peripheral level, in the eye, is not necessary to see things as reddish. This may seem contrary to my proposal. But recall that I am not attempting to root experience of unique hues and opponency solely at very peripheral levels of visual neurophysiology. So it isn't really contrary after all.

Hilbert and Byrne distinguish "the reduction view," the view that the color-blind simply see a reduction of the colors the normal see, to "the alien view," the view that all the colors seen by the color-blind are alien. Worried about claiming that the color-blind misperceive the colors of the world, they prefer a form of the reduction view that says that those who are red-green color-blind see things as yellowish and bluish only.[71] And they think there are no strong arguments for the alien view, because they don't think that experience of all the similarities of unique yellow is necessary for experience of that color. I might agree about that point about *color*, but I think the similarities in question are essential to the *color qualia* involved. So I think there are somewhat alien qualia involved in color blindness.

Justin Broackes has argued that the red-green color-blind can sometimes distinguish red things from green things.[72] But again I am concerned only with the qualia in question. It is implausible to think that the fully red-green color-blind have the same color qualia as the normal and simply miss the difference between the red and the green ones.

This discussion of different views about color blindness puts us in position to properly consider another worry about one aspect of my proposal. I have claimed that you have faint hyperexperience of red when you see an orange chair, which reflects the entwined subsistence and existence that would constitute such an orange in the superworld. I should admit that my view also suggests that you have faint hyperexperience of green when you see a red chair, which may seem more intuitively objectionable. It may seem that we cannot introspect such a green just on the basis of seeing red, that experiencing red involves no mere

---

[70] Kaiser and Boynton (1996: 452–455).

[71] Byrne and Hilbert (2010).

[72] Broackes (2010).

hyperexperience of green, that red seems to be an atomic property in at least that sense.

My first reply is that the essence of such a phenomenal red is bound up with green in the superworld in the way I have suggested, and that to experience such a color is to be confronted with its essence in a way that ineluctably involves green. But it may seem that my claims about introspection and hyperexperience are too strong regardless. So a second reply is required.

Consider this series of cases of what might be called "sensing through to." When you experience a sound of a certain pitch, it is intuitive to say that you can tell by immediate introspection of the nature of that sound that higher and lower pitches are also possible, with various essential relations of similarity and difference to the present pitch. And I think they even seem to have a kind of ethereal presence in your awareness by way of contrast, which you can introspect when you make that judgment about possibility. You are in some sense aware of those alternatives. In hearing the original sound, you in some sense hear through to those alternatives. I think you have mere hyperexperience of the alternatives. Remember that there are less than fully central and paradigmatic forms of phenomenal consciousness and introspection that we expect to occur in such circumstances.

Now consider some cases involving color. If you see a gray, the possibility of pure black and pure white seem quite salient. You see through the gray to black and white. If you see a yellowish orange, then less and more yellowish oranges, and unique yellow, are also salient in that way. And if you see unique yellow, then I think you see through to orangey and greenish yellows, although I admit that is a more questionable case. If you are prepared to go that far, then it may still seem that for instance red cannot be so discerned in green. There may be a relevant difference between some of these color cases and others, which may in turn support a different conception of color-blind experience than I have suggested and a somewhat different modal structuralist account of phenomenal color. But I think the experience of green invokes the faint hyperexperience of both blues and yellows in this way, including unique blue and yellow, and that they in turn invoke the faint hyperexperience of reddish blues and yellows, and so unique red. I think you can see through unique green to unique red in this way, even if it is somewhat indirectly. It may also be that the red afterimage that concentrated and extended experience of green can generate is latent in that green experience in a similar way, although that seems perhaps more analogous to some form of primary imagination. And indeed another possible treatment of your experience of green is that rather than then seeing through to red, you simply have the capacity for primary imagination of red and have that sort of experience when you turn your attention to the essence of the green you experience. But I believe there is something like the experience of the specific absence of red entwined in the experience of green, constituted by neurophysiological structure that underlies the relevant potentiality for red experience, and which can be dynamically relevant to what occurs between stimulation and action in your neurophysiology. Still, I admit that this specific case involves a questionable sort of hyperexperience, and its occurrence is not essential to my general proposal here.

8

There are many other interesting complexities of color vision, even beyond the ways it involves spatial structure that we will consider in Chapters 6 and 7.

One such complexity that involves modal structure occurs in our perception of the surface colors of objects. It is plausible that some of our sense of what the color of an object is depends on our sense of how its color look, its apparent phenomenal color, would change as the illumination changes.[73] Consider for instance one enactivist treatment of such cases: The worldly color of the wall "is a way its appearance varies as relevant conditions change, . . . as ambient light darkens over the course of a day, or as the course of illumination moves, or as the object moves from one sort of lighting . . . into a different sort."[74] "Our ability to perceive the wall's color depends on our implicit understanding of the ways its apparent color varies as color-critical conditions vary. At ground, our grasp of these dependencies is a kind of sensorimotor knowledge."[75] The dependencies allegedly include both movement dependencies, how the apparent color will change as we move, and environmental dependencies, how they will change as the relation of the colored objects to other objects and light changes. The objective color of something, on such a view, is a disposition it has to change its appearance as relevant viewing conditions change, but this is not the usual sort of dispositional analysis of color as involving a disposition to cause experiences in various creatures. I think there is something important to this idea, although one worry is how two individuals with two different forms of know-how regarding the change of color

---

[73] Noë (2006: 123–161); O'Regan (2011: 127–136).

[74] Noë (2006: 141).

[75] Ibid., 129.

appearances, say rooted in different particular objects and lighting conditions and motor activities, would have, on this view, experience of the same objective colors, as they intuitively should.[76] We have been focusing here on basic color qualia. But it is possible that even color qualia involve this second kind of apparent modal structure, involved for example in our sense of the wall's color, in their very being, on top of the essential similarities among different color qualia subsisting and existing on which I have focused. In any case, this is another sort of modal structure of visual phenomenal consciousness that should mirror that of our neurophysiology in some way.

There is another crucial element in our experience of color on any such view, color looks or appearances, and here I disagree with some enactivists. According to radical enactivists like Noë, there are no fundamental or atomic or basic elements of experience such as phenomenal colors. Rather, qualities "are available in experience as possibilities, as potentialities, but not as givens."[77] Consider the phenomenal differences and similarities of colors, which are essential to experienced colors. According to Noë, our experience of colors "is shaped by our implicit grasp on their positions in color space. We experience them as imbued with possibilities of variation, as possessing degrees of freedom in a space of phenomenal possibilities."[78] Even experience of color appearances is itself supposed to consist of know-how. Our know-how relevant to basic color looks is supposed to be our implicit grasp on their positions in color space. In this way, we are supposed to experience them as imbued with possibilities of variation. But, I complain, there are the fundamental color qualia we have been considering, and movement in color space is only a metaphor for real behavioral movement of the kind properly central to enactivism. Particular phenomenal colors do not move through color space in even the way objective colors change looks as illumination varies. But I also believe that even this radical enactivist account of basic color looks is on to something important, which I have been trying to capture in a less metaphorical way by claiming that the very being of actual colors such as we experience would be involved with mere possibility in various ways. So a concessive way to put it is that I have been trying to make literal sense of Noë's metaphor. And I think we can gain some further illumination about one aspect of how this works from Noë's response to an objection. Consider your experience of an opera singer holding a note that goes on and on. It seems like right now it has been going on for a long time, but no set of

---

[76] Thanks to Shane George for this objection.

[77] Noë (2006: 135).

[78] Ibid., 137.

sensorimotor skills seems to capture that, since the past is gone and inaccessible. Noë's response is that the note you hear now is an event with past and future aspects, and that what you hear is not its past aspects but rather its current aspects as surging forth out of the past, as part of a certain trajectory.[79] I think we see color qualia as surging forth out of contrasting possibilities, although this doesn't involve sensorimotor knowledge. I will develop analogous claims about temporal experience in Chapter 7.

---

[79] Noë (2012: 77).

CHAPTER 5

Some Other Senses

The modal structural account of the differences among our sensory qualia has now been sketched. But consider qualia involved in some other senses.

I

Among the basic properties of sounds present in our experience are loudness and pitch. Begin with loudness.

Sound waves are pressure waves, which have a frequency and hence a wavelength, but also an amplitude. The size of the amplitude of a pressure wave might be identified with its degree of loudness, or, alternatively, the production of a certain amplitude by a sound source might be identified with a certain loudness. There are some intuitive pressures towards saying that the loudness of a bell is the same whether it is heard by someone close up or far away, in other words that its loudness is determined by the amplitude of the sound wave where it is produced. But we are interested in sound qualia in particular. So there are more salient pressures to say that the experienced loudness of a sound depends on the amplitude of a sound wave at the ear.

But reality is more complex than this, and blocks all these candidate identifications.

One complication is that there are two kinds of amplitude of such a wave. There is a pressure amplitude and a displacement amplitude, which, however, have a fixed relation. What is called the "intensity" of a sound wave is approximately proportional to the square of the pressure amplitude, and also of the displacement amplitude. The intensity of a sound wave is the rate at which energy flows through a unit area perpendicular to the direction of travel of the wave. So now we have three candidates for objective loudness even if we restrict ourselves to properties of sound waves at the ear.

But none of these environmental properties corresponds well with heard loudness, even at the ear. One familiar and objective measure of what is

https://doi.org/10.1017/9781009514774.005

called loudness is on a logarithmic scale of intensities, in decibels. This corresponds much better to heard loudness than does intensity itself or the amplitudes. But it is only a rough physical approximation of heard loudness, and in addition some explanation is obviously required for why this logarithmic scaling is appropriate.[1] Another complication is that we are more sensitive to certain frequencies of sounds. We are most sensitive to sound waves with frequencies between 1 and 4 kHz. For frequencies above or below that range, our sensitivity, and so heard loudness, declines, until sounds below 20 Hz or above 20 kHz are inaudible to us.[2]

So it is clear that heard loudness depends crucially on the nature of our ears and our neurophysiology of hearing. Adam Pautz has gathered some evidence about neurophysiology relevant to experience of loudness:

> Relkin and Doucet (1997, p. 2738) write that "the perceived loudness of a pure tone appears to be linked both to the number of spikes fired by single neurons and to spatial spread of excitation in the auditory nerve." Langers et al. (2007) used fMRI to look at neural activity further downstream in the auditory cortex. They found that "cortical activity is more closely related to the perceptual loudness level of sound than to its [external, physical] intensity level" (p. 714) and indeed report "a type of non-linearity ... comparable to that reported in psychophysical studies on loudness perception that employ subjective loudness scaling" (p. 716). On the basis of this study and others, Röhl et al. (2011, p. 1494) conclude that "the most simple interpretation would be, that AC [auditory cortex] is fed by ... the auditory brainstem according to the sound pressure level and the bandwith [*sic*] of the stimuli, and an additional component is added which is linearly related to the perceived loudness."[3]

If we interpret these results in accord with modal structuralism, it appears that different levels of experienced loudness of a sound are constituted by different levels of neural activity in areas downstream from the ear. Different possible levels of neural activity at such junctures are the modal environment that helps constitute the nature of the relevant modal filaments reaching in from stimulation of the ear, so that they involve one level of experienced loudness rather than another.

It may be possible for objectivists about loudness to develop more sophisticated and less problematic identifications than the simple ones

---

[1] But for relevant complications, see Schnupp, Nelken, and King (2011: 43–50).

[2] Ibid., 49.

[3] Pautz (2013: 252); see also Langers et al. (2007); Relkin and Doucet (1997); and Röhl et al. (2011). For a summary of auditory neurophysiology, see Bear, Connors, and Paradiso (2016: 370–403) and Kandel et al. (2013: 654–711).

with which we began.[4] But it is reasonably clear that loudness as we experience it depends crucially on the structure of the neurophysiology of hearing, and information rooted in correlation plays no obvious role in this.

What of pitch? Many sounds have no pitch, and the experience of music may involve different features than that of speech or noise. But still all periodic sounds are experienced to have a pitch.[5]

For pure sinusoidal tones, whose environmental correlates have the form of a single sine wave, increasing frequency of compression waves matches increases in perceived pitch. This hence explains well some relevant similarities and differences of perceived pitch. And all complex sound waves can, through Fourier analysis, be seen as consisting of sinusoidal constituents with various amplitudes and wavelengths. Those that are periodic, which repeat their pattern at regular intervals, also have a perceived pitch. The fundamental frequency of such a sound wave is the greatest common whole number factor of the frequencies of its constituents, and corresponds well with the heard pitch of the sound. The same fundamental will cause the same pitch to be heard even if constituents of a sound wave differ, and a complex sound is heard to have the same pitch as a pure sinusoidal wave that has the same frequency as its fundamental frequency.

But a crucial complication is that even if a sound wave is physically missing its fundamental, such a pitch will be heard. So two complex sound waves may share a perceived pitch, yet in reality have no components of the same frequency if they share such a missing fundamental. Hence, heard pitch cannot be straightforwardly identified with an actual physically objective component of a sound wave.

It is reasonably clear that audition performs some sort of analysis of sound waves not unlike Fourier analysis, which is relevant to heard pitch. And the basilar membrane in the ear does perform a kind of rough Fourier analysis of vibrations of the eardrum caused by sound waves, and converts this into neural effects by rubbing hair cells, which are organized by position on the basilar membrane and hence by frequency, up against the tectorial membrane. There are complications, for instance the fact that mechanical tuning of the basilar membrane is only very broad when loud sounds are involved, but that is roughly what happens.

---

[4] See for instance O'Callaghan (2007: 86–88).

[5] For a sketch of relevant properties of sound waves, see ibid., 76–83.

While hair cells are neurons, they don't fire in the usual fashion, but do form excitatory synaptic connections with "spiral ganglion neurons" along their lower ends. Each hair cell is connected to about ten of these ganglion neurons dedicated to it. Some of the ganglion neurons apparently only fire in response to intense sounds, and hence intense vibrations and hair cell movements, while others require only low intensity at their characteristic frequency to fire. At low frequencies, there is a kind of phase locking of firing peaks in the axons of these ganglion cells with that of relevant sinusoidal sound waves.[6] But the important point is that early activity in audition reflects frequencies of sinusoidal constituents, not experienced pitch itself.

The auditory pathway into the brain is complicated,[7] but this frequency-based organization continues up into the auditory cortex.[8] Where is there neural structure that is sensitive to pitch in particular? That must involve equivalent sensitivity to fundamental frequencies of complex sounds, which may be merely "missing fundamentals" in some cases, and to pure sinusoidal waves of the same frequencies. There is one type of neuron in the ventral cochlear nucleus, whose temporal firing patterns seem to correspond to fundamental frequencies, and there is some evidence of some pitch sensitivity in some neurons in the inferior colliculus,[9] but there isn't much evidence of explicit neural sensitivity to pitch until the auditory cortex, and indeed nonprimary areas of that cortex.[10] Still, there is evidence of the right sort of neural structure at that point. For instance, ferrets are capable of recognizing fundamentals as high or low, and in their auditory cortex, there are populations of neurons in which about a third increase their firing rates with increasing pitch and a third increase their firing rates with decreasing pitch.[11] This is the kind of oppositional structure we would expect, given modal structuralism.

Modal structuralism allows that the experienced range of pitch from low to high could be explained fairly simply in the case of pure sinusoidal tones by the physical range of their audible frequencies, or more exactly by the physical range of frequencies of vibrations of the basilar membrane reflected in the frequency organization of the hair cells. A higher pitched pure sound might sound higher simply because of its higher frequency. This sort of relatively upstream modal structure can be conjoined in a relevant modal filament at some downstream neural level with explicit sensitivity to particular pitches. The fundamental frequency of a complex sound might play a

---

[6] Schnupp, Nelken, and King (2011: 75–86).
[7] Ibid., 86–92.
[8] O'Callaghan (2007: 83).
[9] Schnupp, Nelken, and King (2011: 126).
[10] Ibid., 137.
[11] Schnupp et al. (2010).

similar role, in accord with modal structuralism, even when it is missing. There might be the illusion of such a frequency, but even a merely possible alternative frequency related in the right ways to the actual frequencies of constituents might serve according to modal structuralism.

But there is considerably more structure in our experience of pitch than simply an ordering from low to high. We do not simply hear certain pitches as higher than others. We experience varying sizes of differences in pitch, or intervals. Some intervals between pitches are larger than others. And two different *C*s played on a piano are experienced as similar in pitch in a way not accounted for by their basic linear relations of pitch similarity. In one sense they are more similar than any two pitches except those that are identical. And we hear certain intervals between pitches as more consonant than others. Octaves, for instance, spanning two *C*s on the piano, are maximally consonant. Minor seconds seem much more dissonant.

At least in the case of pure sinusoidal tones, some pitch relations of this sort reflect objective features of frequencies.[12] Small whole-number frequency ratios correspond to the octave (1:2), the fifth (2:3), the fourth (3:4), and so on, with decreasing consonance. This is plausibly phenomenologically relevant according to modal structuralism in the same relatively direct way as frequencies are, through what happens in the inner ear.

But much of this experienced pitch structure does not plausibly correspond to any objective structure of sound waves, and must be contributed by neurophysiology in some other way. For one thing, equal pitch intervals even of pure sounds do not correspond to equal frequency intervals, and indeed apparent pitch is neither a linear nor a logarithmic function of frequency.[13] And it is clear that training and individual history affects some aspects of such experienced structure. The experience of music is a complicated phenomenon that may have an evolutionary history,[14] but training surely matters.[15] Those who are embedded in traditional Western music[16] and its offshoots hear the major third, whose frequency ratio is 4:5, as more consonant than the fourth, contrary to the pattern of objective consonance noted earlier. But there are different scales and intervals deployed in different types of music. All scales incorporate the octave, but even the perfect fifth, perhaps the second most consonant interval in general experience, spanning, for instance, *C* on the piano and

---

[12] O'Callaghan (2007: 77).

[13] Zwicher and Fastl (2006: chapter 5).

[14] Beament (2003).

[15] Bull and Back (2003).

[16] Not very traditional. Familiar triadic harmony, which now rules popular music worldwide, was invented by followers of Monteverdi.

the next higher *G*, is not present in all non-Western scales. Humans generally hear musical notes in relation to a tonic and in relation to some scale into which they are trained. When individuals are trained in different scales and ways of organizing music at the same time, they can hear one pitch inside different relevant possibility spaces. And even those most rooted in traditional Western music can hear a given note as the tonic or dominant of different keys.

Even if we focus for simplicity on the shared experience of pitch of those trained in traditional Western music alone, there are considerable difficulties in understanding the full similarity space of pitch structure. And while it is clear that much of this somewhat artificial structure involves the neurophysiology of audition and nothing objective about sound waves, and in a way that modal structuralism can incorporate into the nature of relevant modal filaments wherever in the brain it occurs, on top of more upstream modal structure closer to the objective structure of sound waves, it is far from clear yet how it all works. Animal experiments are of little help. It isn't even clear that octave structure is present in the experience of most animals. Rats apparently have some sensitivity to octaves,[17] but no other animals have been shown to generalize on the basis of octaves, not even songbirds.[18] This is another reason to think that human neurophysiological idiosyncrasies matter to our sound qualia.

Speech perception is a highly specialized human capacity. And it plausibly introduces distinctive phenomenological features into our auditory experience. When you understand a given language, you experience the relevant sound streams as segmented into phonemes like *b* and *p*, and into words. Your experience involves a certain typing of the sounds you hear that depends on your linguistic training. For instance, English speakers have a hard time hearing tone in Mandarin. And the *t* in "tun" and "stun" sound different to monolingual Chinese. And "bed" and "bad" don't sound distinct to Dutch speakers.[19] Developing children seem to lose early discriminatory capacities and become relatively locked into the phonetic typing distinctive of their own language. And adult perception of some speech sounds is "categorical," with boundaries that are more abrupt in experience than differences in stimuli would underwrite, and which differ among languages.[20]

A relatively austere view of these phenomena is that the distinctive phenomenology of speech is simply due to differences in heard sound of ordinary sorts, say to the distinctive pitch and loudness and temporal

---

[17] Blackwell and Schlosberg (1943). [18] Cynx (1993). [19] O'Callaghan (2015).

[20] Ibid., 482–483; Trout (2001). This even occurs in sign language. See Friederici (2017: 158–159).

patterns of speech relative to other sounds. A less austere view is that some of these phenomenological differences are not differences in such raw sound content, but rather are superadded and relatively abstract features yet present in experience. For instance, one sound might be heard more or less directly as a *b*, and another as a certain word. Some reason to believe the less austere view is that some of the relevant features seem quite distant from particular features of raw sound.[21] Consider how we hear as one letter the *i* in "idiom" and "imbue" even though significant coarticulation effects are present, which affect the sounds in question, and we hear *i* across different speakers with very different voices as yet the same letter. Or perhaps the typing in question isn't specifically auditory at all. For instance, the typing in question sometimes involves cross-modal effects.[22] Consider the McGurk effect.[23] If someone pronounces "ga," but you see them pronounce "ba," you will hear them pronounce "da," a kind of compromise. And Liberman even suggested that the objects of speech perception are not sounds at all but articulatory gestures in the speaker.[24]

It seems to me that the second view is most phenomenologically plausible. But in any case, none of these views allows that the qualia associated with speech perception are fixed by features of environmental stimuli in an information-based way, since speakers of different languages hear the same environmental sound streams differently. That is what matters most from our perspective. It is clear that neurophysiological differences are relevant to the constitution of such qualia.

So then the question is how. And the quick answer is that we don't know much about it. Experiments on language-less animals can't yield much insight into these phenomena, and the kinds of experiments appropriately doable on humans are not that enlightening. One thing that seems clear is that detection of relatively coarse levels of detail in the amplitude- and frequency-modulated sound waves that reach the ear allows decoding of speech. There are also thought to be a hierarchy of levels of neural registration of relevant features, focusing on different timescales. And it seems that the cortex is crucially involved, including the upper part of the temporal lobe and also some frontal and parietal areas especially in the left hemisphere.[25] Without more knowledge of the neurophysiology of speech,

---

[21] Remez and Trout (2009).

[22] There are many cross-modal sensory illusions. See O'Callaghan (2019: 19–25).

[23] McGurk and MacDonald (1976).

[24] Liberman (1996).

[25] Schnupp, Nelken, and King (2011: 139–175). For an introduction to the neurophysiology of language, see Bear, Connors, and Paradiso (2016: 686–717) and Kandel et al. (2013: 1353–1372). For a recent, detailed proposal, see Friederici (2017).

there is little enlightenment to be gained from this case about the details of modal structuralism. But it is clear that contrast effects are important in categorical perception, as modal structuralism predicts. And the layering of auditory typing of different sorts of abstraction in such cases is what modal structuralism interprets as the accumulation of structure along a modal filament.

2

There are three other traditional senses: touch, taste, and smell. Touch is too complicated for any brief discussion. But before we consider taste and smell, it will be instructive to first consider something at least analogous to touch: physical pain.

Among the sensory receptor cells present in the skin and most of the body are nociceptors, whose stimulation characteristically generates the sensation of pain.[26] However, the relationship between nociceptor stimulation and pain is not straightforward, because there can be pain felt when nociceptors are not stimulated, and they can be stimulated when no pain is felt.[27] And while certain representationalists – those who hold that qualia are fixed by the nature of stimuli represented – hold that the damaging stimulations registered by nociceptors constitute the nature of pain qualia, still even most representationalists grant that there is another aspect of pain qualia, constituted in another way, which points to the relevance of neural resources of other sorts. Motivation is also involved.

Perhaps the paradigmatic example of such a view is Michael Tye's.[28] His idea is roughly that "a pain in the leg ... is a token sensory experience which represents that something in the leg is damaged, something moreover that is painful or hurts."[29] More exactly, the idea is that "a pain in the leg is a pain that represents accurately if, and only if, it is caused by damage in the leg and causes the immediate desire that it cease, together with anxiety about the relevant portion of the leg."[30] In his "view, it is the representation of a certain class of observational features by our pain experiences, and the role that they play, which gives pains their phenomenal character."[31] But notice that in Tye's proposal, there are two key clauses: There is an output element having to do with the motivational role of pain and an input element having to do with the representation of

---

[26] Bear, Connors, and Paradiso (2016: 437–448); Kandel et al. (2013: 530–555).

[27] Bear, Connors, and Paradiso (2016: 437–438).

[28] Tye (1995a).

[29] Ibid., 228.

[30] Ibid., punctuation as in original.

[31] Ibid., 236.

damage, so even Tye's representationalist view is in fact a mixed input-output view of pain.

This suggests that such qualia are constituted by certain sorts of neurophysiological structure running from the stimulation of nociceptors towards motor activity. In particular, it seems plausible that pain qualia crucially involve in paradigmatic cases not only nociceptor stimulation but also the tendency of their sufferer away from the relevant stimulation, that the nature of the qualia is partly determined in an outgoing way, that the painfulness of painful qualia is in part due to how they push us. More neurophysiology is certainly involved in that.

Nociceptors are the endings of the thinnest nerve fibers, A-delta and C-fibers, which have the slowest conduction velocity of all nerve fibers.[32] And there are basically two kinds. Some respond to high-intensity physical stimulation, but others are polymodal: They respond to several forms of high-intensity stimulation, including low and high temperatures, high-intensity mechanical energy, and pain-inducing chemicals like capsaicin as found in hot chili peppers. But these nerve fibers are associated also with warmth and cold sensations and those of pleasure and itch, so something complex is going on.

They first synapse in the dorsal horn of the spinal cord, which has a very messy organization but probably involves both wide dynamic-range neurons that are responsive to various different forms of sensory input – to touch, pressure, hair movements, temperature, and injury – and also those that are nociceptor-specific in their response. Then there is transmission of pain registration up the spinal cord in a somewhat dedicated way, although temperature sensitivity is also transmitted in the same pathways, and there is recovery in pain sensitivity if these pathways are cut. There can also be sensitization, which is to say enhancement of pain stimuli, at various places in this chain. So even on the input side things are complex. The sensory aspects of pain seem at least in part associated with sensory cortex, farther up the neural trail than nociceptors. But it is most crucial to notice that that is not all there is to pain qualia. The also crucial emotional and affective aspects of pain seem to involve the brainstem, hypothalamus, amygdala, and anterior cingulate nucleus. So the neural basis of pain qualia seems to involve different sorts of neurons in widely distributed regions of the nervous system.

The spatial character of the relevant qualia may be crucially due to the nature and locations of the nociceptors (although the phenomenon of

---

[32] Cervero (2012: 43).

referred pain is a complexity worth nothing), but the awfulness of the pain qualia apparently has much to do with something closer to motor activity. So once again modal filaments seem in play. Since physical pain can be felt when nociceptors register no damage, and indeed when nociceptors are not even stimulated, it would appear that the representationalist elements of accounts like Tye's play no necessary role in some pains. On the other hand, so-called morphine pain, pain that doesn't hurt, may involve nociceptor activation alone. In either case, neural structure active in various ways seems to be all that is required, at various locations along modal filaments.

3

Smell and taste are entwined. Begin with the characteristic intrinsic properties present in our smell experience, and how they reflect relevant neurophysiology.

When a human smells things, volatile molecules – gaseous emanations of source objects which are called "odorants" – encounter a small, thin layer of cells high inside the nostril cavity.[33] This sheet, called "the olfactory epithelium," contains olfactory receptor cells. There is a vast complexity of different kinds of such receptor neurons in your nose. They all end in long, thin cilia, but differ in regard to the receptor proteins that they contain, which determine what odorants will bind with the cilia and induce changes in potential in the neurons. Estimates stretch from about 300 to 400 different types in a given person.[34] There is also considerable genetic investment in this diversity. There is a large gene family that encodes olfactory receptors, the so called G-protein receptors, and consists of about 1,000 different genes that generate an equivalent number of receptor types in different species. There are roughly 300 functional genes in humans in this category.[35] And there is also considerable individual variation among humans in regard to olfactory receptors, which seems to reflect genetic variability among us. For instance, when 500 people were asked "to rate a panel of 66 odors for intensity and pleasantness, they gave the full range of responses – from weak to intense and from pleasant to unpleasant."[36] And when 1,500 young adults were

---

[33] Bear, Connors, and Paradiso (2016: 278–291); Kandel et al. (2013: 712–726).

[34] Batty (2014: 223); Young, Keller, and Rosenthal (2014).

[35] Batty (2014: 223 fn 7).

[36] Spinney (2011).

tested on sensitivity to 20 odors, there were found to be specific insensitivities to all but one of them.[37]

Despite the large number of different sorts of olfactory receptor cells we have, we can distinguish many more odors than that. This is possible because many olfactory receptors are broadly tuned to a range of different chemical stimuli, and it is combinations of activations of receptor types that correspond to different odors. This is called "population coding." Population coding seems to take a relevantly nonadditive form in olfaction. Most odors are blends, which involve many different types of molecules. But there is evidence that some of these molecules inhibit or enhance the effects of others on receptors.[38] And there is evidence that ratios of activation of different sorts of receptors matter to odor experience.[39] The olfactory experience of a blend is not a simple sum of the olfactory experience of its components.

Olfactory receptor neurons synapse on two "olfactory bulbs," which have considerable neural complexity, and then neural stimuli pass on to several areas of the brain, including the olfactory cortex. This rather direct connection to many parts of the brain is different from other senses. There seems to be correlation between spatial patterns of activation in the olfactory bulbs and particular experienced odors, and also between particular odors and particular groups of neurons in the olfactory cortex.[40] But the way in which similarity of odor experience is rooted in similarity of neural activation is at least somewhat different from the way it happens in vision. The olfactory bulbs have significant structure, but its relation to phenomenal structure seems different from that of the locally oppositional structure of the retina or even the primary visual cortex. In the olfactory bulbs, there are spherical neural structures called "glomeruli," on which neurons from the epithelium synapse, and which synapse on neurons that pass on to the olfactory cortex. They involve considerable internal neural complexity, and interact within the bulbs in complex and functionally relevant ways.[41] Each glomerulus is tied to a certain type of sensory receptor cell, but there are some differences across individuals regarding where such a particular glomerulus is placed relative to others,[42] and there is no known relation between similarity in patterns of glomeruli activation and phenomenal similarity.

---

[37] Ibid. [38] Barwich (2020: 190–194 and 196–197). [39] Ibid., 199–200.

[40] Bear, Connors, and Paradiso (2016: 284–290). [41] Barwich (2020: 73–74 and 208–211).

[42] Ibid., 218–232.

Still, this complex neurophysiology at least suggests the serious complexity of the intrinsic properties present in olfactory experience. And they certainly exhibit that.

There are obvious differences in the hedonicity and intensity of smells, but also many different differences in smell quality beyond that. There have historically been many varied systems for classifying basic olfactory properties, which reflect this complexity in different ways. Linnaeus and Zwaardemaker's well-known schemes lumped basic smells into distinct determinable categories. Zwaardemaker's categories, which incorporate Linnaeus', were aromatic, fragrant, ambrosial (or musk-like), alliaceous (garlic-like), hircine (goat-like), foul, nauseating, ethereal (like beeswax), and empyreumatic (burnt).[43] There were also alternative quality-space models, resembling the customary model of experienced colors, some with primary smells that are analogous to unique colors but all with various sorts of similarity relations among determinate smells. Henning's prominent system proposed six primary odors – fragrant, fruity, resinous, spicy, putrid, and burned – occupying the vertices of a quality-space similarity prism, with other odors located on its surface.[44] The difference from the first type of scheme is that in such quality-space accounts, all basic smells have some sort of similarity to others. And they may also allow that within each rough determinable category, there are central or primary smells. There might also be mixed models, with distinct categories that yet include central and paradigmatic smells.

It seems to me that the phenomenological case for such primary smells is reasonably, though not extremely, strong, and that all intrinsic basic smells have at least a kind of similarity to each other that they all lack to colors or sounds, so the quality-space models seem phenomenologically preferable. On the other hand, it is not phenomenologically implausible to claim that the primary smells are not as determinately primary as the unique hues are unique. There are also other sorts of models in which olfactory experience is not analytic in either of these ways, in which instead the various properties of an odor stimulus together generate a unique and irreducible experience with no introspectibly distinguishable features.[45] For instance, it might be that when coffee is smelled, and 800 chemically different compounds consequently enter the nose, all the 300 to 400 types of smell receptors are activated in a way that generates a recognizable pattern to which the nervous system has been trained, and also generates

---

[43] Batty (2014: 224). [44] Ibid., 226 fn 17.

[45] Batty (2014: 227); Wilson and Stevenson (2006).

a single unique sort of quale. But it seems phenomenologically apparent that there are some relations of similarity and difference among smell qualia that must be captured by any plausible view. So this third model at the very least requires supplementation, perhaps by a considerably complex basic quality-space model, which presents qualities that are to the special smell of coffee as pitch and loudness are to the novel qualia presented by understood speech.

It is much more difficult to approach these phenomenological issues empirically than in the case of color, for instance by collecting judgments of just noticeable differences among smells, because of the great complexity involved. But there is reasonably strong experimental evidence for phenomenal complexity characteristic of the quality-space models, if not perhaps for their focus on primary smells.[46] One recent conclusion is that the color quality-space (QS) "arranges a few million qualities in a three-dimensional space whereas the olfactory QS arranges vastly more qualities in a substantially higher-dimensional space. The numerous failures to describe an olfactory QS are merely due to the extremely large dataset required to do so."[47] But perhaps recent techniques provide room for greater determinacy. For instance, there is a quite recent proposal of a 3D odor cube closely related to Henning's odor prism with considerable experimental support.

Let's move now beyond phenomenology to its explanation. There is considerable reason to believe that many differences and similarity in smell qualia are not due to physical differences and similarities in physical odors, as representationalists believe, but are rather due, as modal structuralism implies, to features of our neurophysiology. Begin with the first point.

We've already noted that there are considerable individual differences in what humans experience when smelling given physical odors. And there is strong reason to believe that species differences matter also in this way. Some things smell nauseous or disgusting, in a way that suggests involvement of motivational structures in relevant modal filaments, as with pain. And some of the same things out in the world evidently smell disgusting to some animals but not to others. Representational accounts of that aspect of odor qualia are not very plausible.

But what of the other features of odor qualia? There has been recent general agreement that the similarities of smell qualia do not correspond well to physical similarities of source emanations:

---

[46] For some of the complexity, see Turin and Sanchez (2008).

[47] Young, Keller, and Rosenthal (2014: 8).

> Available evidence indicates that numerous chemical and molecular features . . . can all influence the odorous characteristics of a molecule. However, no systematic description of how these characteristics relate to particular odor qualities has been developed. In other words, chemicals that bear little resemblance structurally can smell the same, and chemicals that are nearly identical structurally can elicit very different perceptual qualities.[48]

And receptors do not respond to intuitive chemical features of odorants, so it has been impossible to predict from studying molecules, independent of knowledge of the structure of sensory receptors, how they will be experienced to smell.[49] But it is possible that on this front there is promising progress. Kepple and Koulakov report that positions of molecules in a certain 10D space of molecule similarity of a new and less intuitive type do allow predictions of perceived odor in a 6D space of odor similarity, in which pleasantness is an important axis.[50]

This development might seem to support an information-based account of such odor qualia. But note the importance of pleasantness in the 6D odor space. And no more than in the case of vision, in which variable activity of sensory receptors plays a role according to modal structuralism, is the activity of sensory receptors in smell ruled out of consideration by modal structuralism. What is ruled out is the immediate relevance of any differential features of external stimuli to qualia that are not reflected in differences in sensory receptor activity. And what the 10D space of molecules seems to promise instead is a more detailed understanding of relevant sensory receptor activity. What's more, we have already noted the significance of ratios of molecules in blends, and also population coding involving both enhancement and inhibition in smell experience, which affect how a given molecule will contribute to an overall smell. There is also significant adaptation in smell experience, so that for instance, over time, the characteristic smell of a room you have entered disappears. And in a blend this may interact with the mechanisms we've already noted to reveal over time different "notes" in the blend,[51] so that over time it smells differently. Sniffing of different strengths and patterns and lengths also affects what molecules reach the olfactory epithelium when and so affect smell.[52] And differences in concentration of a molecule or blend generated in this way can also affect odor not just in intensity.[53] Also, even the verbal

---

[48] Cowart and Rawson (2001); Pautz (2013).

[49] Barwich (2020: 182–186).

[50] Kepple and Koulakov (2017). See also Koulakov et al., (2011).

[51] Barwich (2020: 198–199).

[52] Ibid., 152–154.

[53] Ibid., 110–112.

descriptions of odors, say as emanating from vomit or Parmesan cheese, can affect how they smell, for instance how pleasant they are.[54] So it is quite implausible to hold that odor similarity in general is determined in a representationalist fashion.

There is, however, considerable evidence that experienced smell similarity does correspond well to neurophysiological structure. Subjective similarity space in olfaction

> maps simply onto the multidimensional space represented by neuronal firing in primate cortical taste areas. In particular, if a three-dimensional space reflecting the distances between the representations of different tastes provided by macaque neurons in the cortical taste areas is constructed, then the distances between the subjective ratings by humans of different tastes is very similar …. [And] the changes in human subjective ratings of the pleasantness of the taste, smell, and sight of food parallel very closely the responses of neurons in the macaque orbitofrontal cortex.[55]

And in a study of chemicals that are very physically different but smell similarly, "spatially distributed ensemble activity in human posterior piriform cortex (PCC) coincides with perceptual ratings of odor quality, such that odorants with more (or less) similar fMRI patterns were perceived as more (or less) alike."[56]

This is not to say that the way in which neurophysiology is involved in olfactory experience is not different from other cases we have considered. Smell is structurally shallow, in the sense that relevant cortical and memory synapses are only two synapses away from the sensory receptors.[57] This structural shallowness is related to another distinctive feature of odor perception, the Proust phenomenon, as when Marcel's petite madeleine dipped in tea recalled a whole past world to him.

> Smell has a special capacity to trigger episodic memories—especially emotionally charged memories—because of the direct connections olfaction has with parts of the limbic system involved in emotion and memory. Olfaction is unique in being the only sensory system that projects directly to the amygdala without going via the thalamus. As a result, odour memories are long-lasting and less susceptible to interference.[58]

The complexities are vast, and not well understood. But there are weighty grounds to believe in neural involvement in constituting the experience of intrinsic smell qualia. It is also reasonably clear from the

---

[54] Ibid., 264–266.

[55] Rolls (2007: 843–844).

[56] J. D. Howard et al. (2009: 932); Pautz (2013).

[57] Laurent (1999: 727).

[58] B. Smith (2015: 327–328).

existence of nauseating and delightful smells that there are motivational structures sometimes involved in the relevant modal filaments, as in the case of pain. Modal structuralism seems well positioned to capture all the relevant complexities, while information-based accounts are not.

4

Now consider taste.

It is important to note that olfaction affects the taste of foods. When we have colds, or when our sense of smell is otherwise disabled, food can taste like cardboard. For this reason, it is useful to regiment terminology a bit for clarity's sake. "Flavor" is the general umbrella term. Differences in experienced taste are differences in flavor, however they arise. "Gustation" is taste proper.

Gustation involves specific taste receptors in the tongue, and corresponding primary tastes. The four traditional primary tastes are bitter (as in quinine), sour (as in a lemon), sweet, and salty. But in recent discussions, they have been frequently augmented with savory (or umami), which is a common factor in our experience of tomatoes, mushrooms, peas, parmesan cheese, and seafood.[59] Galen included harsh, pungent, and astringent.[60] And there are some recent suggestions that metallic and fatty may be basic tastes.[61] Each primary taste corresponds to the excitation of dedicated receptor cells, although there are also indications of more than one type of sensory receptor cell for some primary tastes, for instance bitter.[62]

The receptor cells for taste are not neurons by standard criteria, but they synapse on gustatory afferent neurons and other cells and, when activated by an appropriate tasty chemical, undergo a change in electrical potential that may in turn generate an action potential as in a neuron, or otherwise cause them to release neurotransmitters.[63] Sour and salty taste receptors release serotonin, and the other three known types of receptors release adenosine triphosphate, as their primary neurotransmitter. There is also a difference in the way these two classes of taste receptors function. Sour and salty tastants interact directly with ion channels on the relevant receptors by either passing through them or blocking them, whereas bitterness,

---

[59] Ibid., 331. [60] Bartoshuk (1978). [61] B. Smith (2015: 330).

[62] For neurophysiological details, see Bradbury (2004); Chandrashekar et al. (2006); and Kandel et al. (2013: 726–735).

[63] Bear, Connors, and Paradiso (2016: 266–278).

sweetness, and umami involve binding to taste receptor proteins, more or less as in olfaction.

Unlike with color primaries, it has generally been assumed that experience of complex gustations simply involves a combination of different intensities of these basic tastes, that the primary tastes correspond to distinct modalities. But some have held that there are intermediate tastes, analogous to orange or turquoise. Henning's taste tetrahedron was supposed to capture the similarity of distinct basic tastes located across its surface, with the four traditional primary tastes located at its vertices.

In any case, it is clear that all the phenomenal complexity of taste is not captured by these aspects due to gustation. There are at least two other sensory systems that contribute to the experience of taste. Most surprising, touch is involved. The spices that make mustard feel hot and peppermint feel cool work by stimulating the trigeminal nerve.[64] This is like the chemical skin senses, which can register chemical burning on the skin. And of course, as already noted, olfaction plays a large role in taste experience, in the qualia of flavors.[65] There is even evidence that the primary gustatory tastes differ in intensity if olfaction is disabled.[66]

The same olfactory receptor cells are involved in the experience of smells and some aspects of the experience of taste. But the olfactory system doesn't function in exactly the same way in both cases, so there is even a doctrine of the duality of olfaction, which suggests that it constitutes two senses, one relevant to smell proper and the other to taste, despite involving only one set of sensory receptors.[67] This is a very dramatic example of the way context matters to our experience of odors.

How does this work? Vapors pass by two routes to the smell receptors. The orthonasal route is through the nose, as when we inhale. The retronasal route is via the mouth, as when we chew. Our earlier discussion of smell really just concerned orthonasal olfaction. And retronasal olfaction is what is crucial to our experience of flavors and the phenomenology of taste. Clearly there is some phenomenal similarity between an odor sensed by either route, but there is definitely not a phenomenal identity. For one thing, retronasal olfaction involves a kind of location illusion called "oral referral," in that the sensory property in such cases, received through stimulation in the nose, is yet projected in experience onto a position in the mouth, into the chewed food.[68] There can be differences in the intrinsic phenomenal character of the same odor when it is experienced

---

[64] B. Smith (2015: 316). [65] For various cross-modal effects in taste, see B. Smith (2020).
[66] Mojet, Köster, and Prinz (2005). [67] B. Smith (2015: 324). [68] Ibid.

through the two different routes. Cheese that smells bad can taste good. Fresh coffee that smells great can taste merely fair. Cigars can be noxious when smoked by yourself but pleasant when smoked by others.[69] Eating and chewing are exploratory activities relevant to taste, but in this way they are even relevant to the experience of odors.

One question about all this is whether the various factors that go to make up the experience of flavor are distinguishable by those who experience them, at least after training. And the answer seems to be largely no.[70] Even the experience of sweet doesn't seem to involve just gustation.

Clearly the experience of flavor is a very complex phenomenon. And it is complex in ways that show that, despite the straightforward correspondence between distinct sensory receptors and the primary gustatory tastes, neurophysiological structure and not environmental properties are what is crucial to the constitution of taste qualia, in accord with modal structuralism.

There is the way that olfaction helps determine taste, and indeed how different olfactory routes may present one stimulus in different ways. There is the fact that some things taste disgusting to one person or animal and not another. There is the phenomenon of taste contrast, experienced by drinking the bourbon shot before or after the beer. There is the fact that supertasters, who are very sensitive to gustatory qualities, make poor cooks because they prefer very bland food.There is also phenomenal difference despite close chemical similarity in things tasted. Aspartame and its stereoisomer are very chemically similar; they are compounds that differ only in the orientation of two hydrogen atoms. But they taste very different, respectively very sweet and very bitter.[71] Gentiobiose is also bitter, but its very similar stereoisomer – isomaltose – is sweet.[72] And there are examples of phenomenal similarity despite extreme chemical difference. For instance, bitter-tasting things include moderately large organic compounds, large organic acids, small molecules like urea, and some sugars.[73]

Differences and similarities in experienced taste correspond much better to neurophysiological differences and similarities, to difference and similarities among firing patterns involving ensembles of neurons in relatively central parts of the taste system.[74] There is evidence of spatial differences between the neurons active in sweet and bitter experience, and in sour and salty experience.[75] There is evidence that temporal differences in firing

---

[69] Ibid., 327–328.
[70] Spence, Auvray, and Smith (2014).
[71] Pautz (2013); Walters (1996).
[72] Sakurai et al. (2010).
[73] Pautz (2013).
[74] Ibid. collects relevant evidence.
[75] D. Smith et al. (1983).

patterns of certain single neurons correspond to differences in phenomenal taste space.[76] And there is evidence of close correlation between firing rates and taste intensity.[77] It is also plausible that population coding for the specific tastes of specific substances occurs, as in the case of olfaction.[78]

As modal structuralism predicts, the similarities and differences of the qualia presented in our sense of taste have much to do with our neurophysiology of taste and smell, and little if anything to do with the physical nature of the objects out in reality we thereby taste. And the pleasantness and unpleasantness of some tastes suggests the relevance of neural resources distributed along extended modal filaments.

---

[76] Di Lorenzo et al. (2009).

[77] Borg et al. (1967).

[78] Bear, Connors, and Paradiso (2016: 277–278).

CHAPTER 6

Visual Space

Colors present themselves in our conscious sensory experience as in various places and on various objects. This chapter begins our exploration of this complex topic, which will continue in Chapter 7. Here we consider relatively simple aspects of visual space, and associated qualia.

I

Human vision is mediated by the stimulation of photoreceptor cells.[1] And it is natural to think there is rudimentary spatial information present from the beginning in vision, because the rods and cones are arranged spatially in the retina of the eye, which receives light from the world focused in a spatial image, although upside down and reversed. But it is a difficult question how we get from stimulation by photons of spatially arranged photoreceptors to our visual experience, with objects perceived in depth and with various spatial relations, and with edges as well as colors. The standard account is an input-driven, representationalist story, according to which information is processed in a hierarchical way, such that various presumptions are added at each stage of processing, like additional premises. There is supposed to be a hierarchy of levels of unconscious representational states in our neurophysiology, that successively represent more and more complex forms of spatial structure. The theory of David Marr's *Vision* is iconic,[2] and while the details of his account are no longer accepted, the general shape of that account remains popular.[3]

While we experience three-dimensional objects like cars and tables with a robust sort of particularity, existing in an objective three-dimensional space, it is generally agreed that there is also an image-based stage of visual processing, which involves a two-dimensional geometry centered on the

---

[1] Groh (2014: 14–16). [2] Marr (1982). [3] Palmer (1999: 145–461).

https://doi.org/10.1017/9781009514774.006

perceiver, and a distribution of shapes and lines and colors over that roughly flat image of the world.[4] This is often thought to be unconscious.

Spatial structure in that vicinity is our focus in this chapter. Still, on the modal structural conception, information-processing per se is not what is relevant to phenomenal consciousness. Nevertheless, the neurophysiological structures treated in the standard story as early unconscious stages of visual processing remain relevant to the overall modal structural explanation of our visual consciousness. Work on visual perception that discusses unconscious processing can be refocused for our purposes.

Indeed the content of what is the image-based stage according to the standard information-based story is plausibly available in consciousness. Our vision can present us, at least when one eye is closed, with a flat arrangement of colors and shapes, which is a visible image of the world. We sometimes see colors arranged in a flat spatial array. You can move about until the right side of the window you are looking out of seems to correspond with the edge of that building hundreds of feet away, so that the wall of the building is visually adjacent to that of your room. And you can compare their colors at their visual boundary, although of course they don't really touch. You can take your visual experience in that painterly way.

It has been argued that, in the history of vision science and philosophy, the idea that such visual content is conscious only entered with Descartes and his general commitment to the consciousness of all truly mental states, that it was not embraced originally for purely phenomenological reasons.[5] And the history of art does not reveal an ancient and constant commitment to single-point perspective or analogues. And Gibsonians reject the notion that two-dimensional visions are perceptually primary. But all I claim here is that two-dimensional as well as three-dimensional visual structures are available in phenomenal consciousness. That is not inconsistent with any of these other claims. And it does seem hard to deny that humans have long been at least *capable* of the recognition that, for instance, someone's hand looks to them in one sense smaller as they move it out from the immediate front of their face to arm's length, so that it may be at first visually larger in one sense and then visually smaller than some stationary barn, which it only at first occludes.

This is in some ways a nonveridical aspect of visual experience, which hence requires that relevant qualia be present in your visual consciousness. Of course, the state you are in when you examine the color-matching of

---

[4] Ibid., 87–88.

[5] Hatfield and Epstein (1979).

building and window frame at the window's edge is somewhat artificial. But it is phenomenologically clear that even in normal perception, when you are focused intently on the distant three-dimensional building itself and its three-dimensional shape and color, still the image-like two-dimensional arrangement is present in your visual experience and consciousness.

Conscious visual qualia of that relatively simple spatial sort are our focus here. I will call that structure "rudimentary" because it is relatively simple, but that is all I mean by that. I don't claim it is primary in some controversial sense.

It may be puzzling how we can have at once a conscious visual experience of, on one hand, a three-dimensional red building in the distance and the occluding yellow edge of a window quite close by us, and also, on the other hand, a conscious visual experience of a flat visual field in which the red and yellow are touching. But we have already noted that in our visual experience, a bit of color can at once present itself as a color look of a certain specific sort, but also as the surface color of an opaque object that it is understood will present different color looks as we move or illumination changes. There are two sorts of modal structure that overlap in the experienced color, in a kind of palimpsest, one more rudimentary than the other. So too, flat visual space is more rudimentary than complex three-dimensional visual space, but they overlap in a kind of palimpsest in our experience.

The kind of particularity that is ordinarily dominantly present in your conscious visual experience, for instance when you see a building off in the distance, is sophisticated and complex. We will return to that kind of particularity in Chapter 7. But for now we aren't trying to go all the way to that sophisticated form of visual particularity.

More immediately relevant is a kind of particularity analogous to that of spatial points. You can see colors to be present at various locations of flat visual space. This space is two-dimensional rather than three-dimensional, at least to a first approximation. Also, unlike any real space, it is finite. It has a limited fineness of detail. It is centered and oriented in a certain way on you, or perhaps your head or your eyes. And it has more detail in the center than the edges.

We are interested here in modal filaments that are individual color quale in a quite specific sense. They present certain phenomenal colors, but are also experienced as minimal units in that flat visual space. These minimal spatial units are apparently individuated largely relationally, such that this one is near that one and farther from that one. We are interested here in why we experience the apparent spatial structure and resulting particularity of such color quale in flat visual space.

Phenomenal colors and that rudimentary visual field seem made for each other. Just as the range of possible colors seems to be rooted in any instance of any color, so the instance of any color in such a field seems to assure the possible instance of any other at the same position, and each position in the two-dimensional visual field seems to require some sort of filling up with a color. And the phenomenal colors we experience are suited to such a field, whatever other roles they may play. If such a field is one-sided and oriented to the viewer, so too can they be.

But there is an immediate complication. There is a kind of particularity beyond the relational sort characteristic of spatial points, which is present even in the rudimentary flat visual experience that is our focus here. We will eventually see that the visual particularity of three-dimensional objects involves not just spatial but temporal phenomena in crucial ways. But even the two-dimensional space of colored shapes reveals a kind of particularity involving time and change of location.

You can watch two-dimensional color shapes in the visual field move around, hence changing locations in that two-dimensional space. At each moment they consist of various specific minimal bits at specific locations. But they are experienced to have a kind of particularity that does not depend on their locations, in the sense that their locations change. You can even experience these color patches changing shape and size and color. And if at least one of them is suitably transparent, you can watch them move through each other, and hence occupy the same place in flat visual space at the same time.[6]

This more complex particularity clearly depends on the more rudimentary particularity of positions.[7] As a patch of one color moves, its color successively occupies different locations. But the particularity of the patch doesn't seem to be constituted even by that of the locations over time that it occupies, since it might have moved elsewhere than it does and yet remained in experience the same patch.

The modal structural analysis of this phenomenon requires a different sort of modal filament than that which constitutes the experience of a color

---

[6] Pylyshyn (2001). Contrary to Pylyshyn, we can watch many visual patches at once possess this sort of particularity without being able to track all of their identifies over time in an attended way, and it need not be actual objects which we track in an attended way. A case in which something is transparent and something is opaque, so you see one behind the other in flat visual space, is more complex.

[7] This response is developed by Austen Clark (2004). But I myself don't assume that the relevant spatial locations are real.

as at a certain fixed visual location.[8] Even in that second, stable case, there is a kind of complex constitution of one continuing qualia and modal filament by a series of momentary qualia and modal filaments, with other phenomenological structure plausibly required to present the apparent stability of position over time. But consider a simple contrasting case, the perceived steady motion of one minimally discernible bit of color. If it is to be considered a single color quale, it surely consists of a color quale with different individuation conditions than a visibly stationary color quale. It moves in visible space, and could move other than how it does. It hence is constituted by a different type of modal filament, surrounded by a modal environment allowing its range of possible motion. If the simpler kind of immobile color qualia involves what might be called a straight modal filament, surrounded only by alternative possible positions that it itself does not and cannot occupy, then this somewhat more complex sort of color quale involves what might be called a flexible modal filament, which projects the quale at any moment from one location to its next, and indeed which could project it into other locations it will not in fact occupy.

It is useful to consider this case in the context of other cases we have already considered. Already, in the palimpsest of modal structure involved in experiencing a specific color quale as an aspect of an objective color, and experiencing visual spatial structure as at once two-dimensional and three-dimensional, we have considered qualia that are in a sense one and in another sense two. And if the entire nervous system mediating stimulation and action constitutes phenomenal consciousness, we can expect much of what we do on the basis of a certain stimulation to underwrite forms of phenomenal consciousness beyond any intuitively sensory aspect of it, so that such palimpsest structure is quite common. I think in such cases that more than one modal filament is involved. For instance, in an aspect switch when a visual field you are dominantly considering three-dimensional becomes dominantly two-dimensional, it is plausible that there is transformation between one modal filament and another, but such that they both share a more truncated modal filament as a part in which the aspect in question is not resolved one way or the other. Of course, this will be a modal part and so not necessarily a spatial part. The sensory qualia in which we are centrally interested here are often truncated in that way, leading not to particular actions but to particular possibilities of

---

[8] Neither of these phenomena presumes the index or file characteristically deployed by representationalists.

reaction and action, although sometimes tendencies to at least specific types of actions may be a crucial part of the relevant qualia at even the most fundamental and intuitively sensory level.

Now return to our current concern. We will be most centrally interested here in a very simple case, a momentary color quale that presents itself as at a single location at a single moment, what might be called a momentary modal filament. But that there are these various different types of visual filaments – momentary, immobile, and flexible – may present a puzzle. How can there be more than one color quale occupying the same visible location at any given time – one that allows movement over time, one that does not, and one that is momentary? It is because of the same general sort of palimpsest of modal structure as that involved in experiencing a color bit as at once part of a visible field and as the look of a color of an external object. It is a simpler example of that same phenomenon.[9]

So far this idea of flexible and straight and momentary modal filaments is only a rough metaphor. We will work towards fleshing out somewhat better the notion of an experience of somewhat complex forms of particularity like that of moving shapes in Chapter 7, although we will focus there on a case more complex than that of shapes moving in flat visual space. But the first step in this explanation requires a better understanding of the simpler notion of a momentary modal filament, constrained to a location in visual space at a time. That is our next task. Even this case can plausibly involve the composition of modal filaments by others, as the experience of an extended part of the visual field involves at least in part experience also of its bits. But as I said, we are focusing on the spatial bits. Our primary concern in this chapter will be experience of the momentary spatial bits of the two-dimensional visual field, although some complexities will require brief forays beyond this very rudimentary case.

2

Our earlier discussion of color focused on ways in which phenomenal color is not found in reality, and so cannot be delivered by representationalist accounts of sensory content. We will shortly see that some of the detailed spatial features of our visual experience are also of this nature. But the most general and basic features of the spatial aspect of that experience plausibly

---

[9] This suggests that the color bit conceived as in the flat visual field and as on the object should be considered different qualia and modal filaments.

reflect objective spatial reality. Or at least any reasons to doubt that are difficult complexities of physics that won't concern us here.

The focusing of light on the retina of each eye allows a match between the apparent spatial structure of our two-dimensional visual experience and real spatial structure in the eyes, and indeed ultimately real spatial structure in the environment. The cones and rods are arranged spatially on the retina, and their pattern of activity under stimulation is objectively spatial, and matches, though in an intuitively upside-down and reversed way, what we experience, as well as the nature of environmental stimuli. What's more, our motor neurons have a crucial objective spatial arrangement as well, controlling muscles attached to a skeleton that have their own crucial spatial arrangement, although in the case of the motor neurons this is a more than two-dimensional arrangement.

While we will see that some of the spatial structure of even rudimentary and flat visual experience requires a modal structural analysis in much the same manner as nonveridical color vision, still the most basic spatial structure of that experience, its rough two-dimensional structure, involving various locations between other locations along various possibly somewhat curvy lines, what might be somewhat loosely called the topology of that experience, reflects reasonably well the structure present objectively in the spatial arrangement of our rods and cones. And so an information-based account of that aspect of visual experience cannot be automatically dismissed.

Nevertheless, there are strong reasons for a modal structural analysis of even that feature of visual experience. They involve the way in which even the basic spatial structure of our two-dimensional visual experience requires a kind of connection between the two-dimensional spatial structure of the rods and cones and the spatial structure of our motor activity. The rudimentary topological structure of our visual experience isn't in fact due to information of any intuitive sort, but it does involve a correlation in this other way.

On the conception we will be developing, the experience of a color quale as at some position in two-dimensional visual space is a modal filament that connects spatially located cones and rods to spatially located motor activity. There is a correlation among two alternative sets of possibilities, locations of visual stimulation and locations for action, and a given filament, to a first approximation, connects specific locations of those two sorts. When you see a color at a certain location in your visual field, you can move in ways to indicate that position, for instance by putting a finger or thumb up to occlude it, or by moving in its apparent direction, or by ducking or turning away. Of course there is no particular bodily action that is an essential part of the experience of a color quale, but such an experience involves the capacity to coordinate a wide range of possible actions sharing spatial features with its visual location. There is a long tradition in philosophy, stemming from at least Locke, of maintaining that depth in visual experience involves an interaction of vision and action. But what is in question here is the role that such interaction plays in constituting the experience of even two-dimensional visual locations. Berkeley also developed one version of this idea.[10] But my version is somewhat different.

There are two clusters of reasons to believe there is such a role for the interaction of visual stimulation and motor action.

The first involves inverting lenses.[11] People can wear goggles with inverting prisms, lenses, or mirrors that cause light to be received at the retina in a way that is upside down, or left-right reversed, or both, relative to the way in which it would have been received without the glasses. The spatial structure constituted by activity of the relevant sensory receptor cells is hence also upside down, left-right reversed, or both. If the spatial structure of the cones and rods constituted in itself the two-dimensional structure of visual space, one would, when wearing such glasses, see the world upside down or left-right reversed or both, immediately and continuously. But that is not what happens.

Rather, at first, vision is seriously disrupted, presenting no coherent space. Then adaption occurs in stages.[12] In the first stage, visual experience is inverted, so there is conflict between visual experience and other sorts of spatial experience, for instance between the location at which something is seen and at which it is heard or felt. Then the spatial structure of visual experience becomes dominant, so sounds are heard at the false places where their sources are seen, much as in ventriloquism. But if the person wearing the goggles moves around and acts in the world, then eventually their experience becomes normal again, even though they have the inverting lenses on. There is an important complexity, however: The reinversion occurs in a piecemeal fashion, in different realms of activity as different sorts of activities are engaged in.[13] Finally, when the glasses are

---

[10] Berkeley (1710). But for difficulties, see Grush (2007a).

[11] Noë (2006: 7–11, 91–96); O'Regan (2011: 34–41). My treatment of inverting lenses is as such enactivists suggest. For an alternative, see Rock (1966). For the original experiments, see Stratton (1896, 1897). For somewhat more recent experiments, see Kohler (1964) and Taylor (1962).

[12] This summary of stages is due to Noë and based on Kohler and Taylor.

[13] There is also evidence for adaptation in the absence of action. See Howard, Craske, and Templeton (1965). Harris (1965) suggests that it is not visual space but proprioceptive space that inverts in the

removed, there is once again visual chaos until, more quickly, readaptation occurs, and things return to normal.

The natural explanation of these phenomena is that a certain sort of connection between the incoming spatial structure of sensory receptors and the outgoing spatial structure of motor activity is constitutive of locations in visual space. The ordinary connection of that sort is disrupted when inverting goggles are first put on, and so is visual space. But eventually an unusual and indeed inverted correlation or connection is established through adaptation, and the world is seen as it originally was, even though the pattern of spatial stimulation on the retina itself remains spatially inverted from the normal case.

In modal structural language, the modal filaments constituting color quale with certain visual locations are at first disrupted, and then reformed. One might wonder if it is clear that the structure of alternative spatial possibilities for sensory receptor stimulation, that aspect of the modal environment of such a modal filament, and indeed that spatial aspect of received light, is irrelevant to the spatial nature of the experience on this conception. No, because there is a kind of spatial structure preserved even across inversions of visual experience. But some of the structure of two-dimensional visual space is also due to connections with motor activity. One might wonder if in the end all that really matters is the space of action. But no, because then there would be no disruption of visual space at first, and only slow adaptation. A relevant modal filament gathers along its length from stimulation to action the modal environment, involving alternative possible stimulations and actions, that constitute it as an experience with a certain location in the visual field.

There is also a second phenomenon that supports such a modal structural understanding of the basic structure of flat visual space against a simple information-based account rooted in retinal positions.

Stable spatial position of stimulation on the retina does not correspond well with stable position in two-dimensional visual space. If you stand still and look at something, its visual position doesn't seem to change, even its position in two-dimensional visual space. But your eyes move many times each minute, changing the place on the retina where relevant stimulation is received. These eye movements are of many types.[14] But focus on a particular and frequent type of movement, called a "saccade."[15] These

---

final stage, so things are visually inverted but it isn't noted. But this seems phenomenologically implausible. See also Hurley (1998).

[14] Palmer (1999: 520–526). [15] Ibid., 523–524.

are very rapid and abrupt eye movements that are ballistic in nature, in the sense that once such a movement is initiated, it cannot be altered.

We don't perceive these movements, because perception is in some way suppressed during the movements.[16] But nevertheless, after a saccade, your retina is receiving stimulation from a given stimulus in a different location than it did before. Yet you experience no difference in visual location. Some correction is made, that allows for position constancy in experience. This certainly limits the significance of retinal position to visual position.

There are various theories about how this happens.[17] The "afferent theory" is that after the eye achieves a new position, the difference in position is registered by neural activity in a way that compensates for the change in retinal stimulation. The "efferent copy theory" is that the very neural action that initiates the ballistic motion of the eye also leads to anticipatory compensation for the initiated change in retinal stimulation. There are also "direct" theories that hold that the compensation is due to perceptual variation introduced by the motion, for instance compensation that subtracts from experience motion that occurs uniformly across all parts of the visual field. And there are various complicated variants of these ideas.[18]

It is reasonably clear that at least much of the correction is in accord with the efferent copy theory. You can note this yourself by gently pushing on one of your eyes, which makes things seem to move, because the movement of your eyes is then not neurally initiated in the usual way that would lead to correction. But it really doesn't matter from our perspective which of these views is correct. What matters is that the correction happens. The type of modal filament constituting a given color quale with a given location in the visual field tolerates not only some variation in what motor activity it supports, but also some variation in which combined forms of eye position due to saccadic movement and location on the retina can initiate it. While it is possible that this involves a spatial map somewhere in the brain that captures visual position in a stable way, that need not be so. What is required is a modal filament that gathers together the relevant structure along its length.

The various theories of visual position constancy may however suggest some objections to my overall proposal. For instance, the afferent and efferent copy theories may seem to involve different sorts of unconscious perception. Or they may otherwise involve neurons in the constitution of something other than phenomenal consciousness.

---

[16] This does not imply that it is completely eliminated. It might exist in something like the manner of blindsight.

[17] Palmer (1999: 339–342).

[18] One proposal is Grush (2007b).

But remember that the autonomic nervous system is not a part of what constitutes phenomenal consciousness on my view either. What neural resources matter to phenomenal consciousness are those that mediate between sensory receptors and motor action. And the kinds of neural signaling that the efferent copy theory of position constancy involves are not of that form. They rather initiate certain specific channels between sensory receptors and motor action rather than others. One might object that saccades are motor activity, but still they are not of the voluntary and relatively gross intuitive kind on which we are principally focused, and indeed they are analogous on this model to dedicated forms of motor activity under the control of dedicated stimuli that are distinct from a unified nervous system and unified phenomenal consciousness, although in this case there aren't even intuitive sensory stimuli involved. To apply such a response in the case of the probably false afferent theory would require that some stretch receptors in the eye muscles not count as sensory receptor cells, and such a denial is itself certainly deniable. Still, in any case, the direct theories more clearly involve sensory reception of some kind. But once again it isn't really motor action that is controlled by these stimuli on such accounts, but rather adjustments in the form of neural channels. And the phenomena required by the direct theories could involve a lesser form of consciousness, analogous in some ways to blindsight, which might be true in the case of the afferent theory as well. If necessary, such a treatment of the efferent copy theory is also available. And indeed another objection to my view based on the efferent copy theory probably requires such a response. Even on that theory there is a neurophysiological difference between different combinations of registered eye position and retinal position that are experienced as a single spatial location, while my suggestion has been that such differences in stimuli generally introduce differences in phenomenal consciousness. But, I reply, such differences needn't occur specifically in visual experience. Your eye after a saccade may be felt in a different place. And different degrees of consciousness are also available in this case as well.

So the bottom line is that we are in the clear regarding these objections based on saccades. We can live with any of that.

3

The modal filament that constitutes a given visual quale is a path of causal influence that passes through the very active arena of the nervous system from stimulation towards action. But it need not be a spatial filament or path, except perhaps at its beginning. It is merely an independently variable aspect of what intervenes between stimulation and motor action. It would be relevantly inflected during saccades as if metaphorically bent and focused by a lens, and this sort of lensing will be important again in Chapter 7 during our treatment of more complex forms of visual particularity. But the crucial point for now is that the filament gathers modal structure relevant to its phenomenal features along its length. We discussed such structure relevant to phenomenal color earlier, and have already considered here some modal structure relevant to its most basic spatial individuality, which correlates relatively peripheral aspects of the nervous system. But there are other sorts of relevant structure we need to explore, which depend on what comes between stimulation and action.

We have seen how opponent-processing is important to our color experience. And there is also an analogous structure of neural activity that helps constitute additional spatial structure in two-dimensional visual space, beyond its rough topology. But once again it is important to think of it in the right way, which is the modal structural way.

This sort of structure is best understood in the case of our visual experience of straight lines of various orientations, lines where there is an experienced difference in luminance or color. But such structure is also plausibly relevant to our experience of two-dimensional visual size and shape. A quick schematic sketch of how it works in the clearest case is as follows.[19]

Photoreceptors are inhibited in their firing by the reception of light. But they can synapse in an inhibitory or excitatory way on other neurons. Beyond the first layer of photoreceptor cells, there is another layer of neurons in the retina, in which each is inhibited by the photoreceptor directly above and excited by the photoreceptors that surround that first one. The effect of this is that such a second-level neuron is highly stimulated by a dark circle with a central hole of the right size, a donut shape.[20] There are also such neurons that respond in the opposite way. This is one kind of opponent organization, a center-surround organization. When large numbers of such contrast-sensitive neurons are present, say in various lines across the retina, yet higher levels of neurons can become differentially responsive in similar ways to the presence of other sorts of spatial patterns in light stimulation, for instance the presence of a dark or

---

[19] Groh (2014: 74–80).

[20] It is only second level in one sense, since other neurons may intervene between it and sensory receptor cells.

light line with a certain orientation. And it turns out that the way in which various receptive neurons are organized in us is systematic, such that neurons that prefer horizontal lines are closer to those that prefer nearly horizontal lines. This is another sort of organization reminiscent in some ways of opponent organization, but in which there are many different alternatives opposed in different degrees. The spatial organization of such neural fields mapping differences in orientation is not constitutive of the nature of experience itself on the modal structural view. But rather, for what amount to reasons of efficient engineering, that spatial organization corresponds to differences in connectivity of neurons that intervene between sensory stimulation and action that are so constitutive. More details: The neurons of the retina itself have a complex layered structure. The cones and rod cells are, somewhat unintuitively, in the layer farthest from the lens of the eye. These photoreceptors connect to ten different sorts of "bipolar cell" neurons, which have long axons that project into a layer organized into a series of ten different strata, with each bipolar cell synapsing on just some strata. At the innermost side of this layer, there are twelve types of ganglion cell neurons, where most have dendrites reaching into one distinct strata. Amacrine cells are neurons that inhibit responses within a stratum or between strata. There are at least twenty-seven types. The ganglion cells have axons that then synapse on different parts of the brain, signaling into the interior neural maze, in the nonsemantic sense of signaling that the fan of modal structure should admit.[21]

So far the moral is merely that things in even the retina are complex. But the particular sorts of retinal structure relevant to the orientation of lines work as follows.[22] Sensory receptors and bipolar cells produce what are called graded potentials, which are continuous changes in electrical potential. Ganglion cells produce the more familiar spike discharges. But both bipolar cells and ganglion cells have the kind of center-surround sensitivity noted in the sketch earlier. Axons of the ganglion cells synapse on cells of the lateral geniculate nucleus, or LGN, of the thalamus, which we encountered earlier. LGN cells have a similar center-surround organization to ganglion cells. They respond to similar stimuli. The LGN is a complex three-dimensional structure. It has two layers consisting of large neurons and four consisting of small neurons, alternating between layers responsive to one eye and the other. But what is most interesting to us is what happens at the next neural stage, in the striate visual cortex, called "V1."

---

[21] Werblin and Roska (2007).

[22] Palmer (1999: 146–157).

Hubel and Wiesel famously identified in V1, on the way to their Nobel Prize: (i) neurons that seem to function as edge detectors, responding maximally where dark edges on light, (ii) neurons that seem to function as line or bar detectors, responding maximally to dark or light lines, and (iii) "end-stopped" neurons whose response to such lines, edges, or bars falls off if that line, bar, or edge is above a certain length. And such detecting neurons are organized in a relevant way within V1. The striate cortex in each hemisphere of the cerebrum corresponds to one half of the visual field, and is organized so that nearby areas of the visual field correspond to nearby areas of the cortex, although within each local region of V1 there are ocular dominance slabs, dominated by one or the other eye. The cortex has six layers of neurons, organized at a fine grain into hypercolumns of about a millimeter on a side, which run down through all six layers. Within each hypercolumn, there are cells arranged along one spatial axis that are differentially sensitive to different orientations of lines in a systematic way. Closer cells are most responsive to more similar orientations. There is also at least some reason to believe that a perpendicular direction in the hypercolumn corresponds to a progression of size-scales, from small to large.[23]

These various details fit our sketch. But as I said, it is important to think of this neural structure in the right way.

The orthodox view of what happens in all this is information-based. On that view, a neuron is a line-detector for a line with a certain orientation in a certain region of visual space just because its firing is causally correlated with that kind of sensory stimulus, because it carries the information that there is such a visible line.

I have no strong objection to the existence of information processing of this sort, or to its importance to vision. But there are debilitating problems with this idea as the basis of an account of visual phenomenal consciousness. In that role, the effects that such a firing has also matter. If we were wired to react to some neural state that correlates with visual line input with behavior appropriate to pain at the tip of the big toe on the left foot, it is hard to believe the state in question would constitute the presence of a line in phenomenal consciousness. And it isn't plausibly accidental to our

---

[23] I presume the orthodox view, as in Hubel (1995). Another tradition, rooted not in neurophysiology but psychophysics, maintains that neural activity in V1 involves spatial frequency analysis rather than edge or line detection, something closely analogous to Fourier analysis. See R. L. and K. K. De Valois (1988). Palmer (1999: 186–192) suggests a synthesis, with spatial frequency analysis prior to orthodox edge detection. Another complication is that the size tuning I am suggesting is discussed in the tradition of spatial frequency analysis and not the orthodox tradition.

conscious visual experience that there are different neurons that respond to the full range of different alternative orientations of lines, and arranged in such a way as to map intuitive similarities. But this systematic structure would play no obvious role in simple information-based stories of phenomenal consciousness, which deploy only necessary correlations between each neuron and the relevant type of stimuli. There are, we should note, more complex and plausible accounts of visual representational content that deploy similarity structure as well as informational correlations in a crucial way,[24] but that is a move in the direction of modal structuralism. Another problem for simple information-based accounts is that the information about line orientation must be integrated farther up the neural hierarchy with other bits of information to form overall visual space, and I am aware of no good accounts of how through causal correlation such a phenomena could occur. Causal-correlational information does not "flow" in any natural way, as is often loosely said by philosophers. The firing of a higher-level neuron could be correlated with the firing of such a first-order edge detector, but how could it be at once correlated with that and the firing of another first-order color detector? And even if it was doubly correlated in some way, wouldn't it simply represent just those firings and not what they in turn represent? Perhaps the idea is that there is supposed to be somewhere in the brain some neuron or complex of neurons whose firing is causally correlated with some type of complex environmental stimulation involving both color and orientation, and what happens in V1 is simply a mechanism on the way to that representation-determining correlation downstream. But that fancy and very specific higher-level correlation is not what has been observed. Nor is it particularly antecedently plausible, as some have claimed, that some sort of temporal resonance among different neurons representing different neural features constitutes a representation of the whole. Why would that have any unified behavioral effect?

But, in any case, my purpose here is not to deny informational representation in all important psychological roles, but rather to develop a modal structural account of sensory qualia. I claim that the proper way to think of the activity in V1 for our purposes is this: In some region of V1, corresponding to some rough location in visual space, there is neural activity that contributes to the relevant modal filament one sort of alternative, in contrast to various merely possible alternatives; in particular it contributes being oriented this way rather than in one of the contrasting

---

[24] See Neander (2017).

ways. This is a kind of opponent structure, although involving many alternatives. There is hence a refinement of the overall modal structure of the modal filament, or more exactly to the collection of minimal modal filaments, which constitutes the relevant visual quale, and so refines the rudimentary spatial structure that such experience would otherwise possess in the absence of such a tweak.

What the V1 neurons do was discovered by looking at correlations with stimuli, but that is not what constitutes their contribution to conscious visual experience on this understanding. Rather, given the way in which visual sensory receptor cells and motor activity are already loosely "topologically" linked without this additional contribution of modal structure, V1 neurons add what is required to stiffen and straighten up our visual experience so it contains such straight lines and the possibility of such lines. We hence go beyond mere topology. And there need be no higher-order representation of lower-order representations or their contents on such a view. The unity of visual experience is assured by the fact that the individual who has it senses and acts as one animal, although the detailed nature of that unity depends on many specific neural facts about them.

It also seems antecedently plausible that the same general sort of extended opponent-structure we have found to be present in line-orientation "detectors" also helps provide two-dimensional visual experience with further structure. For instance, there is reason to believe that visual shape perception is hierarchical, that a given flat triangle is registered abstractly as being a plane figure without a hole, and more specifically as a triangle, yet more specifically still as a triangle with certain angles and certain sized edges.[25] And there is evidence that these different levels of shape are registered at different levels in the neural hierarchy of visual perception.[26] Any systematic contrasts and similarities at one such neural level would invite a similar interpretation to that we have given neurons responsive to line orientation.[27] And the fact that registration of different levels of abstraction occurs at different neural levels is also quite friendly to the modal structuralist conception, which gathers together relevant structure along modal filaments.

---

[25] Green (2015). This concerns shapes with depth as in Marr's 2 1/2-D sketch, but many relevant experiments involve presentation of plane figures. For a general introduction to shape perception, see Palmer (1999: 362–407).

[26] Biederman (2013); Green (2015: 26–27).

[27] As far as I know, there is no clear evidence that bears on this one way or the other at this time. My proposal suggests this happens but doesn't require it. Perhaps that is a prediction.

4

There are certain standing biases in our perception, such that we tend to interpret visual stimuli in just one of a variety of different possible ways. This may generate visual illusions when these biases don't accord with local reality. This is a source of more examples that may suggest that information-based representation is not the root of our visual consciousness, and so more support for neurally based alternatives like modal structuralism. But on the other hand, familiar stories about this biasing deploy stored information in a way that may seem to present an objection to my view. So these cases require our consideration in more than one way.

If I may be allowed a case beyond two-dimensional vision, consider a group of three circles in a row, shaded so that the outside circles seem brighter at the bottom and the middle one seems brighter at the top. They look as if the middle one protrudes and the outer ones are concave, when viewed right side up. But that look is reversed when the group is viewed upside down. This shows that some features of our spatial perception are not simply determined by their environmental causes, since the causes in this case are the same despite a difference in what is perceived. And another point to focus on is that this case shows that our visual system in some sense presumes that light comes from the top.[28] As what seems to be on top shifts, so does our experience.

One immediate question is how all this fits with orthodox representationalism about consciousness. While the outermost stimuli in such cases may carry information about what is out in the world, still any information processing due to our natural presumption about light sources breaks such links, it would seem. So it is a puzzle how such a visual content can accrue to mental representations on such a view.

Perhaps the idea is that such "processing" simply is a mechanism that creates, eventually, a direct informational correlation between some upper-level representational state and what it represents out in the world. But notice again that in the case described, that correlation does not exist when the pattern is viewed upside down. In fact, since the stimulus in this case is two-dimensional and involves no depth even when viewed right side up, that correlation is absent however the pattern is viewed.

Causal-informational accounts of content are often buttressed with an account of "normal" causes to deal with cases like this. Perhaps the statistically normal cause of such a stimulus matches what is experienced

---

[28] Groh (2014: 48).

when it is right side up. But then we would expect that the visual experience you have on the basis of the stimulus would vary as your statistical environment varied, and without any variation in neural structure or what you would do on the basis of such a stimulus, even what you would say about it. And that is not plausible, since you can report how the thing looks. Since you cannot change your history, it may seem that normal causes are in particular statistically normal causes in your evolutionary or learning history. But then your present physical duplicate with a different evolutionary or learning history would have different experience, yet they would speak and act exactly as you do on the basis of your current experience. This is quite implausible.[29]

It is more common in such a case to grant that the state of our neurophysiology has something to do with this illusion in another way, that it somehow encodes the presumption that light comes from the top, which is added in like a premise to convert an original unconscious representation of the stimuli, through something like inference, into another experienced representation with a different content. But representational, information-based accounts of this premised content would seem to require neural sentences that covary in relevant ways with the environmental conditions presumed, which would be replaced by alternative neural sentences expressing different presumptions if other contingent conditions obtained. And there is no reason to think this is a plausible account of how the necessary presumptions are encoded in the brain. The kind of information resting on causal correlations can no more easily be stored than it can flow, despite loose talk by philosophers that suggests otherwise. Perhaps one relevant idea is that there are fixed mental words that, when in a special perceptual role, covary with environmental stimuli, and then can be stored otherwise but retain their content, perhaps partly because they will be treated as having the same content later on. But good luck finding determinate mental words of that type within realistic neurophysiology, and also innate registrations of the right mental sentences to carry the presumptions.

So while I have no objection to thinking of perceptual biases in such representational, information-based ways if that can plausibly be done, still we don't at present understand such conceptions well enough, I claim, for them to constitute a strong objection to modal structuralism as an account of conscious visual experience incorporating such biases. It is clear that

---

[29] Mendola (2008) includes a more extended critical discussion of such etiological accounts of normal causes.

neurophysiology matters, as modal structuralism predicts, and not at all clear that such informational-representation does as well.

There are alternative information-based views with some popularity. While we have already noted difficulties for the most popular alternative view, that "the representation of objects and scenes is not preceded by prior representational stages,"[30] since there is little experimental evidence for the specific, sophisticated, higher-level representational states such information-based accounts require, many now go in the other direction, and require yet more elaborate intervening representations. There is currently serious interest in Bayesian theories of perception.[31] On such models, perceptual systems execute unconscious statistical inferences from stimulations to experiences, by assigning probabilities to hypotheses about environmental configurations, and updating these prior probabilities in light of sensory stimulation in the traditional Bayesian way. This apparently presupposes that there are unconscious mental representations of probability assignments over external configurations. It would be particularly hard to get explicit representations of that very fancy sort of standard informational stories about representational content.[32] But once again, I have no qualm about the use of Bayesian methods to discover and represent perceptual biases of this sort, or indeed any sort of information present in these cases if it is really there. But the modal structuralist about phenomenal consciousness can avoid this whole difficult and problematic tussle.

On my view, the nature of these perceptual biases and what constitutes us to have visual experience with relevant qualia may have nothing whatsoever to do with information or representation based in information. These features of our visual experience are introduced in much the same way as the opposition of color hues, by modal structure in what neurophysiology intervenes between sensory stimulation and action. That may or may not generate illusion depending on how the world goes. It is of course not an accident if through evolution, or through sensory adaptation as to inverting glasses, we are reasonably well suited to our local environments in these ways, so that resulting perceptual illusions are not frequent or debilitating. But that is not because our mental or neural states bear information rooted in correlation to the environment. Of course, our

---

[30] Orlandi (2013, 2014). [31] Knill and Richards (1996); Rescorla (2015).

[32] Perhaps it is more charitable to think of these probabilities as mapped by something like intensity at each node of a Bayesian net, with the strength of links between nodes representing strength of conditional probabilities, with these general correlations delivered by evolution. But even the nodes themselves are implausible representational elements given standard accounts of representation.

principal interest is not perceptual content but perceptual qualia. But it is antecedently plausible that the relevant visual qualia would be the same if our neural states were the same in such a case, whatever is out there and however it is statistically correlated with our neural states. And we have seen no good reason to doubt that initial plausibility.

5

But let's dig a little deeper. There are many regular spatial illusions that infect even two-dimensional visual experience, which involve a mismatch between actual environmental causes and qualia, and so between information and qualia.[33] These are hence, at least on the surface, more reason to prefer a modal structural account to orthodox representationalist alternatives deploying information, since information rooted in correlation must be dominantly veridical. But once again there may seem to be representation-based accounts of these illusions that constitute objections to modal structuralism. So let's consider these cases.

In our experience of the Kanizsa triangle, three dark circles with appropriately arranged cut-out wedges, which are similar to three Pac-Man icons with their mouths facing in, create the illusion that there is a phantom triangle that spans them.[34] And our experience of lengths is sensitive to orientation. Consider the T-illusion, whereby a straight upright placed at the middle of a straight base of equal length looks longer, and indeed longest if oriented just to the right or left of vertical. Then there are various illusions involving the interaction of angles and straight lines: the Hering illusion, the Müller-Lyer illusion, the Zöllner illusion, the Poggendorff illusion.[35] There are also size-contrast illusions. In the Ebbinghaus illusion, two circles of the same size look respectively larger or smaller if surrounded respectively by four evenly spaced smaller or larger circles. In the Delboeuf illusion, two circles of the same size look respectively larger or smaller when encircled by smaller or larger circles.

These illusions are reasonably well explained by "natural scene statistics." This predicts the nature of a visual experience on the basis of statistical links between retinal image features and the statistical distribution in the environment of objective features that cause such image features, so that "perceptions generated by light patterns projected onto

---

[33] For a survey see Howe and Purves (2005).

[34] Groh (2014: 85).

[35] Masrour (2016) discusses a three-dimensional angle illusion.

the retina are determined by the probability distributions of the possible sources of those projections"[36] in the "natural scene."

But there are different ways to think of these facts. These are perceptual biases, and so the same alternatives are available for the explanation of these biases as we discussed in the previous section. And the information-based alternatives remain problematic. Natural scene statistics might be deployed as an account of normal causes, meant to refine informational accounts of representation. But that is initially implausible, because it implies that experience would vary if things out in the world changed in a way that had no causal effect on the perceiver and what they would say. And if the statistics that are relevant in determining normal causes are those in some learning or evolutionary history, two creatures who are current physical duplicates but have different evolutionary or learning histories would implausibly have different conscious visual experiences while acting and speaking in the same way. One might deploy natural scene statistics instead as a proposal about what sorts of presumptions are built into our perceptual information-processing.[37] But in the previous section we saw various difficulties that must be faced by such conceptions. So these representation-based developments of natural scene statistics seem to provide no more serious objection to a modal structural understanding of perceptual illusions than we discovered in the previous section. On balance, given that illusions don't fit well with veridical representation, and given the difficulties of information-based ways of understanding natural scene statistics, such standard illusions in fact provide further grounds for a view of the general type proposed here.

Also, there are other spatial aspects of two-dimensional visual experience that might be considered illusions, but don't fit at all well with natural scene statistics. For one thing, this visual space is two-dimensional, and finite, and has a limited grain of detail that is more determinate in the center than at the periphery. Things in some sense get smaller in this experience as they recede, while they do not become objectively smaller as they recede. And we only see one side of things at a time. But none of those peculiarities are due to statistically normal environmental conditions.

---

[36] Howe and Purves (2005: 3).

[37] Ibid., 103–105 propose a statistical model of sensory processing on the basis of natural scene statistics, whereby a neuron encodes probabilities of objective correlates, and such action across many neurons somehow sums together to visual experience, or perhaps rather (page 104) in which a neuron functions "as an estimator of the value at a point in the probability function." But I find their proposal somewhat obscure.

On the modal structural view, these facts are principally due to the way our vision is mediated by the two-dimensional arrangement of sensory receptors on the retina. But it is also clear that experiences of relative visual sizes do not correspond well to the relative sizes of projections of stimulations on the retina, that there is a kind of size constancy that inflects even our two-dimensional visual experience as things recede. This requires a treatment like perceptual biases.[38] More generally, it is clear that the geometry of visual space is not only unlike the three-dimensional geometry of real space but also unlike the two-dimensional geometry that reality projects onto the retina.[39] But we already expect this deviation from the spatial structure of the retina, because of the way that modal structure accounts for features like the orientation of straight lines, and because of the modal structural treatment of perceptual biases.

Other spatial illusions of two-dimensional visual experience involve grouping. Consider a 4 × 4 matrix of dots, with the columns a bit more widely spaced than the rows, and the second and fourth row slightly lighter than the first and third.[40] You can see this as involving principally four rows of same-color dots, or four columns of spatially close dots, and switch between these ways of seeing more or less at will. These are phenomenologically distinct experiences. However, since you can maintain these perceptual groupings even if you turn on your side, it isn't quite as things horizontal or vertical in your experience, as rows and columns in that sense, that you see the dots.[41] You see them simply as grouped together one way or the other. And that difference corresponds to nothing objective in the world. So informational representation seems an implausible explanation of what happens in this case. Such perceptual grouping does often plausibly involve an expectation that a group will persist over time, indeed even over changes in membership, and to move together.[42] But you most plausibly expect the entire structure of the grid to persist and move in that way, including both rows and columns, in this case. So that can't explain the phenomenological difference between the two experiences.

I think this is more work for modal structure. But this phenomenon must involve a kind of palimpsest of modal structure, much as when you see a color bit as at once part of a two-dimensional color experience and as the presentation of the color of an opaque surface. There is a hierarchy of individuals experienced, the entire structure as a whole, which is seen as consisting of either individual rows or columns of dots at will, and then

---

[38] Palmer (1999: 314–332). [39] Masrour (2015); Wagner (2006). [40] Peacocke (1983).
[41] Green (2016). [42] Ibid.

most specifically the individual dots. A plausible modal structural account of the particularity present in two-dimensional visual experience must allow for this sort of hierarchy.

6

Fanciful cases may also be revealing.

One that has been proposed as showing a difference between the informational content and qualia of spatial experience is Doubled Earth.[43] On Doubled Earth, everything is doubled in size, but relevant laws are similarly adjusted so everything remains over time isomorphic to Earth. Oscar here and Big Oscar there are otherwise identical twins in otherwise identical environments. They are alleged by some to have the same phenomenal experience that yet differs in representational content, that differs in what information it carries about the world.[44]

My proposal is consistent with this claim, but also with its denial. The explanation of this is what is most relevant here.

One question is what it means for a certain phenomenal experience to present as involving a certain size. If visual experience involves a number of minimally discriminable bits, then in one sense another experience is larger if it involves more bits. And if that's all there is to it, then Oscar and Big Oscar do not differ in spatial qualia. They both also plausibly neurally register size differences within their individual experience in something like the manner we have seen they plausibly register line orientations. But that neurophysiological structure too would be preserved across them. However, it might be that the visual experience of an elephant presents as larger than that of a gnat, even if it involves the same number of minimally discriminable bits and the same sorts of neural registrations of size differences. It is fairly clear that this difference in qualia is something that we can conceive. We can make judgments about lengths and imagine them doubled, and indeed we can experience visual lengths and imagine experiencing them doubled, and so we can understand the case in that sense. But this might, or might not, involve an illusion of possibility, such as when we imagine that Hesperus and Phosphorus are distinct when in fact they are just the one planet Venus that cannot be distinct from itself.

The basic issue about this conception of visual size is, according to a modal structural account, whether the absolute size of the spatial array of sensory receptor cells and of the spatial structure of motor activity of a

---

[43] B. Thompson (2010). [44] Ibid., 156.

creature can make a difference to the phenomenal size of their experience.[45] My proposal can allow that it does and allow that it doesn't. It depends on what aspects of the spatial structure of sensory receptor cells and motor neurons matter. I don't know if this difference between gnat and elephant experience reflects a real possibility. But perhaps some reason to doubt it comes from a case in which the sensory array remains the same size but the motor array doubles in size, and yet through various tricks with mirrors and screens the discrepancy between motor size and sensory size is never noticed. Then if size matters it would apparently matter in an inconsistent way. But perhaps instead the size of the motor array or the sensory array alone is what matters to phenomenal size.

So far I have just claimed that modal structuralism is consistent with whatever is the truth. But about a second and somewhat analogous case I think there are grounds for greater clarity. And perhaps this analogy is grounds for a similar treatment of the case we just discussed.

Consider Mirrored Earth, in which everything, including yourself, is as the actual world looks in a mirror.[46] It seems clear that your mirror twin has a different spatial experience than you do, inverted as in a mirror. And a modal structural account locates the difference between you and your twin in the actual spatial arrangement of sensory receptor cells and motor activity. Yours mirror those of your twin.

But some may think that the difference here is also a modal illusion. For instance, some claim that you are not intrinsically different from your mirror twin, because if the space in which we were to reside were nonorientable, you could become fully congruent with your mirror twin by moving through the world, much as travel on a two-dimensional Möbius strip can seem to flip two-dimensional figures that would otherwise be mirror-inverted into congruence.[47] And it may seem that intrinsic differences must explain these alleged differences between your experience and that of your mirror twin. I agree that any such difference must be explained intrinsically, but think that such wildly hypothetical possibilities as what might happen in nonorientable spaces are irrelevant to understanding our own intrinsic spatial structure and hence that of our experience. Indeed, the Mirrored Earth case seems to be evidence of that.

Mirror inversion is interesting in many ways.[48] It would seem that mirrors invert in a particular way, not up-down, but left-right, or perhaps in-out. And our visual experience may well present a spatial structure that

---

[45] Or less plausibly the absolute size of the intervening neurophysiology.

[46] G. Lee (2006).

[47] Baker (2012).

[48] See Mendola (2021: chapter 4).

corresponds to that. And since inversion of your look in a mirror doesn't seem to involve any physical difference in the physical relations between various parts of your reality and corresponding apparent positions on the mirror, this phenomenon can seem mysterious. I believe that this apparent fact about mirrors is due to the peculiar way we *think* of left and right, up and down, and in and out, which is due partly to our natural visual and actional orientation forward, partly to the left-right symmetry of our bodies, partly to our natural life as upright walkers in gravity,[49] and perhaps partly to peculiarities of our spatial language. We continue to think of the dimensions this way even when floating in space, unmoving, and with our heads facing sideways. I believe that we don't think of the three intuitive dimensions of space in quite the same way, even though in reality they are quite the same. And quite possibly these peculiarities even infect the spatial structure of our visual experience, not just our fancier cognition. If so, this is another aspect of visual experience that is not plausible due to informational representation, although neural phenomena are clearly relevant.

7

Another way in which our visual experience of space involves qualia that are not plausibly due to objective spatial features of objects is evident when we consider different ways in which the same space is presented in different senses. Some of the most striking cases of this sort involve a comparison with other animals. But there are other cases worth considering.

We already noted Molyneux's problem, and experimental work, that suggests that there is a difference between the visual and tactile experience of shape. It takes the formerly blind who have come to see a little time before they understand what felt shapes correspond to what seen shapes, which suggests there is a difference. There is also an intuitive phenomenological difference between what an edge feels like and looks like. But there is of course only one edge out in the world in such a case. These differences in experience are plausibly due to neurophysiological differences.

There are different sorts of sensory receptor cells that play the analogous role to photoreceptor cells in different senses. Consider some details of the somatosenses, which we haven't yet much discussed. There are muscle spindle receptors that respond to the length of muscle fibers, and Golgi

---

[49] Denyer (1994) provides another story of this type, though focused as an explanation of peculiarities of our spatial language.

tendon organs that respond to the force supplied by the muscle to tendons. They are set up in a complicated oppositional structure matching the mechanical oppositional structure of the joints and limbs. There are also skin senses, which involve a variety of different specialized sensory receptor cells that respond to different sorts of stimuli, for instance vibration or pressure, and have various locations.[50] And the organs of balance in the ear involve hair cells of a sort resembling those that mediate hearing.[51] These organs include three fluid-filled semicircular canals oriented in various ways, which respond to motion of the head because the fluid drags behind the motion of canals, hence engaging hair cells. And there are also the utricle and saccule filled with little stones that function in a similar way to detect linear motion.

The way in which spatial content accrues to our experience in these different senses differs from the very beginning. In vision, different photoreceptors are located at different locations within the retina. They are arranged in a spatial map, which corresponds in ways we have discussed to spatial structure in visual experience. And light that preserves some spatial structure in the objects of the environment is focused on the retina in such a way as to preserve this environmental structure, although inverted. But the various somatosenses are quite different. The skin sense receptors have locations that might be spatially mapped in something like the way that visual receptor cells are mapped, although the surface of the body is continuous in a way that the retinas are not, and there is no particularly good analogue in the case of skin receptors for the focusing of light on the retina. And there are different skin sense receptors that mediate different sorts of experiences which are not unified in quite the obvious way as visual experience. And while the muscular sense organs are also located at particular places, they seem rather to primarily differentially register particular muscle and joint action, and they are arranged in some cases in a crucial oppositional structure. The organs of balance directly register movement and orientation in space, but the only positional information they encode is right ear versus left ear. The locations of hair cells within these organs seem subsumed into their coding of motion.

There are also differences in what sorts of spatial structure are registered after immediate stimulation of these various senses, and in the way our neurophysiology works to do this. Because of the arrangement of our two eyes and their lenses and higher-order neural activity (even, as we have seen, activity at least close to the motor periphery), experience of a kind of

---

[50] Groh (2014: 64–65). [51] Ibid., 180–184.

three-dimensional spatial structure of external visible objects is somehow induced in us on the basis of two-dimensional spatial structures of retinal stimulation. This three-dimensional spatial structure is at least in the first instance egocentric in form, focused on the self, although we do also have an experience of the relatively stable spatial locations of objects as we move around. There is a long and indirect route, perhaps through the two-dimensional visual experience we have been discussing, to any experience of the objective spatial relations of seen objects in the world. But in the muscular senses, the oppositional structure of the sense organs allows a somewhat more direct sensitivity to limb position, and so to some relatively abstract three-dimensional features of our own spatial structure. But as the fan of modal structure might expect, the resulting spatial experience seems to retain the kind of spatial opposition of for instance one finger movement as against an opposite motion, which is relatively close to what the relevant sensory receptor cells most directly register. The skin senses, however, can quite directly register the location on the skin of stimulation, though it seems they allow much greater spatial resolution in some areas, for instance the lips and the palms of the hands, than others.

It is also undoubtedly important that there is integration of the spatial structure available through different senses. For instance, there is the phenomenon of visual capture. Sounds will be heard coming from the location of certain mouths on a television screen even if the television has only one speaker. And sounds seem to come from the moving mouth of a ventriloquist's dummy. And so vision seems to dominate hearing in some cases. But even within the working of each sense alone, the constitution of spatial experience is quite neurophysiologically complex, and differs from that involved in other senses. And this plausibly constitutes corresponding phenomenological differences.

In the information-processing tradition, it is sometimes said that there are differences in the way that information may be "coded" across the different senses. I haven't completely suppressed this language even in my recital of relevant neural and psychological data, though I think it reflects neurophysiological differences that should ideally be expressed in a different way when applied to the issue of phenomenal consciousness. Firing patterns of neurons, in other words time variance in the succession of roughly similar-sized action potentials, sometimes seems important. This is neither quite digital nor quite analog in traditional ways. Firing is quantized, but on the other hand the patterns may continuously vary, and also take complex forms. There are also graded potentials, as noted earlier. But the main point is to contrast this case with another case. It also matters which neuron is firing and "where" it is, either spatially if it is a sensory receptor cell or motor neuron, or what it is connected to and how. Groh, in the informational tradition, suggests a difference between mapped coding, as when it matters which neuron fires, and metered coding, as when it matters how a neuron fires. There hence might be two different ways in which spatial structure might be encoded. She suggests that vision and the skin senses initially use mapped coding for spatial structure, and the muscular senses initially use metered coding.[52] But of course it also no doubt matters what muscular sense organs fire and where they are. So what do the muscular senses encode and how do they do it? Do they just encode contraction and force or also the oppositional structure in one limb or joint, taking into account their position?

On the modal structural understanding, we don't have to answer these questions, which may indeed not have determinate answers, while we can accept that there are relevantly different ways such as those noted here in which neurons work in mediating different senses. Rather, all of the complex neurophysiology involved in our experience of a limb position is part of what constitutes the relevant modal filaments. And I've noted ways in which both sorts of "encoding" seem to be reflected in our experience of limb position.

The spatial experience we have of our bodies, once the skin senses and the sense of balance and the muscular senses are integrated into one experience, is quite complex, but with a complexity different from that of visual experience. We experience a kind of rough arrangement of our limbs and torsos and heads, and various axes on which various limbs might be moved, and of continuous locations on our skin where various things are happening. But the experienced structure of this space is quite different from that of any visual space. Vision is in humans a distance sense. We see things out there in the world beyond us. Hearing is the only other obvious distance sense we possess. This makes it perhaps unsurprising that our experience of space in touch is different from our experience of space in vision. The differences are due to differences in our relevant neurophysiology, even when it is stimulated by the same parts of the world.

8

Beyond the two-dimensional aspect of visual spatial experience on which we have focused in this chapter, there is visual spatial experience of a more

---

[52] Groh (2014: 155–158).

complex sort, which involves correspondingly complex phenomenology. You have experience of facing and oriented surfaces in depth out there in your three-dimensional environment, but in a kind of space oriented on you. And you experience those surfaces as belonging to three-dimensional objects with sides you don't currently see. Indeed you experience those objects as things that can retain some sort of objective spatial position even as you move around and towards them.[53]

There are spatial qualia distinct from fully objective features of the world present in at least experiences of the first two of those sorts. Certainly the neurophysiology involved in these sorts of experience works in a very complex way which we can't consider at all adequately here.[54] But there is no particular reason to believe that what is understood about it cannot be reformulated in a modal structural guise. And one point about three-dimensional vision to note for purposes of comparison with other senses is the way in which we can zoom in on something by walking towards it, and see more detail than we otherwise would see. This phenomenon does not have very exact analogues in other senses, and this plausibly affects the kind of visual experience we have of space. Even in the other paradigmatically distance sense, hearing, in which we do move our heads or move around to hear better, the detail hence introduced doesn't seem to be spatial detail of the same sort as visual zooming allows. And a creature who had only body senses such as ours would have a very different conception of space, I think. I doubt for instance they would have the conception that space might be infinitely divisible, which is natural to those who can zoom in and out through vision. Other human peculiarities that plausibly affect our experience of visual spatial structure involve the way in which we are especially sensitive to facial differences and structure.[55]

Chapter 7 provides a somewhat more detailed discussion of sophisticated kinds of particularity that appear in our visual experience. The most intuitively dominant of those kinds is especially complex. It involves a rich modal structure, and interaction with other senses in certain ways. But even within two-dimensional visual experience, there appears the somewhat sophisticated form of particularity noted at the beginning of this chapter, through which you can experience a certain two-dimensional visual shape or color patch moving and even changing color and shape.

---

[53] Marr (1982) and Palmer (1999: 85–92) suggest analogous levels of visual perception.

[54] For an introduction, see Palmer (1999: 199–461).

[55] For provocative neurophysiological findings, see Chang and Tsao (2017) and Quiroga et al. (2005).

We will be working towards a better understanding of this case also. But it undoubtedly involves the fact that our neurophysiology of vision is not only quite directly sensitive to spatial and shape properties in the manner we have already discussed but also to motion.[56] And this must involve distinguishing between the motion of things seen and the apparent motion of retinal images due to the constant movement of the eyes. Watching an apparently spatially stable two-dimensional shape over time and watching a moving two-dimensional shape over time are in this way entwined though complementary capacities.

---

[56] Palmer (1999: 465–518).

CHAPTER 7

Time, Cause, and Particularity

There goes a green car. And here comes a brown one, after the school bus. Those sedans might have taken different routes, or come by in a different order. They might have both been red. And one or the other might have been parked in the garage. Each presents such a commonplace but modally robust sort of particularity in our visual experience. We have considered more rudimentary forms of visual particularity present in phenomenal consciousness, in a kind of palimpsest. But this is another form.

It is familiar to assume a representationalist, information-based account of this robust experienced particularity, of that aspect of visual phenomenal consciousness. On such a view, you experience the robust particularity of some vehicle in front of you, as opposed to that of a qualitatively identical vehicle somewhere on the other side of the globe, simply because it is the one you are currently seeing. But the modal structural view requires a different sort of account of this particularity, which would be available also in nonveridical cases like hallucinations and dreams, and even if the world itself contained no such robust particularity and it was always a mere illusion of our experience.[1] It requires particularity qualia.

Experience of robust visual particularity is the topic of this chapter. But to be in a position to discuss it, we must first consider our experience of two other sorts of things – things worth treating for their own sake in any case – namely, time and causation. And we must refine our sense of the way actions and act-tendencies can play a role in constituting modal filaments and hence qualia. These will crucially contribute to our goal. Some aspects of the overall account to be sketched here cannot currently be rooted in details of relevant neurophysiology, about which little is known. But we will be able to see that internal structures of our

---

[1] Papineau (2021: 93–106) suggests the terminology "quasi-objects" for roughly this sort of thing, but I am skeptical of his arguments that they can be cleanly distinguished from traditional intentional objects.

https://doi.org/10.1017/9781009514774.007

neurophysiology are plausibly sufficient for such experience, as modal structuralism requires.

I

The first step in our progress towards an account of robust visual particularity involves temporal experience. We begin with some phenomenological points about that.

There are a variety of temporal properties and relations present in our temporal experience.

There is an order and direction of occurrences in time. There is also simultaneity of occurrences. Those are perhaps the most basic temporal features present in our experience.

But we also experience relative duration, and not merely when the occurrence over time of one event is contained in another that lasts longer. Even two time periods that do not overlap may present themselves in our experience as of different relative duration, perhaps with a merely ordinal, or rather instead something more closely approaching a cardinal comparison possible between their lengths. And there is some experience of absolute duration, comparable between experiences or among possibilities. It can make sense to say that everything one experienced seemed to happen twice as slowly, or at least somewhat more slowly, than it does in reality, and that that constitutes a different experience, as of a different and slower world.

There are also complex properties that we experience that are only in part temporal. Experience of movement is experience of spatial change over time. Experience of velocity is experience of rate of spatial change. And experience of acceleration is experience of rate of change in velocity.

We will begin with the simplest case, our experience of temporal order and sequence. There is a variety of different models of how temporal order and sequence appear in experience.[2] Say you hear a tune, or watch a croquet ball roll through a wicket. There is experience of an ordered succession of notes or of positions of the ball on the grass.

On the "Cinematic" conception of how this works, there is no immediate *experience*, in any straightforward sense, of temporal properties or relations, beyond perhaps simultaneity. Our experience over time is rather a temporal series of momentary stills. On this conception, *cognition* of motion or temporal order doesn't involve recognizing temporal order that

---

[2] Arstila and Lloyd (2014: 75–81).

is somehow already present in the content of sensory experience, but rather involves memory. The temporal order that is perceived or thought but not experienced is equal to the temporal order of the states of experience one has, which is at least in normal cases the same as the temporal order of the worldly phenomena perceived through those states. Each individual moment of experience is of something that involves no temporal succession or order at all, but the moments of experience themselves are ordered in time.

But there are alternative views. On the "Retentional" conception, for instance of Husserl,[3] there is in fact *experience* of temporal properties other than simultaneity, even in the most immediate sense. And, on this conception, it occurs so to speak within the content of a moment or minimal temporal bit of sensory experience, much as the experience of spatial structure may. There are in fact three aspects of that temporal content of a momentary experience according to Husserl. There is *retention,* a kind of primary memory that still reverberates in consciousness, of the preceding note of the tune one is hearing or of the preceding position of the ball. There is also the so-called *primal impression*, which is of the current note or position of the ball. And there is finally a *protention*, which points towards the future as a kind of experienced expectation of a succeeding note or position. Temporal experience is a series of such moments, but each has such a three-part temporally complicated content. The moments need not strictly last a temporal instant; there might be a small finite time it takes to experience even a temporally minimal bit of experience. And alternative models of this same general type, which are sometimes called "atomistic," may of course not involve the specific Retentional conception characteristic of Husserl himself.[4]

There is also the "Extensional" conception.[5] On this model, each distinguishable episode of experience is temporally extended in a way that exactly parallels the temporal content of that episode, and hence incorporates temporal features in that straightforward way. And there are no moments or minimal temporal bits of experience with either a still or quasi-Husserlian content.

Each of these three clusters of views faces characteristic difficulties.

According to Cinematic views, there is no true experience of temporality. But this does not accurately capture the phenomenology of our temporal experience. And indeed the hypothesized stills are apparently

---

[3] Husserl (1991). Gallagher and Zahavi (2014) discuss later developments in Husserl's view.

[4] G. Lee (2014b).

[5] Dainton (2014).

not available in our experience. One way to see that these views misrepresent the phenomenology of time experience is to note that those with damage to the neural region V5 sometimes suffer from cerebral akinetopsia, and are motion blind. They see the world in a series of stills such as the Cinematic view proposes. But they have a radically different experience of the world than normal humans.[6]

This understanding of the implications of the Cinematic view may suggest yet a fourth view, or at least a variant. Perhaps we have temporal experience much as the Extensional view suggests, but that extended experience consists of point-like temporal atoms. While there is *experience* of temporal succession, contrary to the Cinematic view, it is nevertheless made up of extensionless atoms of temporal experience, contrary to the Extensional view. But this view apparently differs from the Extensional view only so to speak metaphysically, with regard to whether temporal experience is made up of independent atoms or not. But we are presuming a variant of the adverbial account of perception and experience, according to which there are not independent entities *experiences* of which these contrasting claims make sense, although there are people with experiences of various sorts. So this difference with the Extensional view is not directly relevant in our context. Still, there might be a question regarding whether the content of our experience presents temporal points or not. That would be one sort of variant.

Still, Extensional views, as well as such variants that allow the presentation of temporal atoms in temporal experience, also face telling objection. If you hear a constantly ascending tone, the Extensionalist will say that the experience of each of the notes traversed is contained in the experience of the ascension. But notice that the experience of those individual notes will no more be an experience of the motion of the tone, on this treatment, than on the Cinematic view. And that seems phenomenologically inaccurate. On the Extensional view itself, no individual tones move in this case, but they are simply at most point-like abstractions from the extensional whole.[7] This doesn't capture how the ascending tone is always experienced as moving.

But there is a reply which the Extensionalist, although not the variant views noted, might adopt. Recall that the Extensional view also holds that there are no temporal moments in experience, no experiences without temporal extent. And that may be thought to help with this point. Just as there are those who hold that the continuum of space is not constituted by

[6] Dainton (2023). [7] Pelczar (2010).

extensionless points which are partless, but rather that there are no bits of space that aren't extended, the fan of the Extensional view insists that there are no minimal moments of experience.[8] Hence perhaps there are no stationary notes at all experienced in this case. But there is also a counter-reply. Consider an ascending tune that consists of several notes, each that persists for a bit. And consider relatively short bits of time that constitute experience of the full extent of a single such note. That the relevant stills are temporally extended makes no real difference to the intuitive problem in play. They wouldn't be experienced on such a view as a hiatus in ongoing motion. But that is the way they are experienced in reality. On Extensional views, it might be said that there is another overlapping experience which constitutes that of two neighboring notes. But that is closely analogous to what the Cinematic theory says about such a case, which is phenomenologically inadequate. The size or existence of temporal atoms is really irrelevant to the issue in question. The Extensional view is no more phenomenologically accurate than the Cinematic view. It fails to capture the kind of motion that can be present in even a minimal moment of experience.[9]

We seem pushed in our account of temporal phenomenology towards the remaining class of views we distinguished,[10] and Husserl's original version is quite phenomenologically sensitive. Still, it too faces characteristic objections.

Derrida focused critically on the notion of the primal impression.[11] He argued that the primal impression cannot be fully distinct from the retention if it can be attached to it in the way Husserl requires. But if it is not fully distinct from the retention, then it would need to have a kind of temporal density in itself. The very being of the primal impression must, it seems, be merged with that of the retention and indeed the protention, though Husserl requires that those things be distinct. So while perhaps the primal impression could be a kind of hypothetical limit of analysis between the retention and protention, it is not in itself a distinct element. The nonpresence of the past and of the future is somehow bound up in the experience of the present, which cannot be fully distinguished from them. I think this is a cogent objection to Husserl. But also, from the other direction, Henry complained that the presence of immediate consciousness

---

[8] Dainton (2014: 112–115).

[9] For a thorough introduction to options, see Chuard (2020).

[10] A contemporary representative of this class in the representationalist tradition is G. Lee. I am a "weak atomist" in his terminology.

[11] Derrida (1973).

is not properly captured by Husserl's analysis simply because the primal impression is not sufficiently distinct in that analysis from the retention and protention.[12] To be fully present, it cannot have even as much to do with the protention and the retention as Husserl allows. I think there is also something to this objection, which in conjunction with the first objection seems to put Husserl's view in a troubling bind. A third problem is that the link between the retention of one moment and the primal impression of the preceding moment, and indeed the protention of the yet preceding moment, is crucial, but Husserl's model is not developed enough to explain it.[13]

But the Retentional view can be elaborated in such a way as to solve all three of these problems at once, and by using resources quite like those with which we are already familiar. We turn now from a focus on temporal phenomenology towards a way in which it is rooted in neurophysiology, in accord with the modal structural conception.

Recall that the experience of a certain shade of red involves hyperexperience of all the colors, and that the very being of an actual red such as we experience would be, if such a thing really existed, entwined with the mere possibility of all the colors, in a superworld. Experience of a primal impression also involves in a somewhat similar way hyperexperience of the relevant retention and protention, and an actually existing temporal moment of the kind of which the primal impression is an experience would involve the mere but still real possibility of the analogous objects of the retention and protention, entwined in a superworld.[14] But in this case, the less present aspects of hyperexperience are rather more present than in the case of color experience. They are closer to fully present experience in somewhat like the way that our mere hyperexperience of the now unseen backside of a car, or the part of a cow now hidden by a fence, are closer to full presence than a green that merely contrasts with the red of some seen cherry. And I think the same general sorts of fully actual neurophysiological structures that modal structuralism deploys to deal with mere hyperexperience of a cow partly shielded by a fence and the now hidden side of a car will also work in this case, without any questionable metaphysics of subsistence required.

Our neurophysiology is in a state that assures the real possibility of experience of the other colors when we actually experience red, and in

---

[12] Henry (1975: 139). [13] Dainton (2023).

[14] Although there are detailed complexities in this case, reflected in hyperexperience, that probably aren't captured by the simple superworld structure.

some important sense we cannot experience red without being capable of that other experience. This constitutes the neurophysiological basis for mere hyperexperience of the other colors. And our neurophysiology is in a more highly potentiated state, which is yet closer to full actualization, in the case of the hidden side of the car or the hidden parts of the cow. Likewise for the retention and protention that accompany a primal impression. And there is a corresponding sort of mere hyperexperience in both these cases, which is closer to full presence than in the case of hypothetical colors. Like an actually experienced color, the content of the experience which is the primal impression seems to rise up out of the lesser being of the content of the retention, but also fall down into that of the pretension, with which its very being is entwined. It presents a kind of unity less than identity with the object of the relevant retention and protention. But, unlike the case of contrasting colors, the object of the protention and the retention seem to have, while a less than full, still a quite substantial degree of being. They present themselves as unusually real possibilities yet less than fully actual. In temporal consciousness, there is an unusually strongly potentiated state of one's neurophysiology that would realize the experience that was the content of the immediately past experience which is now the focus of retention, and of the expected future experience which is now the focus of protention. The states in question are very close to actuality in the modally layered structure of our neurophysiology when the primal impression in question is had, largely because our neurophysiology has recently passed out of the activated state corresponding to the full experience of what is now merely a retention, and is about to pass into the fully activated state corresponding to the protention. And this constitutes the relevant mere hyperexperience. Notice that this all happens in such a way that one would expect the primal impression of the past moment to fit like a glove with the retention of the next moment.

This elaboration of Husserl's conception not only solves its characteristic problems but also plausibly provides insight into how the very basic Husserlian structure of time consciousness is rooted in our neurophysiology. Since that Husserlian structure is richer than the simple temporal order of events plausibly present in reality, where protentions and retentions have no obvious analogue, this qualitative feature of our temporal experience is nonveridical, and hence a poor fit with representationalist conceptions of sensory qualia. But it fits well with modal structuralism.

There is a further question about the nature of the protention in particular. As the nervous system enters the state that constitutes consciousness of the next primal impression, it most likely passes from a state that includes some of the necessary and characteristic preconditions for that next primal impression. It is unlikely that all those conditions arise instantaneously together. So perhaps the protention is constituted by those preliminary conditions. Or perhaps, on an alternative conception that we might call the "prediction model," it includes conditions that are a kind of prediction or anticipation of a future state, a prediction that might be falsified, and not merely by interruption before the entire sequence constituting the next primal impression is in place. That's the way I put it earlier, and later in this chapter we will consider some reasons to think that the second, prediction model of the protention is correct. But, on the other hand, some of the cases at the end of the next section may provide some support for the alternative, "preliminary conditions model."

Besides the very basic temporal structure of experience we have been discussing, it is also clear that there is a kind of experience of temporal duration, much as there is experience of spatial size and length. This is prone to subjective variation and distortion. The dinner may seem to drag on or be done in a flash. There are different views of how this might work. It has been proposed that there is a kind of basic quantum of felt duration that provides a metric for experiences, of which two dinners that each last an hour may seem to include many or few.[15] But it is more plausible that there are various ways in which felt durations of various lengths may be experienced without reference to a given fixed quantum, because any particular quantum we might pick seems phenomenologically arbitrary.

The neurophysiological basis for the distension and compression of felt duration is not well understood. But it is natural that information-based models have been developed, which require explicit neural representation, at a particular time in the same somewhat loose sense of "particular time" in which any kind of information-carrying neural firing pattern occurs at a particular time, of extended temporal durations. It is natural on such a conception to presume there can be differential neural representation at a single time of different temporal durations. For instance, there might be some neuronal comparison of the output of some incoming neural chain with a time-delayed version of the output of the same chain, comparing what is happening with what was happening.[16] And that comparison might yield a neural representation of the time delay involved. Despite the fact that these models are information-based, modal structuralism simply requires a somewhat different understanding of what any such duration-sensitive neurons do in constituting relevant modal filaments

---

[15] Merino-Rajme (2014).

[16] G. Lee (2014b: 16–20); Mauk and Buonomano (2004).

and hence qualia. These modal filaments might also gain temporal structure in such cases in a way analogous to that in which they gain spatial structure through differential motor activity.

Modal structuralism does provide a different account of the neural basis of the most basic form of temporal structure, as we saw previously. But it is also possible that even this basic form of temporal structure involves some aspects, for instance differences between protentions and retentions, that are refined by registration at a time. For example, there is experimental evidence that one can sometimes perceive two events as not simultaneous without a perception of the temporal order in which they occur, so perhaps experience of temporal order at least sometimes requires explicit registration of that order at a time.[17]

There is also the experience of more complex time-involving phenomenon, such as motion, velocity, and acceleration. Focus on the relatively simple case of motion. Information-based cognitive science has developed reasonably plausible accounts of explicit motion representation by neurons at a time, of the same general sort as those discussed in the preceding paragraphs. And such proposals can also be given a modal structuralist reading.[18] There is even some knowledge of relevant neurophysiology. Some retinal ganglion cells and corresponding cells in the LGN seem to be motion sensitive, and feed neurons that are more clearly motion sensitive in V1, which in turn feed other cortical areas that integrate input from across V1.[19]

2

There are illusions of motion that can tell us somewhat more about how our experience of temporality is constituted by our neurophysiology.

For instance, there is the waterfall illusion.[20] If you stare at a fixed rock jutting out from a waterfall for a bit, and then look off at some stationary area, you will experience the illusion of that area moving upward. This illusion is due to adaptation in the visual system, to its diminished response to constant stimuli. It has in this way a natural neurophysiological explanation.

There are also illusions involved in more basic aspects of temporal experience. For one thing, there is evidence that temporal experience is not, as it appears, temporally continuous. In human vision, the eyes saccade, jump from a focus on one region to another, about three times a second. During the saccade, there is active inhibition of vision. We don't experience the blur of motion that might be seen while the eyes saccade.

[17] Mitrani et al. (1986).

[18] Palmer (1999: 465–518).

[19] Ibid., 481–484.

[20] Ibid., 470.

But nor do we experience the gaps that occur due to this suppression.[21] How is this consistent with the modal structural conception? Partly because the level of temporal detail of nervous system activity, like the level of its spatial detail, that is noticeably relevant to differences in conscious experience is limited. Partly because there are limited degrees and forms of consciousness that even suppressed blurs, like blindsight, might plausibly enjoy. And in saccadic suppression, there is reason to think that the sharp and vivid images that arise before and after a saccade simply dominate and hence mask those that occur during the saccade.[22]

This phenomenon is related to another illusion of motion, called illusory motor reversal. Sometimes in a movie a fast spinning wheel with spokes will look as if it is rolling backwards. The explanation is that movies are discontinuous stills moving by at a certain speed, and when the motion of spokes fits that speed in a certain way, the nervous system "presumes" that the closest spokes in successive stills are one moving spoke, even though they are different spokes. But the most interesting thing about this phenomenon from our perspective is that an illusory motion reversal of this type can occur not merely while watching a movie but in actual vision of the turning of actual spoked wheels. And this is apparently because of a similar timing interaction with the "stills" of vision between saccades.[23]

There are also the phi phenomenon and color phi phenomenon.[24] In the phi phenomenon, successive appearances of a dot at different locations on a screen, when they are timed properly, can create the experience of a continuous dot moving between the two locations. In the color phi phenomenon, such an illusory traveling dot can even be seen to change color midstream.[25] There is filling in, as in the filling in that masks the blind spot in the eye. It is especially interesting that in experience it seems that the dot is moving and the color is changing in the proper direction before the end point of the change is reached. Another phenomenon of this general sort is the cutaneous rabbit illusion, in which taps are applied rhythmically to different parts of the arm, say five taps to the wrist and five to the elbow, but the subject feels that something jumps regularly up their arm, missing some of the actual taps but encompassing others.[26]

Dennett provided vivid names for two different ways these illusions might work.[27] Focus on the cutaneous rabbit illusion. On the

---

[21] Holcombe (2014: 181).

[22] Palmer (1999: 523–524).

[23] Holcombe (2014). And see his discussion of the jittery motion illusion.

[24] Mölder (2014: 220–221).

[25] Kohlers and von Grünau (1976).

[26] Geldard and Sherrick (1972).

[27] Dennett (1991: 116–117).

"Stalinesque" model, experience incorporates the illusion because there is a time delay between visual processing and experience, and so all the actual taps can be unconsciously processed before the relevant experience is created that smooths them out. On the "Orwellian" model, experience incorporates all the actual taps, but we don't remember it. Dennett claims there can be no empirical decision between two such theories, in service of the claim that there is no fact of the matter between them. Akins has proposed that an empirical test that might be relevant in the case of the phi phenomenon is intercepting the information processed before the experience but after the spots have been flashed.[28] But it is open to Dennett's fans to insist that this won't reveal the relevant content but will interrupt the normal processing.[29]

My view, while it eschews talk of information processing, suggests something of a compromise between the two models, namely that all the actual taps are experienced in some way, but not in a mode that allows for memory, while the smoothed-out illusory taps are more dominantly and memorably experienced.[30] Another relevant point is that the fully experienced taps suggest a time delay between relevant stimuli and experience, which may be specially relevant to the nature of protentions and may provide some support for the preliminary conditions model noted previously.

3

Our experience of time, with its retentions and protentions, is crucial for our experience of the robust particularity of visual objects. It is also important to the locking together of different moments of experience deployed in Chapter 3. We turn now to a second important ingredient of our experience of visual particularity.

Our neurophysiology of action plays a crucial role in constituting basic qualia even beyond those we considered earlier. I will first introduce the relevant qualia, and then explain how our action is relevant to such conscious experience.

---

[28] Akins (1996).

[29] Mölder (2014: 228).

[30] The proper treatment of some of these temporal illusions is arguably related to the following issue: According to special relativity, there is no absolute simultaneity and even no fact of the matter about the temporal order of events with spacelike separations, which some believe present difficult problems for accounts of temporal experience, in which there is apparently absolute temporal order and absolute simultaneity. See G. Lee (2007). But I believe that the differences in question are below the scale of detail relevant to any dramatic and noticeable features of our temporal experience. This is more or less to reject Lee's Premise 4.

First is a phenomenological claim: Causal powers are among the intrinsic properties of things present in our conscious phenomenal experience of ordinary concrete objects, although not quite in the way of colors or tones.[31] When you hit a ball, run into a tree, walk or row or swim into the wind, lift a heavy stone, peddle uphill, or try to push the negative poles of two strong magnets together, you experience causal powers of such things in a quite direct way.

Such experience is not to be confused with the difference between the trajectory of your actual visual or tactile sensory experience while swimming into the wind and such sensory experience you expect you would have if the wind didn't exist, nor with the passive experience of pressure and effort characteristic of certain aspects of touch. What I mean is beyond that. It is an experience of how things in the world, including yourself, are actively deflected, from a course they otherwise would have taken or a situation they would otherwise have maintained, by the specific action of other concrete things in the world, and also of how you must yourself actively push and rebalance to counter such action by those other things. It is an experience of how what might initially seem possible, the magnets coming together with little forcing on your part, is in fact not possible, because of the concrete way certain things are. Although these dynamical sorts of powers are a central case, they are not the only powers we seem to experience. The loose electric wires and the torpedo fish look thus and so, but when you grab them, you find that they have what might be a quite surprising power to shock.

Such a power, for instance the charge or mass or momentum of some concrete observed thing, would be, if it really exists, a kind of modal property that is yet located at a particular place in the world. We have seen that even colors as we experience them are a kind of localized modal phenomenon. But a power is another sort of localized modal phenomenon, different in form and modal structure from a color. Powers, at least when active, do things, make a difference, are a kind of restriction on the rest of the world. The powers of concrete entities somehow make the actual world beyond those entities be within a certain range of possible ways. They rule out certain mere possibilities so that actuality is within a limited range of possibilities that remain. They are restrictions on what is possible in reality, on real possibility in some sense, given that such a concrete thing with such a power exists. Different powers are different such restrictions. So, despite the fact that they are in themselves intrinsic

---

[31] For further development of this conception, see Mendola (2021: chapter 7).

and irreducible basic concrete properties, at least in our experience, these powers can be characterized and compared by regard to the contents of the restrictions that they make.

We are familiar by now with modal aspects of experience. But if such powers are actually out there in reality, then they would plausibly involve a kind of modal structure beyond that of the superworld. There are several complexities. For one thing, powers that would exert different restrictions on the rest of the world depending on the presence of various different background conditions, or are chancy in their operation, have possible outcomes that are in some sense real possibilities, but of a less ontically weighty sort than causal outcomes mandated by a power, outcomes that are certain and indifferent to the presence of extra background conditions even before such a mandated outcome is actually achieved. For another thing, it is clear that all these possible outcomes are more salient and closer to actuality than the very ethereal and remote possibilities of alternative colors relevant to the essence of some instanced phenomenal color. But one important aspect of powers of this sort might be captured within the superworld in this way: The superworld is structurally constrained such that the powers instantiated in it must be coherent with the rest of the world, such that relevant actual parts of the superworld must be as the instanced powers demand. There are so-called metaphysical possibilities beyond the superworld. But on this conception it would not be metaphysically possible for a power to be instantiated in reality and the rest of reality to violate the restriction it makes.

The experience of such powers introduces considerable complexity into hyperexperience. There is often in these cases a suggestion in experience of the truth of a counterfactual conditional. Such conditionals are claims like "If I were rich, then I'd buy a Maserati and drive away" or "If I had a river, I would skate away." For instance, if there is an unstably balanced object, you may have a kind of virtual experience as of what would happen if it were to fall.[32] And when a wave crashes in and slows your progress into the deeper water, there is sometimes analogous virtual experience of what would have happened if it hadn't crashed in. What's more, there is sometimes one kind of mere hyperexperience of a range of alternative things that might happen, or of one thing that certainly will happen, consequent on the presence of a power, and even hyperexperience of alternative things that would happen should still other things happen that

---

[32] Siegel (2010: chapter 5).

might happen. There are many complications, which I will not attempt to trace here.

We do not experience such powers in quite the way we experience green. But they are very concrete properties of ordinary particulars that we do in some recognizable sense experience. We experience concrete things in the world doing things, and the properties by which they do them are their causal powers.

Perhaps some follower of Hume will object that we experience no powers, only widespread regularities of things that have other concrete properties, and project a kind of expectation based on that experienced regularity onto the world. But remember that we are not now talking about what is truly out there in reality, but merely about what we experience, which that Humean story is an attempt to deliver. And besides, in many cases we have no conception of what other intrinsic properties of things there might be to enter into Humean regularities except their powers. The intrinsic properties of electrons as we conceive them are perhaps exhausted by what they do and can do, and even ordinary visible magnets have one crucial property of that sort. And in any case, Hume's specific story about how we come to such conceptions is implausible. The first time some young child careens headlong into a tree or drops a brick on their foot or burns their hand on a stove or even tries to push strong magnets together, there is no general regularity that they have experienced that can plausibly explain in suitable detail their current experience of such powers.

Two electrons, or even two bricks dropped on your bare feet, have causal powers among their basic concrete intrinsic properties, which are one sort of modal property. And in some cases we experience these powers. That is my first controversial claim about powers. But there is also a second, about how we have experience of such powers.

You can hit, push, walk, row, and swim with greater or lesser force. In pushing into a wave or the wind more or less forcefully, you do different things, in a way which can take account of differences in the way you expect that the wind or the water will push back, will resist your action. I think your experience of the force of the wind or wave is in part constituted by how forcefully you push against them, irrespective of any differences in the passive pressure you might feel in the presence of the wind and wave. Imagine two twins. They are duplicates in regard to aspects of sensory experience that can be captured without regard to causal experience, including passive pressure sensation. But imagine that the first acts with twice the force, into a world that is twice as forceful. I think the twins hence have different forms of conscious experience, constituted as different by the difference in their motor activity.

I further think that it is crucial to such experienced powers of humans that we have just the kinds of motor capacities we have, with muscles attached to a skeleton in such a way that we can move and apply mechanical force in three dimensions, and can apply different amounts of mechanical force. I think that if we had different basic action capacities, say capacities to change the experienced phenomenal colors of things at will and with different degrees of "force," as those things resist color change with analogous degrees of "force," we would have different sorts of causal experience than we do. We also experience powers, like those of torpedo fish, which are not of the dynamical sort we can ourselves directly exert. But the way in which we experience such alternative powers, as for instance involving electrical forces that are stronger and weaker, is linked to dynamical effects they have.

And I yet further think that this is not to say that it is precisely and merely what we actually do or could do out in the world that is crucial. Rather, as in the case of our experience of colors and tones, the kinds of neurophysiology involved in the control of our motor activity and hence the constitution of this causal experience is crucial to the nature of that experience as well, to the similarity and differences of what might be called the causal qualia that it involves. I don't think the world into which we act in fact has exactly the kind of structure we experience in this causal way, although that must be a long story for another day. I think the nature of that causal experience is partly constituted by the neurophysiology that controls our action. I believe that there are what might be specifically called "motor qualia," that our experience and phenomenal consciousness involves a kind of motor experience, experience of what we are doing, and in a particular guise that can introduce the differences characteristic of qualia between what is real and how it seems to us.[33]

It would of course be a large task to properly convince you that we have causal experience involving motor qualia, and that we have it because we act. I can't do that adequately here, though I have discussed these issues elsewhere.[34] There are many complications. We'd properly need to consider the effects of paralysis, beyond our earlier treatment. And it would be another task, beyond our current neurophysiological knowledge, to properly show how this is all reflected in the modal structure of our motor

---

[33] For more, see Mendola (1997: 143–165) and Mendola (2021: chapter 7).

[34] In the references indicated in the preceding note, and in Mendola (2008: 9.5).

neurophysiology. And there is neurophysiology dedicated to direct compensation for changes in load on muscles which is outside of any ordinary conscious stream between sensation and action, and would need to be treated by modal structuralism in something like the way it treats the neurophysiology involved in control of visual saccades according to the afferent copy theory. But those are details. What I have provided here is merely the outline of the story I would eventually need to tell in a convincing way to explain the experience of causal powers.

Fortunately, our current goal is something else, an account of the particularity apparently present in our visual experience. So, though I will deploy what I am calling a kind of causal experience in that account, you need not presume in what follows that either my tendentious description of these phenomena or my story about how we have such experience is correct. Hume's story might do just as well for our current purposes.

4

There are several ways in which motor activity, and the neurophysiology which controls it, are plausibly important in the constitution of modal filaments and hence sensory qualia. Consider the tendency to flee that is part of our experience of pain. Not just alleged motor qualia are at issue. So while not enough is now known to flesh out adequately the story about causal experience I just told, we should consider some relevant details about motor activity that are known.

You have at a time a specific range of motor capacities, capacities to move your limbs and body and tongue and mouth in various ways. They are constituted by the muscular-skeletal structure inside your skin and the rough features of the neurophysiology that controls it. What can be said about that neurophysiology?

First, it is at least reasonably clear and uncontroversial that the neurophysiological control of human action has a roughly hierarchical form.[35] There is a rough division between motor planning and execution. The association cortex generates what amount to action goals. It then influences the basal ganglia, the supplementary motor area of the cortex, and the premotor area of the cortex, which execute further planning functions such as specific movement selection reflecting the goals, also incorporating in that selection learned patterns of action and external stimulus monitoring. Such areas also perform such functions as the proper control of

---

[35] Gazzaniga, Ivry, and Mangun (2002: 445–498); Jeannerod (1997); Rosenbaum (1991).

switching between movements. The lateral cerebellum is involved in preparation of specific motor patterns. After such a planning phase, action execution begins with the activation of muscles by the motor cortex, with the participation of the intermediate cerebellum, which helps implement and correct movement patterns.[36]

There are of course many controversies and complexities regarding this rough story. For instance, the hierarchical organization is of a surprising kind that allows that we retain characteristics of our handwriting even when writing in media that require very different fine muscle movements.[37] And the traditional view that the cerebellum is involved principally in motor control is challenged by the fact that it is also involved in judging the timing of sounds or differences in pitch, the movement of visual patterns, and differences between similar sounding words, and that loss of the cerebellum yet allows eventual recovery, particularly in the young, of a large degree of movement coordination.[38] But we can presume that something in the rough vicinity of this common understanding is correct. As modal structuralists, of course, we must interpret it in a certain way, not in the traditional way as involving representations of different sorts of motor activity, but as helping to constitute modal filaments that are constitutive of relevant qualia.[39]

A second relevant point, though it is not as well confirmed as the general hierarchical organization just noted, is that there appears to be what we might call a kind of horizontal organization of motor neurophysiology as well, involving the kinds of contrasting alternative possibilities that we have come to see as crucial to qualia similarity and differences. Much of what happens at specific layers in the hierarchical structure seems to involve such contrasts between alternatives, in a way characteristic of modal structuralism. There is also some evidence that the entire interconnected musculoskeletal network is controlled and balanced as a whole by neural structures in the primary motor cortex.[40]

To sum up, what little is known about the neurophysiology of action comports well with modal structuralism.

5

We have been considering various precursor resources for our goal that are also important in their own right. We also have available the account of

---

[36] Gazzaniga, Ivry, and Mangun (2002: 497).

[37] Rosenbaum (1991: 229).

[38] Bower and Parsons (2003).

[39] This is contrary to Rescorla (2016).

[40] Murphy et al. (2018).

rudimentary visual particularity from Chapter 6, which is another crucial resource. But now we arrive at the central topic of this chapter.

Say you watch a white station wagon back out of the garage, or see a blue truck rush down the road. Visual experience of the robust but commonplace particularity of those vehicles involves complex aspects of hyperexperience.[41] Those aspects are not all fully present in experience, but they weave together aspects of temporal, spatial, and causal hyperexperience, and motor action, in a complex way.

Chapter 6 discussed a very simple form of two-dimensional spatial particularity, delivered by momentary modal filaments. There we didn't even properly consider somewhat more complex two-dimensional cases, stationary modal filaments that present a color bit as persisting at the same location in your visual field over time, and flexible modal filaments that present a color bit as moving in that field and as indeed capable of other movements than those it in fact makes. And there may be the experience of extended visual shapes moving and changing in your visual field as well. But now we will leap straight to a still more complex form of visual particularity, in hope it provides some indirect illumination about these lesser cases as well.

The visual experience of three-dimensional objects is a very complex matter.[42] But it is clear that we see ordinary visual objects like trucks and houses as having various three-dimensional geometric and mereological properties, not just as bare particularities.[43] It is clear that the visual particulars that we predominantly experience are bounded and cohesive, in the manner of cars and balls, but that we can also experience swarms of bees that move together, and also holes, as having a kind of visual particularity.[44] And it is clear that we can experience such visual particulars to persist despite change, not only in their positions but also, to some degree, in their shape and the spatial configuration of their parts.[45] Just watch someone walking.

In the representationalist tradition, it is now customary to say that these phenomena involve a specific "file" for representations of a particular such environmental object and its features,[46] though it could be clearer what that means. If there are specific neural states that might be plausibly claimed to constitute such files, then they could alternatively provide a neural anchor for a modal filament, itself a collection of other modal

---

[41] For more on the apparent metaphysics involved, see Mendola (2021: chapter 5).

[42] For an introduction, see Palmer (1999: 199–407).

[43] E. J. Green (2017).

[44] E. J. Green (2018).

[45] E. J. Green (2019).

[46] Recanati (2012).

filaments, that would be experience as of such a visual particular. But such a local anchor isn't necessary according to modal structuralism. And since such a modal filament is internal to the person, and could be connected to a variety of different real objects in the world or none at all, the kind of particularity that is directly present in experience cannot depend on the particular object out in the world with which that experience happens to be connected. To get it linked with one particular object in the world, more complications are required, in a way that is probably contrary to the information-based tradition, since the complications plausibly involve more than just an informational link with one particular object. And this generates many issues about how two people, or even a single person over an extended period of time, can share thought or experience of one particular environmental object. But, fortunately, most of these complexities are beyond our remit here.[47] We will simply be concerned with the way in which modal filaments can constitute experience of the kind of particularity visual objects often present in our experience, whether or not they are really out there or anyone else sees them.

The visual experience of a moving object at a time includes a protention, though it is possible that it is the kind of protention that involves a prediction that might be falsified. We expect certain sorts of temporal progressions of seen objects in general, and are surprised by others. Despite the holes and swarms of bees noted, we generally expect such objects to follow continuous trajectories, to continue to exist when occluded or unperceived, to be unified and largely impenetrable, to be able to causally influence other objects, to have motions explicable by simple mechanical forces and a mass or weight that will lead them to fall if unsupported, and to have a roughly determinate shape or range of shapes that will only change in continuous ways.[48] And we have yet more specific expectations regarding specific objects, for instance our own idiosyncratic cars.

Consider the changes over time in visual experience of such seen objects. We always see something from some distance and direction or other. One's immediate visual experience of a concrete object has a place in a range of visual experiences of that thing one would have over time if one moved in various ways, and one ordinarily has an expectation of how one's visual experience of something would change if one moved in various ways. To each object that one currently experiences with a particular shape and size and orientation, and with a particular spatial relation to one, there corresponds a series of visual experiences one would have of the object if

---

[47] But see Mendola (2008) for hints.

[48] This list is adapted from Bermúdez (2003: 82–83).

one moved in various ways, along various paths in the world. And one often has at least a rough expectation of how all that would go, due to general familiarity with the way ordinary objects in the world work, and sometimes to specific familiarity with particular objects. This is so even if the object is not within one's immediate experience, assuming that the object is stable and unmoving over the relevant temporal interval. You can be familiar with how those things in the next room would look if only you went in there. And you are also in some sense aware that there is a different series of visual experiences which you would have of the object if it itself were moving along some path.

I think that one's sense of the various paths which some object or oneself might take through the world is constituted by two things, in ordinary cases. Such paths might appear in for instance visual experience from various vantage points, as visually continuous spatial paths that could be traversed by objects over time. But one's sense of such paths is also one's sense of how one or an object would travel if it or one did various things, or merely coasted along. For instance, one's sense of a particular path forward from where one is in part one's sense of what would happen if one performed a series of motor operations, running forward and turning slightly left. And this is also a path by which one expects that causal influence can pass through the world, of the trajectory a ball or a planet might take, perhaps by bouncing for the left turn. What's more, we can zoom in close to things and see more spatial details than we otherwise would, or pan out and see a larger part of the world. In imagination, or in reality, we can zoom, pan, and rotate about. In all of this, there is an important interaction of motor activity and visual experience, or more exactly of possible motor activity taking into account any causal qualia on one hand and possible visual experience involving visual qualia on the other.

Some of these combined trajectories of vision and motor action actually occur. But many are reflected in your expectations only. But there are ways in which such expectations are reflected in conscious experience, at least if what we decided about the nature of hyperexperience earlier is correct. Conceivably, instead this all might involve largely a kind of disciplined primary imagination, with a link between the primary imagination of certain individual movements and of certain visual experiences. But it seems to me that, as we have considered in analogous cases before, not fully present parts of hyperexperience often involve these same phenomena. When one has a visual experience of a car, one has one kind of mere hyperexperience of the range of visual views of the car one might have if one did various things, rooted in a rather robust potentiation of both visual and causal experience not completely disanalogous to that involved in the protentions and retentions of temporal experience, although somewhat less robust. And even when an object is not in one's current visual experience, one has a sense of what one would see of it if one moved thus and so which might conceivably involve a more extended sort of hyperexperience. Indeed, one has such a sense of how everything in the world as one conceives it would look altogether, irrespective of its organization into particular objects, although such objects seem to be important anchors that organize one's sense of it all. So maybe all the relevant hypotheticals are reflected in one's mere hyperexperience of the world as a whole, as in one of the more radical models considered earlier but not adopted. But any of those models, including the more moderate and I think plausible ones, involve a complexity in the structure of hyperexperience we have not sufficiently developed before, which involves not just different degrees of presence but a linking between specific motor and visual aspects of experience in our experience of a three-dimensional particular.

As far as I know, little is known about the way our neurophysiology in fact roots such modal structure. But Mel showed in the 19080s that a parallel-distributed-processing-style neural net, utilizing Hebb-style learning rules, could come to reflect the way one's vision of a three-dimensional world would shift for a given shift in one's vantage point, in effect by associating motor commands that move the net in various ways through the world with particular changes in the view those moves would induce.[49] And of course computer-aided imagery has come a long way since then. And in rats, neural activity in the hippocampus and medial entorhinal cortex is not only correlated with location, head direction, speed, and proximity to boundaries in spatial navigation, but also is more generally correlated with continuous, task-relevant factors in other sorts of behavior.[50]

We noted in Chapter 6 other ways in which even the most rudimentary spatial structure present in visual experience involves correlations with motor activity, and it is clear that such correlations are quite significant in still other ways as well. Consider for instance visual size. If your truck has the picture of a big faucet painted on a door, from exactly what vantage point do you see the true and objective visual size of that picture? If any vantage point is to be privileged in this way, it would seem to be when you are right on top of the truck, when you have your eye right at the image. But when you are so close to almost anything, its visual size fills your entire

---

[49] Mel (1986). For more, see Mendola (1997: 313–322).

[50] Aronov, Nevers, and Tank (2017).

visual field and overflows out beyond it. So it would seem in fact that things do not have any objective visual size, or at least none that is experienced through vision. Still, in motor activity and touch, we seem to be in direct contact with the true sizes of things. The conception of an objective spatial object with an objective size apparently must be the conception of something whose objective size is available at all only through another sense than vision, unless perhaps motor activity and touch contribute to even visual size in some way we haven't yet considered, or visual spatial experience involves a more absolute registration of size than we considered in Chapter 6, which however would not fit easily with the obvious fact that things look smaller as they recede from us.

There are many issues regarding objectivity of one sort or another involved in the matters now in play. We have been discussing experience and its qualia. But leaving the more radical, and I think implausible, models of hyperexperience aside, we are certainly capable of thoughts of relatively complex concrete scenarios, with three spatial dimensions as well as a temporal dimension, and with various objects and properties in various locations, that plausibly outrun anything that would fit within the experience of a moment.[51] We are capable of conceiving various paths through such a world part, and various persisting concrete objects that take various paths. And no matter how radical a form of hyperexperience we presume, there is in any case a gap between the kinds of hyperexperience involving vision and motor activity and touch we have been discussing, and these objective scenarios, even beyond those involving the other senses we aren't currently discussing. That gap is this: Vision and visual experience always involves a vantage point, which fails to capture the back sides of things in the same way as the facing sides, and views them only from a certain angle and distance. But the objective world is independent of any particular vantage point. Call visual experience of something from a vantage point a "look." There is not a very exact fit between the content of a look and any plausible objective structure in the world itself. And hyperexperience involves various looks in particular.

It is true, we have noted, that the experience of colors involves a palimpsest of spatial structure, so that a color look that is experienced as visually close to another color look in your visual field can at the same time be experienced as the color of an external object that is far away. But there are other forms of objectivity that do not appear within experience in the same direct way as that second color. Hyperexperience flows out of and is

---

[51] Mendola (2008: 214–215).

anchored in the actual current and fully present experience of the person who enjoys it, and that actual experience involves looks. Also, what I merely visually hyperexperience of the backside of the car is what I expect I would see if I walked behind it. It too is a look. It is not fully objective.

It is even not unreasonable to complain that the very looks of which I am speaking have objective aspects that misrepresent hyperexperience. The looks in question may also seem to capture our sense of what someone else might see in a similar position, or even a kind of objective vantage point on things that could be occupied by many different individuals or by none, who would see the same thing if they were there. Relevant motor aspects present in individual hyperexperience may also have analogous objective correlates that are similarly important. And it is somewhat doubtful that a palimpsest imposed on our experience can perform the miracle of entirely transubstantiating immediate visual experiences into a plausible visual reality as seen from a given objective point, or the experience of another person, since the fact that it is our experience that contains the structure of looks will still remain within any such palimpsest. Perhaps these worries are important. But even if they aren't, even if as I think the looks in question are quite subjective and the gap with the objective world they involve quite large, we can somehow come to consider these experiences in other and somewhat more objective ways, for instance as experiences that might be had by others in a world in which we do not exist.

The complexities are vast. But there is much in Kant and his commentators that may help flesh out the relevant objectivities.[52] Strawson read Kant's Refutation of Idealism as suggesting that a certain sort of self-consciousness required the conception of objects as distinct from oneself, that the self and those other things are in that way correlative.[53] To be aware of oneself as having certain experiences seems to require a capacity to be aware that sometimes one's experience is as of objects distinct from oneself, and that could exist unperceived. Jonathan Bennett read it as rather showing the need to distinguish between one's memories of something and the facts of it.[54] Some readings of Wittgenstein's Private Language Argument stress the importance of the capacity to understand the possible difference between one's own experience and that of other people.[55] John Campbell suggests that the phenomena noted by Bennett and Strawson are connected with the possibility of thinking of oneself

---

[52] Ibid., 215–217.

[53] P. F. Strawson (1966: 125–132).

[54] Bennett (1966: 202–218).

[55] Mendola (2008: Part III) discusses various understandings of the Private Language Argument.

from the outside, as an object that moves through the world.[56] In the *Opus Postumum*, Kant suggested that a proof of the objective reality of the categories requires the subject's awareness of their own corporeal nature as a self-mover exerting force and subject to external forces.[57] There are many interesting and difficult issues. But I will focus here on just four specific phenomena in this general vicinity that are specially relevant to the kind of particularity present in our visual experience, and which plausibly can be present in relevant looks and so hyperexperience even if other forms of objectivity cannot. We yet know little about relevant neurophysiology. But at least we can get in position to see that such experience could be constituted in accord with modal structuralism, by internal neural resources only, that information-based representation is not crucial.

6

Begin with cases like this: You experience some speeding car as the same thing despite its changes in position over time. You also have a sense of how it could be seen to move or alter shape and yet remain the same thing. It isn't tied to its current position over all of time, to where you now see it, so it can possibly exist at other locations at other times. It certainly could exist at the positions it will later occupy. But it isn't even tied to any of its actual future locations, since it could be seen to go in other directions and yet remain the individual car you see now. You could watch it as it does so. What's more, although it did take a certain trip through the world to get to the place where you currently see it, which you may have watched closely, still it could have taken a different route than that to get where it is, and yet have been the very thing you currently see.

As you concentrate on that car you see right there in front of you now, you can yet consider how it might in the past and future have occupied or occupy very different positions than it actually has or will, while remaining the same particular car. Your focus on that one particular car now remains a sufficient anchor for these hypothetical conceptions.[58] The car's individuality as you experience it now is independent of its actual past and future trajectories in this way, although you can have a hypothetical form of primary imagination of possible past and future trajectories it might have taken to here or might take out of here, and also plausibly sometimes have analogous forms of mere hyperexperience.

---

[56] J. Campbell (1999).

[57] Förster (2002).

[58] For the idea that attention is an important root of reference, see J. Campbell (2002).

That is our first step towards understanding the particularity of the car we visually experience. But there is also a second step nearby. The car might not only be in some other place in the future than it actually will be, or have been in other places in the past than it was, as we have already noted. It also might have taken another route to that merely possible future position than any that would pass through where it actually is now. It might never have been where you see it now at all. You focus on it through seeing it where it is now, but can consider many different places where it might be in the future, and also many different ways it itself could get to that hypothetical future place and yet be the very car it is, some of which never involve being where it is now. And you can also run such a series of hypothetical scenarios working into the past rather than into the future. What's more, at least aspects of such a complex modal structure are, I think, not just available to you in experience through primary imagination, but also in your mere hyperexperience whenever you watch a speeding car and have a visual experience of its peculiar particularity. These modal phenomena are highly relevant to the kind of particularity that the car presents to you even in your current visual experience of it, which is a richly modal kind. This is due to that particularity partially reflecting quite hypothetical aspects of your potentiated experience, which still sometimes suffice for a rather evanescent and ethereal form of hyperexperience.

So-called haecceities are metaphysical beasts introduced by Duns Scotus, bare particularities with a very robust modal nature. Say there are two otherwise identical model rockets in front of you, one to the right and one to the left. Still, each rocket might have been in the position of the other. And if one had been in the garage, or the other, those would have been two different possible situations. What can explain this? The rockets share all ordinary properties except position, but they might exchange positions. So we might conclude that each rocket has a specific haecceity, which distinguishes the two here in front of you, and which each of them carries with it into different hypothetical scenarios.[59] The forms of experience we have been discussing are beating on the doors of the presence of particularities in our visual experience as robust as those delivered by haecceities. But notice that we have been merely making new use of capacities crucial to our visual experience of persisting objects over time, and of three-dimensional objects that look different if we travel in various ways around them. And notice also that the experience as of such robust

---

[59] See Mendola (2021: chapter 5).

particularities, of haecceities if you will, does not at all imply that they exist in reality, on this conception. The kind of particularity present in your visual experience of the car might be an illusion of that experience, due to the psychological capacities just noted, and any link to one or another real car might be merely through that car's current position, where you see it now.

With haecceities we entered into our second step towards understanding visual particularity. But to understand this step properly, some background may be helpful. There are two natural axes of contrast between accounts of the particularity present in our experience.

Some hold that the basic particularity in experience is that of objects, that experienced properties are attributed to objects. Others hold that the basic particularity in experience is that of locations, that experienced properties are attributed to locations. In regard to visual experience, on this axis we have been splitting the difference by deploying a palimpsest. Colors are present as at visual locations, but also as properties of robustly particular external objects of the sort that now concern us.

But another axis of contrast is between those who hold that really existing external objects or locations contribute the relevant particularity, and those who hold that experience is merely as of such particularity, which might for instance also be present in experiences that are misperceptions, dreams, hallucinations, or imaginings. We might call the first party "particularity externalists" and the second "particularity internalists," since the first camp thinks that external objects out in the world contribute the relevant particularity to our thought. On this second axis of dispute, I am of the second, particularity internalist party.[60]

In various instances of the phi phenomenon, flashing lights in different locations are seen as or not seen as continuous moving objects. They can even be seen, when the lights are different colors, as the presentation of moving objects changing colors. And no actually existing and persisting external objects need be involved. There is one flash on a screen and then another. There is experience as of a certain kind of persisting particularity in the absence of that persisting particularity in the world. This is an intuitive example, along with other examples involving dreams and hallucinations and imaginings, of robust particularity in experience without such particularity in the world.

I hope these ways in which our visual experience of the particularity of ordinary objects like cars is independent of, yet entwined with, our

---

[60] Austen Clark (2000) favors actual locations, and Matthen (2005) actual objects.

experience of their locations is now sufficiently clear. We can now take the next step in our progress, which introduces yet a third aspect of visual particularity, by considering a prominent contrast.

Siegel has developed an influential account of the particularity present in visual experience worth our consideration.[61] The visual experience you have when your friend Francis looks sad to you involves both, in her view, singular content that can only be made true by states of Francis, and also nonsingular content that could be made true by any Francis-duplicate looking sad. Seeing Francis and seeing the duplicate would involve the same phenomenal character, closely related to that nonsingular content, she believes. Yet such phenomenology and such nonsingular content involves a kind of visual particular, which might also be present in a hallucination.[62] This is not quite, as far as I can see, the kind of particularity internalism I favor, since Siegel claims that hallucinations could not have robustly singular contents,[63] apparently not even so robustly particular as to distinguish hallucinatory objects appearing to the left and to the right of you, a point to which we will return. But she is at least a fellow traveler on issues of particularity internalism. And it is another aspect of Siegel's view that is most important for us here, and with which I am in complete agreement.

Siegel thinks that visual experience can represent, or I would say present, the subject-independence of objects, or even apparent objects. This specifically involves, she claims, two principles. The first is SI: If the subject changes her perspective on the object, then the object will not thereby move. The second is PC: If the subject substantially changes her perspective on the object, her visual phenomenology will change as a result of this change.[64] But the important point for us is that her claim is supported by the Doll Case.[65]

There is a doll that apparently "moves with movements of your head as if you were wearing a helmet with an imperceptible arm," and in addition "seems to move with your eyes as well" even when you keep your head still, and even is such that "when you close your eyes, you continue having a visual experience as of a doll," and finally is such that when you try with your eyes open "to put an opaque object right in front of the doll to block it from your view," it continues to be seen.[66] Creepy! According to Siegel, this strange doll does not have sufficient independence from you to present

---

[61] Siegel (2010). [62] Ibid., 141–174. [63] Ibid., 153–155.

[64] Ibid., 178–179. And see page 197 on filling out the consequent of PC. [65] Ibid., 184–185.

[66] Ibid.

the relevant sort of particularity that is present in your visual experience in normal cases. Seeing this scary doll is more like seeing stars when you are hit on the head or phosphenes than seeing a normal visual object. The doll case involves a lesser sort of particularity phenomenologically than is present in an ordinary case, Siegel claims, and hence involves different qualia.

I think this is right. It might be objected that the kind of experience in question seems too complex to be the immediate kind of visual experience we are focused on. But we have already considered a number of ways in which at least hyperexperience is modally rich and complex, in fact some that are quite like the one now under consideration.

There are other cases in this general vicinity that may seem problematic for me.[67] Searle has suggested that there is a difference in visual phenomenology when we perceive something or just imagine it, because, in the first case, the object seems to cause our experiencing.[68] Matthen suggests that pictures and imaginings and dreaming do not involve the kind of content provided by motor-guiding vision in normal cases, and hence do not involve a feeling of presence as in actual seeing, are not "asserted." And he thinks that no demonstrative thought about the individuals in such cases is possible, with the exception of thought about persons because of our capacity for face recognition.[69] But I reply that in hallucinations and dreams, all of these factors might be present in experience, and allow a kind of demonstrative thought, and that in experience of pictures, at least the phenomena Matthen cites might be present. I think in this way that even in the absence of real particulars out in the world, we might have experience as of the robust particularity that visual objects may present to us in normal cases.

But to get all the way to that particularity, we have at least one more step, a fourth step, to take. I think there is a richer modal aspect of the visual experience of even such a creepy doll than Siegel allows, and also indeed a richer modal aspect of any hallucinatory doll that behaves more normally. This allows an hallucinatory doll to be distinguished from a second qualitatively identical hallucinatory doll experienced off to the side, and even supports the possibility of a different situation in which the two dolls are interchanged in experience. And such a structure is also relevant to ordinary perception. If there are two actual identical-looking cars now in front of you, you can focus in some sense on both through your current

---

[67] Consider also Farkas (2013); Kroon (2013); and Masrour (2013).

[68] Searle (1991: 184); Siegel (2010: 195 fn).

[69] Matthen (2005: 320–321).

visual experience, and then consider, in the way we considered earlier, how they might each occupy various hypothetical positions in the future and in the past, until you end up with a scenario where they are now exchanged in position, yet retain their individuality. You can also focus on two cars that you merely imagine are now to the left and right of you, and perform the same operations.

One set of cases that evoke the right intuitions here are Austin's two tubes cases.[70] Someone looks through what they take to be a binocular microscope at what they take to be two spots on a slide, which are in fact one spot on which they are twice focused because of some deceptive fiber optics. I think they can conceive a different situation in which those two spots are switched. And that sort of modally rich particularity is I think generally present in our visual experience, even though in the tubes case it involves mere illusion. Alternatively, someone might look at what are in fact two spots which they experience and take to be but one. I think that in that case the real modal differences between the particulars that they see are not present in their experience. Sometimes we experience modal illusions, and some visual experiences of particularity involve them. In fact, it may even be that the modally rich forms of visual particularity we have been discussing here, which require metaphysical resources as strong as haecceities, are quite generally mere illusions of our experience.

In the absence of current knowledge of the neurophysiology relevant to some of the later steps in our progress here, we have been making a largely defensive case, that internal resources like neurons should be able to account for our experience of robust visual particularity in accord with modal structuralism. But there is also some evidence that some neurophysiological differences would constitute differences in such experience, which is a prediction of the view. Consider the sorts of individuality that seem present in our experience of sounds and smells. There is a rudimentary spatial particularity of sounds, involving a direction from the perceiver and perhaps a rough distance. But another sort of particularity may present itself when we hear two notes that emanate from different places as part of one chord.[71] And there is the kind of particularity distinguished by what is called monaural auditory scene analysis,[72] so that we can hear a continuing tune as distinct from what momentarily masks it. There is also a somewhat more robust particularity that adds together structures of these sorts, so that for instance we hear as distinct two otherwise identical sound streams that are over here and over there. And while most such sounds are heard to

---

[70] Austin (1985). [71] Matthen (2005: 286). [72] Bregman (1990).

occur at a particular location, say where the bell rings, there are also sounds that can be heard to move through space like rolling thunder or echoes. Odors sometimes present themselves as roughly at the nose. But by moving around and sniffing, we can follow an odor through the air, or experience it lingering after its source is removed, although its location in such a case is only somewhat vague and unspecific. And in a somewhat similar way, we seem able to experience such a smell concentrated at and emanating from its environmental source, which can indeed be hence experienced to be the thing that smells, such that the smell in question is *its* smell and a qualitatively similar smell belonging say to another sneaker is a different one, two particular smells which it wouldn't make sense to say were switched among sneakers in some other possible world.

Both sounds and odors involve palimpsests of different sorts of experienced particularity than those present in vision, due it seems to different sorts of sensory neurophysiology.

CHAPTER 8

Summary and Objection

So here we are:

1

The MOUDD theory of the phenomenal consciousness associated with our sensory experience – MOdalized Up, Down, and in Details – has three main components. The first component is a modalized conception of neuralia like us with unified nervous systems, and itself has three parts: (i) Our phenomenal consciousness is crucially that of agents, in whom possible stimulations and possible actions are linked in various possible ways. (ii) Our phenomenal consciousness is crucially that of a living organism on earth, constituted by living cells of familiar types, and involving both metabolism and replication, which also involve forms of modal structure. (iii) Our phenomenal consciousness involves the complex modal dynamics of our neurophysiology. In particular, it often matters that actual neural firing occurs in a context in which other sorts of firing are really possible given a relatively fixed state of our neurophysiology. Our exploration of this component involved pressing up from below, understanding the physical basis of phenomenal consciousness by probing its modal structure.

The second main component of the MOUDD theory is the whole nervous system model. On this view, what constitutes our phenomenal consciousness is all of our nervous system that connects sensory receptor cells and motor action, conceived in an extended way as including the muscles and skeletal structure crucial to our action. Our development of this model involved pressing down from above to understand the modal structure of phenomenal consciousness itself, and in particular the existence of complex hyperexperience, various forms and degrees of phenomenal consciousness, and the way in which the state of our nervous system is

https://doi.org/10.1017/9781009514774.008

relevant to what we would experience under merely hypothetical conditions of stimulation.

Pushing at once down from above and up from below, we can see that our phenomenal consciousness is a complex modal structure that is constituted by the complex modal structure of our nervous system, that is identical to that complex modal structure of our nervous system when it is considered at some appropriate level of abstraction. The third component of the MOUDD theory involves detailed explanations of the modal structure of sensory qualia by appeals to the modal structure of relevant sensory neurophysiology. More exactly put, the modal structure of our sensory neurophysiology explains our conscious sensory qualia because the actual modal structure of that neurophysiology explains the apparent modal structure of those qualia. For the most part, in reality there are no properties of the sort that our qualia apparently present to us; they are illusory.

2

But there are dualists who insist that any such physicalist story about our phenomenal consciousness must be mistaken. They usually object because of a cluster of familiar arguments to which we should attend.

It seems to many that there is an "explanatory gap" between any physicalist account of our nature and what it's like to be us,[1] as shown by these alleged cases: (1) We can understand everything there is to know about the physically constituted neurophysiology and biology of a bat and not know what it is like to be a bat, that can experience through a kind of sonar things we cannot imagine hearing.[2] (2) We can conceive that there is something which is our exact physical duplicate but which it is not like anything to be, whatever it insists, which is called a philosophical "zombie." In this way, the physical can be conceived to be distinct from the phenomenal, and some think that what can be conceived to be distinct is distinct, at least subject to certain exceptions that don't apply in this case.[3] (3) We can conceive that there is something which is our exact physical duplicate but is a "color invert," for instance who has an experience of a unique red quale when we have an experience of a unique green one, and indeed has their entire experience of color qualia flipped along the blue-yellow axis of the traditional color wheel. (4) Mary, raised in a black-and-white room, can come to know everything there is to know about the

---

[1] Levine (1983). [2] Nagel (1974). [3] K. Campbell (1970); Kirk (1974).

physical basis of seeing red by working hard on her (unillustrated) neurophysiology texts, without knowing what it is like to see red.[4] Inside the room, she knows everything there is to know about the physical basis of seeing red, but she has to leave the room to find out what it is like to see red. So that seems to be a new fact, beyond all the facts she knew inside, which is not constituted by the physical, since inside she already knew all the relevant physical facts.

Each of these arguments depends on the existence of an epistemic gap between physical reality and facts about what it's like, between the physical and the phenomenal.[5] Each argues from the existence of this epistemic gap to the existence of a metaphysical gap. The arguments hence share a general form.[6]

This form is controversial. Through the influence of Kripke and Putnam, it is now widely believed that while H-O-H is identical to water, and Hesperus to Phosphorus, and while such identities are necessary, so that there are no metaphysical gaps in these cases, still there can be epistemic gaps that lead someone to believe that there is water in a glass and not H-O-H, or that Hesperus has risen but Phosphorus hasn't.[7] That may well suggest that these dualist arguments have little probative weight. Still, there are semantic theories that drive a relevant wedge between these less controversial cases of identity or constitution and those involving the physical constitution of phenomenal consciousness.[8] Some dualists favor these. Yet, on the other hand, such theories are quite controversial, and there are popular alternative theories that treat all these cases in a similar way,[9] and hence drain the dualist arguments of their force. So it is probably best in this short space for me to avoid appeals to controversial semantic theories, and just face the dualist cases head-on. And that is what I will do. But still, for clarity's sake, here is what I think about this issue. With authors such as Chalmers and Nagel,[10] I believe that the constitution of water by H-O-H and the constitution of our phenomenal experience by our neurophysiology are different sorts of cases. In the first case, there is a kind of semantic entailment from the nature of H-O-H and its playing a certain role in our world to its being water. The identity of water and H-O-H is a posteriori but necessary, but the a posteriori aspect is merely that it was an empirical discovery what the microstructure of water happened to be, from among many possible microstructural candidates which might

---

[4] Jackson (1982).

[5] Chalmers (2010: 103–139).

[6] But see Robinson (2016: 15–16).

[7] Kripke (1980).

[8] Chalmers (2010: 41–205).

[9] Block and Stalnaker (1999).

[10] Chalmers (2010); Nagel (2002).

have entailed the manifest nature of water. But with Nagel and unlike Chalmers, I still think that our neurophysiology necessarily constitutes our phenomenal experience. I think there is a kind of entailment from having a neurophysiology like ours to having phenomenal experience like ours, though it is not analytical or definitional in the manner of the water case. Nevertheless, unlike Nagel, I do not think that this requires that there be some sort of new concept of what underlies both the physical and the mental that explains the relevant entailment. The somewhat novel concept of modal structure plays a different role here. I think that our understanding of our physical neurophysiology is already sufficient to see that it semantically entails the existence of relevant modal structure, and that we can see on a mixture of empirical and conceptual grounds that that structure must be sufficient to constitute our phenomenal experience. I do think with Nagel that this requires a new conception of the elements of mental life, but the notion of a modal filament provides it. Another disanalogy with the case of water, as I see it, is that there isn't a simple but necessary theoretical identity between a certain specific type of physical state and a certain type of phenomenal state. With the functionalist, I think that what constitutes your phenomenal experience might be somewhat physically different from what actually constitutes it while yet the phenomenology remains the same. And I think that things quite different from humans may be at least in some extended sense phenomenally conscious, since we are only the paradigmatic case. To suitably justify such a view would be a lengthy matter, but this short discussion may be helpful in understanding the nature of my claims here.

But now on to the cases: Fortunately, we aren't talking here about bats but only humans. And I claim, by pushing at once both down from above on the nature of phenomenal consciousness and also up from below on the nature of our neurophysiology, by exploration of both their modal structures, to have closed the explanatory gap to at least some degree in the case of humans, to have made it comprehensible how our physical structures constitute our phenomenal consciousness. Our discussion of color qualia in Chapter 4 was partly organized around qualia inversion cases. But there remain Mary and the zombies.

Zombie cases seem of quite limited argumentative weight. We can certainly say the words that describe them. But it is easy through the mediation of words to speak also of round squares and reddish greens and even water constituted by XYX and not by H-O-H. Once we turn to concretely imagining in a more serious way all the details of such cases, it is far from clear that they are available. I don't think you are really capable of thinking coherently that some living human animal in front of you screaming from some serious injury is not really feeling pain. If you are talking someday to two identical twins, real ones who aren't even physically identical in all respects, I don't think you will be capable of coherently thinking that it is like something to be one of them but not the other. So I will focus on what seems to me a more persuasive case, namely Mary in her black-and-white room.

3

The resistance I favor to dualist arguments based on Mary has close affinities with what is called "the phenomenal concepts" strategy[11] and also with what is called the "knowledge by acquaintance" strategy.[12] But it has idiosyncrasies. So it is probably best to go straight to what I think, and ignore the complexities of its relations to standard views. One of the idiosyncrasies, you will notice, is that this response wouldn't help the representationalist. Maybe that is another argument against representationalism. While we will focus on the case of Mary and not worry much about bats here, it will also be useful, as we proceed, to bear in mind an alternative case, where someone on the basis of vast knowledge of bat neurophysiology attempts to understand what it is like to be a bat. But here's what I have to say about Mary. There are no unicorns. Mary may "see" a unicorn, but then she is hallucinating. Yet she can focus her thought and perception on that unicorn she "sees," and tell us all about it. She can also focus on her color quale of unique red when she comes out of her black-and-white room and looks at a suitably red apple. According to the form of physicalism developed here, that unique red she sees is also a kind of illusion, an illusion that she can likewise focus on and describe, but there is in reality nothing that is unique red, although there is illusory experience of unique red. A sophisticated bat could presumably focus on its analogous but to us unknown illusions.

Aristotle thought such unique red was right out there on the objects, but he was wrong about that. There is nothing out in the world of ordinary objects that corresponds exactly to that unique red quale. While the conservatives will say under these conditions that red exists but our experience misrepresents it, the radicals will say with Galileo that there really is no red in the world, just surface spectral reflectance. But it really doesn't matter for our current purposes whether the color red exists in some sense.

---

[11] Hill (1997); Loar (1990); Papineau (2002).

[12] Conee (1994).

What matters is the what-it's-like, the way in which unique red is presented to unbound Mary. There is nothing in reality which is that specific sort of unique red, red just as she experiences it, phenomenal unique red, neither out in the world nor as part of dualist spiritual stuff. Still, she "sees" it. What is real is Mary's state of experiencing such a phenomenal red, a red on which she can focus, but only in the way some hallucinator may focus on a unicorn.

What then did Mary learn when she came out of the room? She already knew all about the neurophysiology of vision. She already knew all about what things in the environment had various sorts of surface spectral reflectances, and how that would interact with the neurophysiology of vision. She already knew about all her own relevantly possible neurophysiological states. But she had yet to come to quasi-perceptual acquaintance with illusory unique red, and to know which objects with what surface spectral reflectances in the world would seem that way, and which forms of neurophysiology would correspond to that particular experience.[13]

How does this situation allow for this new knowledge? In particular, what is it that Mary comes to believe when she comes to believe that to be in a certain neurophysiological state is really the same as to experience that unique red?

Notice that this is not the kind of thing most animals could believe. Such a belief is a complex cognitive achievement, and plausibly requires concepts of a rather sophisticated kind.[14] There is the quasi-physical and neurophysiological concept she had of the state in question all along. And then there is another concept she has, which is complex in a different way, after she comes out. She had experienced many things even in her room, and she had the concept of what it was to experience something in general. But by coming out of the room, she gained a new sort of experience, because its sensory content was new, and she could now think of experiencing like that, and so had available a new rather specific experience-based concept. Does the concept incorporate the new thing she experiences? Not

---

[13] Marianna, in Nida-Rümelin (1998), initially has acquaintance with phenomenal red but doesn't connect it up with objects she knows to be red. Hence the need for the second and third clauses of this sentence.

[14] This point undercuts current worries about the phenomenal concepts strategy felt by some of its originators, including Hill and Tye, although it is relevant that I do not accept Tye's historical conception of concepts or his view that the Burge phenomenon is applicable. See Hill (2009); Sainsbury and Tye (2012); and Tye (2009).

literally, since that doesn't really exist. But it involves it in the way the concept of believing in unicorns involves unicorns.[15]

What Mary comes to believe and know when she comes out of the room is that two different concepts, a neurophysiological concept of her state of experiencing, and the experienced-based concept she gains after leaving the room, are coreferential. What can't she know about a bat? She knows all the facts in a sense. But she is incapable of forming relevant experience-based concepts, and so of believing similar claims about concept coreference.

You might prefer a version of this sort of response to Mary that focused not on concepts of experiencing red but on concepts of red, that explains how Mary learns in particular what red looks like. In fact, on my view, Mary lacks *two* experienced-based concepts before she leaves the room. One is of phenomenal red, and another is of what it is like to experience phenomenal red. The second is plausibly dependent on the first, but the first is nonveridical, since nothing in reality is like that. In reality, possession of that first concept in any very full sense requires illusory experience of a certain sort that not all creatures can have. The second is not a fully transparent concept, which allows us to see by introspection deeply into conditions that are required for its instantiation in ways that dualist arguments demand. But it is somewhat translucent, in that it allows us to understand something of the modal structure involved in the apparent property of which it is an experience, and hence get clues about the proper modal structural account of its own instantiation, clues that I have exploited in this book.[16] There are dualists who claim that we can see by our concept of the experience of phenomenal red that such experience cannot be constituted by neurophysiology in the way I have proposed, that such an experienced-based concept and such a neurophysiological concept cannot be coreferential. But this requires a kind of fully transparent concept of our experience of phenomenal red, which is not plausibly available to us simply by introspection of such an experience.

It is important to notice that there are two different things, the experience of unique red and the unique red that is experienced. And on the view developed here, there is no unique red in reality, but there is experience of unique red. This difference is sometimes overlooked, partly because

---

[15] This undercuts the dualist arguments in Chalmers (2010: 251–275 and 305–336) when conjoined with the remarks about zombies.

[16] This distinction is due to Goff (2017: 102). A concept is transparent when it reveals the essence of something, and translucent when it reveals a significant aspect of its essence but not all of it.

"experience," like "thought," can be ambiguous regarding whether it refers to a mental state with a content or to just the content of that state.[17] But the difference is crucial, I claim, to a proper understanding of Mary.

All this is not to say that the familiar case-based arguments against physicalism about phenomenal consciousness do not express a legitimate demand. I think the physicalist should provide some intelligible explanation of how the physical can constitute conscious phenomenal experience. But I claim to have provided here the rudiments of just such an account.

---

[17] There are views not unlike mine that fail to note this ambiguity. One example is Balog (2012). And because I think there is in reality no unique red, I count as an illusionist in the sense articulated by Frankish (2016), while I think phenomenal consciousness itself exists and is in that sense no illusion.

Bibliography

Abramov, I. and J. Gordon. 1994. Color appearance: on seeing red – or yellow, or green, or blue. *Annual Review of Psychology* 45, 451–485.

Akins, K. 1996. Lost the plot? *Mind & Language* 11, 1–43.

Albert, B., R. Heald, and A. Johnson, et al. 2015. *The Molecular Biology of the Cell*, sixth edition. New York: Garland Press.

Allen, K. 2009. Inter-species variation in colour perception. *Philosophical Studies* 142, 197–220.

Aranyosi, I. 2013. *The Peripheral Mind*. New York: Oxford University Press.

Armstrong, D. M. 1968. *A Materialist Theory of Mind*. London: Routledge & Kegan Paul.

Aronov, D., R. Nevers, and D. W. Tank. 2017. Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature* 543, 719–722.

Arstila, V. and D. Lloyd. 2014. *Subjective Time*. Cambridge, MA: MIT Press.

Austin, D. F. 1985. *What's the Meaning of "This"?* Ithaca, NY: Cornell University Press.

Azzouni, T. J. 2013. *Semantic Perception*. Oxford: Oxford University Press.

Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

1997. *In the Theater of Consciousness: The Workplace of the Mind*. Oxford: Oxford University Press.

Baker, D. J. 2012. "The experience of left and right" meets the physics of left and right. *NOUS* 46, 483–498.

Balog, K. 2012. Acquaintance and the mind-body problem. In S. Gozzano and C. Hill (eds.), *New Perspectives on Type Identity*. Cambridge: Cambridge University Press, 16–43.

Barinaga, M. 1999. Learning visualized, on the double. *Science* 286, 1661.

Bartoshuk, L. 1978. History of taste research. In E. Carterette and M. Friedman (eds.), *Handbook of Perception, VIa*. New York: Academic Press, 3–18.

Barwich, A. S. 2020. *Smellosophy*. Cambridge, MA: Harvard University Press.

Batty, C. 2014. Olfactory objects. In D. Stokes, M. Matthen, and S. Biggs (eds.), *Perception and Its Modalities*. New York: Oxford University Press, 222–246.

Bayne, T. 2008. The unity of consciousness and the split-brain syndrome. *Journal of Philosophy* 105, 277–300.

https://doi.org/10.1017/9781009514774.009

2009. Perception and the reach of phenomenal content. *Philosophical Quarterly* 59, 385–404.

2010. *The Unity of Consciousness*. New York: Oxford University Press.

2020. Conscious thought. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 142–163.

Bayne, T., J. Hohwy, and A. M. Owen. 2016. Are there levels of consciousness? *Trends in Cognitive Sciences* 20, 405–413.

Beament, J. 2003. *How We Hear Music*. Martlesham: Boydell.

Bear, M. F., B. W. Connors, and M. A. Paradiso. 2016. *Neuroscience*, fourth edition. Philadelphia, PA: Wolters-Kluwer.

Bennett, J. 1966. *Kant's Analytic*. Cambridge: Cambridge University Press.

Berger, M. and F. Gerstenbrand. 1981. Phantom illusions in spinal cord lesions. In J. Siegried and M. Zimmermann (eds.), *Phantom and Stump Pain*. Berlin: Springer, 66–73.

Berkeley, G. 1710. *An Essay Towards a New Theory of Vision*. Dublin: Rhames.

Bermúdez, J. L. 2003. *Thinking Without Words*. New York: Oxford University Press.

Biederman, I. 2013. Psychophysical and neural correlates of the phenomenology of shape. In L. Albertazzi (ed.), *Handbook of Experimental Phenomenology: Visual Perception of Shape, Space, and Appearance*. Malden, MA: Wiley-Blackwell, 417–436.

Billock, V. A. and B. H. Tsou. 2010. Seeing forbidden colors. *Scientific American* 302, 72–77.

Blackwell, H. R. and H. Schlosberg. 1943. Octave generalization, pitch discrimination, and loudness thresholds in the white rat. *Journal of Experimental Psychology* 33, 407–419.

Block, N. 1978. Troubles with functionalism. *Minnesota Studies in the Philosophy of Science* 9, 261–325.

1990. Inverted earth. *Philosophical Perspectives* 4, 53–79.

1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18, 227–247.

2007. Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences* 30, 481–548.

2008. Consciousness and cognitive access. *Proceedings of the Aristotelian Society* 108, 289–317.

2010. Attention and mental paint. *Philosophical Issues* 20, 23–63.

2013. The grain of vision and the grain of attention. *Thought* 1, 170–184.

2014. Rich conscious perception outside focal attention. *Trends in Cognitive Science* 18, 445–447.

Block, N. and R. Stalnaker. 1999. Conceptual analysis, dualism, and the explanatory gap. *Philosophical Review* 108, 1–46.

Borg, G. H., H. Diamant, L. Ström, and Y. Zotterman. 1967. The relation between neural and perceptual intensity. *The Journal of Physiology* 192, 13–20.

Boring, E. G. 1942. *Sensation and Perception in the History of Experimental Psychology*. New York: Appleton.

Bourget, D. 2010. Consciousness is underived intentionality. *NOUS* 44, 32–58.

Bower, J. M. and L. M. Parsons. 2003. Rethinking the "lesser brain". *Scientific American* 296, 51–57.

Bowmaker, J. K. and H. J. A. Dartnall. 1980. Visual pigments of rods and cones in a human retina. *Journal of Physiology* 298, 501–511.

Bradbury, J. 2004. Taste perception: cracking the code. *PLoS Biology* 2, 295–297.

Bregman, A. S. 1990. *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.

Bricker, P. 2001. Island universes and the analysis of modality. In G. Preyer and F. Siebelt (eds.), *Reality and Humean Supervenience*. Lanham: Rowman and Littlefield, 27–55.

2006. Absolute actuality and the plurality of worlds. *Philosophical Perspectives* 20, 41–76.

Broackes, J. 2010. What do the color-blind see? In J. Cohen and M. Matthen (eds.), *Color Ontology and Color Science*. Cambridge, MA: MIT Press, 291–404.

2011. Where do the unique hues come from? *Review of Philosophy and Psychology* 2, 601–628.

Bull, M. and L. Back (eds.). 2003. *The Auditory Culture Reader*. Oxford: Berg.

Burge, T. 2010. *Origins of Objectivity*. Oxford: Clarendon Press.

Byrne, A. and D. Hilbert. 2003. Color realism and color science, with commentaries. *Behavioral and Brain Sciences* 26, 3–64.

2010. How do things look to the color-blind? In J. Cohen and M. Matthen (eds.), *Color Ontology and Color Science*. Cambridge, MA: MIT Press, 259–290.

Campbell, J. 1997. A simple view of color. In A. Byrne and D. Hilbert (eds.), *Readings on Color*, Vol. 1. Cambridge, MA: MIT Press, 177–190.

1999. *Past, Space, and Self.* Cambridge, MA: MIT Press.

2002. *Reference and Consciousness*. Oxford: Clarendon Press.

Campbell, K. 1970. *Body and Mind*. London: Doubleday.

Cao, R. 2012. A teleosemantic approach to information in the brain. *Biology and Philosophy* 27, 49–71.

2014. Signaling in the brain: in search of functional units. *Philosophy of Science* 81, 891–901.

Carroll, S. B. 2005. *Endless Forms Most Beautiful*. New York: Norton.

Carruthers, P. 2000. *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press.

Cervero, F. 2012. *Understanding Pain*. Cambridge, MA: MIT Press.

Chalmers, D. 1995. Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2, 200–219.

1996. *The Conscious Mind*. New York: Oxford University Press.

2002. Does conceivability entail possibility? In T. Gendler and J. Hawthorne (eds.), *Conceivability and Possibility*. Oxford: Clarendon Press, 145–200.

2004. Epistemic two-dimensional semantics. *Philosophical Studies* 18, 153–226.

2010. *The Character of Consciousness*. New York: Oxford University Press.

2012. *Constructing the World*. Oxford: Oxford University Press.

Chalmers, D. and F. Jackson. 2001. Conceptual analysis and reductive explanation. *Philosophical Review* 110, 315–361.

Chandrashekar, J., M. A. Hoon, N. J. P. Ryba, and C. S. Zuker. 2006. The receptors and cells for mammalian taste. *Nature* 444, 288–294.

Chang, L. and D. Y. Tsao. 2017. The code for facial identity in the primate brain. *Cell* 169, 1013–1028.

Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.

Chirimuuta, M. 2014. Minimal models and canonical neural computation. *Synthese* 191, 127–153.

Chuard, P. 2020. Temporal consciousness. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 188–207.

Churchland, P. 1986. Some reductive strategies in cognitive neurobiology. *Mind* 95, 279–309.

1995. *The Engine of Reason, The Seat of the Soul*. Cambridge, MA: MIT Press.

Clark, Andy. 1996. *Being There*. Cambridge, MA: MIT Press.

Clark, Austen. 1993. *Sensory Qualities*. Oxford: Oxford University Press.

2000. *A Theory of Sentience*. Oxford: Oxford University Press.

2004. Feature-placing and proto-objects, and reply to commentators. *Philosophical Psychology* 17, 451–477 and 563–589.

Cohen, J. 2009. *The Red and the Real*. New York: Oxford University Press.

Cole-Adams, K. 2018. *Anaesthesia*. Berkeley, CA: Counterpoint Press.

Conee, E. 1994. Phenomenal knowledge. *Australian Journal of Philosophy* 72, 136–150.

Coombs, J. 1993. The possibility of created entities in seventeenth-century Scotism. *The Philosophical Quarterly* 43, 447–459.

Cosmelli, D., J.-P. Lachaux, and E. Thompson. 2007. Neurodynamical approaches to consciousness. In P. D. Zelazo, M. Moscovitch, and E. Thompson (eds.), *The Cambridge Handbook of Consciousness*. Cambridge: Cambridge University Press, 731–772.

Cowart, B. J. and N. E. Rawson. 2001. Olfaction. In E. Goldstein (ed.), *The Blackwell Handbook of Perception*. Oxford: Blackwell, 568.

Crane, H. D. and T. P. Piantanida. 1983. On seeing reddish green and yellowish blue. *Science* 221, 1078–1080.

Craver, C. 2007. *Explaining the Brain*. Oxford: Oxford University Press.

Crick, F. and C. Koch. 2003. A framework for consciousness. *Nature Neuroscience* 6, 119–126.

Cynx, J. 1993. Auditory frequency generalization and a failure to find octave generalization in a songbird, the European starling. *Journal of Comparative Psychology* 107, 140–146.

Dainton, B. 2000. *Stream of Consciousness*. London: Routledge.

2014. The phenomenal continuum. In V. Arstila and D. Lloyd (eds.), *Subjective Time*. Cambridge, MA: MIT Press, 101–137.

2023. Temporal consciousness. *Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu.

Dennett, D. 1991. *Consciousness Explained*. Boston, MA: Little, Brown & Co.

Denyer, N. 1994. Why do mirrors reverse left/right and not up/down? *Philosophy* 69, 205–210.

Deonna, J. and F. Teroni. 2020. Emotional experience. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 102–123.

Derrida, J. 1973. *Speech and Phenomena and Other Essays on Husserl's Theory of Signs*, D. B. Allison (trans.). Evanston, IL: Northwestern University Press.

Derrington, A. M., J. Krauskopf, and P. Lennie. 1984. Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology* 357, 241–265.

DeValois, R. L. and K. K. DeValois. 1988. *Spatial Vision*. New York: Oxford University Press.

Di Lorenzo, P. M., J.-Y. Chen, and J. D. Victor. 2009. Quality time: representations of a multidimensional sensory domain through temporal coding. *Journal of Neuroscience* 29, 9227–9238.

Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.

Edelman, S. 1998. Representation is representation of similarities. *Behavioral and Brain Sciences* 21, 449–498.

Egan, A. 2006. Appearance properties? *NOUS* 40, 495–521.

Ehrenreich, B. 2018. *Natural Causes*. New York: Twelve.

Eliasmith, C. 1996. The third contender. *Philosophical Psychology* 9, 441–463.

Ellenberger, H. 1970. *The Discovery of the Unconscious*. New York: Basic Books.

Elpidorou, A. 2015. Review of *The Peripheral Mind*. Notre Dame Philosophical Reviews, https://ndpr.nd.edu, 2015.04.26.

Evans, G. 1982. *The Varieties of Reference*, ed. by J. McDowell. New York: Oxford University Press.

Falkowski, P. G. 2015. *Life's Engines*. Princeton, NJ: Princeton University Press.

Farkas, K. 2008. Phenomenal intentionality without compromise. *The Monist* 91, 273–293.

2013. Constructing a world for the senses. In U. Kriegel (ed.), *Phenomenal Intentionality*. Oxford: Oxford University Press, 99–115.

Fields, R. D. 2004. The other half of the brain. *Scientific American* 297, 54–61.

2008. White matter matters. *Scientific American* 301, 54–61.

Fodor, J. 1990. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.

Förster, E. 2002. *Kant's Final Synthesis*. Cambridge, MA: Harvard University Press.

Frankish, K. 2016. Illusionism as a theory of consciousness. *Journal of Consciousness Studies* 23, 11–39.

Friederici, A. 2017. *Language in Our Brain*. Cambridge, MA: MIT Press.

Freud, S. 1962. *The Ego and the Id*, J. Riviere and J. Strachey (trans.). New York: Norton.

1963. *General Psychological Theory*. New York: Macmillan.

1965. *The Interpretation of Dreams*, J. Strachey (trans.). New York: Avon Books.

1989. *Introductory Lectures on Psycho-Analysis and New Introductory Lectures on Psycho-Analysis*, J. Strachey (trans.). New York: Norton.

Gallagher, S. and D. Zahavi. 2014. Primal impression and enactive perception. In V. Arstila and D. Lloyd (eds.), *Subjective Time*. Cambridge, MA: MIT Press, 83–99.

Gazzaniga, M. S. 1970. *The Bisected Brain*. New York: Appleton-Century-Crofts.

Gazzaniga, M. S., R. B. Ivry, and G. R. Mangun. 2002. *Cognitive Neuroscience*, second edition. New York: Norton.

Geldard, F. A and C. E. Sherrick. 1972. The cutaneous "rabbit": a perceptual illusion. *Science* 178, 178–179.

Gennaro, R. 1996. *Consciousness and Self-Consciousness*. Philadelphia, PA: John Benjamins.

Gertler, B. 2001. Introspecting phenomenal states. *Philosophy and Phenomenological Research* 63, 305–327.

Gibbs, W. W. 2003. The unseen genome. *Scientific American* 297, 47–53.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin.

Glennan, S. 1996. Mechanism and the nature of causation. *Erkenntnis* 44, 49–71.

2017. *The New Mechanical Philosophy*. Oxford: Oxford University Press.

Godfrey-Smith, P. 2007. Information in biology. In D. Hull and M. Ruse (eds.), *The Cambridge Companion to The Philosophy of Biology*. Cambridge: Cambridge University Press, 103–119.

2016a. Animal evolution and the origins of experience. In D. L. Smith (ed.), *How Biology Shapes Philosophy*. Cambridge: Cambridge University Press, 51–71.

2016b. Individuality, subjectivity, and minimal cognition. *Biology and Philosophy* 31, 775–796.

2016c. Mind, matter, and metabolism. *Journal of Philosophy* 113, 481–506.

2017. *Other Minds*. New York: Farrar, Straus, and Giroux.

2020. Gradualism and the evolution of experience. *Philosophical Topics* 48, 201–220.

Goff, P. 2017. *Consciousness and Fundamental Reality*. Oxford: Oxford University Press.

Goff, P. and S. Coleman. 2020. Russellian monism. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 301–327.

Goodsell, D. S. 2010. *The Machinery of Life*, second edition, second corrected printing. New York: Copernicus Books.

Gordon, I. and S. Hayward. 1973. Second-order isomorphism of internal representations of familiar faces. *Perception & Psychophysics* 14, 334–336.

Green, E. J. 2015. Layered view of shape perception. *British Journal of Philosophy of Science* 68, 355–387.

2016. Representationalism and perceptual organization. *Philosophical Topics* 44, 121–148.

2017. Attentive visual reference. *Mind & Language* 32, 3–28.

2018. What do object files pick out? *Philosophy of Science* 85, 177–200.

2019. On the perception of structure. *NOUS* 53, 564–592.

Green, S. and R. Batterman. 2017. Biology meets physics. *Studies in History and Philosophy of the Biological and Biomedical Sciences* 61, 20–34.

Greenfield, S. A. and T. F. T. Collins. 2005. A neuroscientific approach to consciousness. *Progress in Brain Research* 150, 11–23.

Griffiths, P. 2001. Genetic information: a metaphor in search of a theory. *Philosophy of Science* 68, 394–412.

Griffiths, P. and K. Stolz. 2007. Gene. In D. Hull and M. Ruse (eds.), *The Cambridge Companion to The Philosophy of Biology*. Cambridge: Cambridge University Press, 85–102.

Groh, J. M. 2014. *Making Space*. Cambridge, MA: Harvard University Press.

Grush, R. 2007a. Berkeley and the spatiality of vision. *Journal of the History of Philosophy* 43, 413–442.

2007b. Skill theory v2.0: dispositions, emulation, and spatial perception. *Synthese* 159, 59–92.

Hardin, C. L. 1993. *Color for Philosophers*, expanded edition. Indianapolis, IN: Hackett.

1997. Reinverting the spectrum. In A. Byrne and D. Hilbert (eds.), *Readings on Color*, Vol. 1. Cambridge, MA: MIT Press, 289–301.

2014. More color science for philosophers. In D. Stokes, M. Matthen, and S. Biggs (eds.), *Perception and Its Modalities*. Oxford: Oxford University Press, 379–389.

Harman, G. 1973. *Thought*. Princeton, NJ: Princeton University Press.

1990. The intrinsic quality of experience. In J. Tomberlin (ed.), *Philosophical Perspectives*, Vol. IV. Atascadero, CA: Ridgeview Publishing, 31–52.

Harold, F. M. 2001. *The Way of the Cell*. New York: Oxford University Press.

Harris, C. 1965. Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological Review* 72, 419–444.

1980. Insight or out of sight?: two examples of perceptual plasticity in the human adult. In C. Harris (ed.), *Visual Coding and Adaptability*. Hillsdale, NJ: Lawrence Erlbaum, 95–149.

Harrison, B. 1973. *Form and Content*. Oxford: Blackwell.

1986. Identity, predication, and color. *American Philosophical Quarterly* 23, 105–114.

Hatfield, G. and W. Epstein. 1979. The sensory core and the medieval foundations of early modern perceptual theory. *ISIS* 70, 363–384.

Heidegger, M. 1962. *Being and Time*, J. Macquarrie and E. Robinson (trans.). New York: Harper & Rowe.

Helmholtz, H. 1954. *On the Sensations of Tone*, A. J. Ellis (trans.), second edition. Mineola, NY: Dover.

Henry, M. 1975. *Philosophy and Phenomenology of the Body*, G. Etzkorn (trans.). The Hague: Martinus Nijhoff.

Hilbert, D. 1987. *Color and Color Perception*. Stanford, CA: CSLI.

Hill, C. 1997. Imaginability, conceivability, possibility, and the mind-body problem. *Philosophical Studies* 87, 61–85.

2009. *Consciousness*. Cambridge: Cambridge University Press.

2015. Review of *The Peripheral Mind*. *Mind* 124, 312–317.

Hobson, J. A. 1988. *The Dreaming Brain*. New York: Basic Books.

Holcombe, A. O. 2014. Are there cracks in the façade of continuous visual experience? In V. Arstila and D. Lloyd (eds.), *Subjective Time*. Cambridge, MA: MIT Press, 179–198.

Horgan, T. and J. Tienson. 2002. The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers (ed.), *Philosophy of Mind*. New York: Oxford University Press, 520–533.

Howard, I. P., B. Craske, and W. B. Templeton. 1965. Visuomotor adaptation to discordant exafferent stimulation. *Journal of Experimental Psychology* 70, 189–191.

Howard, J. D., J. Plailly, M. Grueschow, J.-D. Haynes, and J. A. Gottfried. 2009. Odor quality coding and categorization in human posterior piriform cortex. *Nature Neuroscience* 12, 932–939.

Howe, C. Q. and D. Purves. 2005. *Perceiving Geometry*. New York: Springer.

Hubel, D. 1995. *Eye, Brain, and Vision*. New York: Scientific American Books.

Hurley, S. 1998. *Consciousness in Action*. Cambridge, MA: Harvard University Press.

2007. Neural dominance, neural deference, and sensorimotor dynamics. In M. Velmans and S. Schneider (eds.), *The Blackwell Companion to Consciousness*. Oxford: Blackwell, 640–656.

Husserl, E. 1991. *On the Phenomenology of the Consciousness of Internal Time*, J. Brough (trans.). Dordrecht: Kluwer.

Jackson, F. 1975. On the adverbial analysis of visual experience. *Metaphilosophy* 6, 127–135.

1977. *Perception*. Cambridge: Cambridge University Press.

1982. Epiphenomenal qualia. *The Philosophical Quarterly* 32, 127–136.

Jeannerod, M. 1997. *The Cognitive Neuroscience of Action*. Cambridge, MA: Blackwell Science.

Jekely, G., F. Keijzer, and P. Godfrey-Smith. 2015. An option space for early neural evolution. *Philosophical Transactions of the Royal Society of London, Series B, Biological Science* 370, 1684.

Joesch, M. and M. Meister. 2016. A neuronal circuitry for colour vision based on rod-cone opponency. *Nature* 532, 236–239.

Kaiser, P. K. and R. M. Boynton. 1996. *Human Color Vision*, second edition. Washington, DC: Optical Society of America.

Kandel, E. R., J. Schwartz, T. Jessell, E. Siegelbaum, and A. J. Hudspeth. 2013. *Principles of Neural Science*, fifth edition. New York: McGraw-Hill.

Kauffman, S. 2000. *Investigations*. New York: Oxford University Press.

Keller, E. F. 2000. *The Century of the Gene*. Cambridge, MA: Harvard University Press.

2002. *Making Sense of Life*. Cambridge, MA: Harvard University Press.

Kennedy, J. M. 1993. *Drawing and the Blind.* New Haven, CT: Yale University Press.

1997. How the blind draw. *Scientific American* 289, 76–81.

Kepple, D. and A. Koulakov. 2017. Constructing an olfactory perceptual space and predicting percepts from molecular structure. arXiv 1708.05774. http://doi.org/10.48550/arKiv.1708.05574.

Kim, J. 2011. *Philosophy of Mind*, third edition. Boulder, CO: Westview Press.

Kind, A. 2018. Imaginative presence. In F. Dorsch and F. Macpherson (eds.), *Phenomenal Presence*. Oxford: Oxford University Press, 165–180.

2020. Imaginative experience. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 124–141.

Kirk, R. 1974. Zombies vs. materialists. *Proceedings of the Aristotelian Society* 48, 135–152.

Knill, D. and W. Richards. 1996. *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.

Koch, C. 2004. *The Quest for Consciousness*. Englewood, CO: Roberts.

Kohler, I. 1964. Formation and transformation of the perceptual world. *Psychological Issues* 3, 1–173.

Kohlers, P. A. and M. von Grünau. 1976. Shape and color in apparent motion. *Vision Research* 16, 329–335.

Koulakov, A., B. Kolterman, A. Enikolopov, and D. Rinberg. 2011. In search of the structure of human olfactory space. *Frontiers in Systems Neuroscience* 5, 1–8.

Kriegel, U. 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford: Oxford University Press.

Kripke, S. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Kroon, F. 2013. Phenomenal intentionality and the role of intentional objects. In U. Kriegel (ed.), *Phenomenal Intentionality*. Oxford: Oxford University Press, 137–155.

Kuehni, R. G. 2010. Color spaces and color order systems: a primer. In J. Cohen and M. Matthen (eds.), *Color Ontology and Color Science*. Cambridge, MA: MIT Press, 3–36.

Langers, D., P. van Dijk, E. Schoenmaker, and W. Backes. 2007. fMRI activation in relation to sound intensity and loudness. *NeuroImage* 35, 709–718.

Laurent, G. 1999. A systems perspective on early olfactory coding. *Science* 286, 723–728.

Lee, A. Y. 2023. Degrees of consciousness. *NOUS* 57, 553–575.

Lee, G. 2006. The experience of left and right. In T. Gendler and J. Hawthorne (eds.), *Perceptual Experience*. Oxford: Oxford University Press, 292–315.

2007. Consciousness in a space-time world. *Philosophical Perspectives* 21, 341–374.

2014a. Experiences and their parts. In D. Bennett and C. Hill (eds.), *Sensory Integration and the Unity of Consciousness*. Cambridge, MA: MIT Press, 287–321.

2014b. Temporal experience and the temporal structure of experience. *Philosopher's Imprint* 14, 1–21.

2015. Review of *The Conscious Brain*. *Philosophical Review* 124, 163–167.

Leng, G. 2018. *The Heart of the Brain*. Cambridge, MA: MIT Press.

Levine, J. 1983. Materialism and qualia: the explanatory gap. *Pacific Philosophical Quarterly* 64, 354–361.

2001. *Purple Haze*. New York: Oxford University Press.

2010. Review of Kriegel, *Subjective Consciousness*. Notre Dame Philosophical Reviews, https://ndpr.nd.edu, 2010.03.14.

Lewens, T. 2007. Adaptation. In D. Hull and M. Ruse (eds.), *The Cambridge Companion to the Philosophy of Biology*. Cambridge: Cambridge University Press, 1–21.

Lewis, D. 1975. Languages and language. In K. Gunderson (ed.), *Mind and Knowledge*. Minneapolis, MN: University of Minnesota Press, 3–35.

1986. *On the Plurality of Worlds*. Oxford: Blackwell.

Liberman, A. M. 1996. *Speech: A Special Code*. Cambridge, MA: MIT Press.

Loar, B. 1990. Phenomenal states. *Philosophical Perspectives* 4, 81–108.

Lockwood, M. 1989. *Mind, Brain, and the Quantum*. Oxford: Blackwell.

Loebrich, S. and E. Nedivi. 2009. The function of activity-regulated genes in the nervous system. *Physiological Review* 89, 1079–1103.

Loewenstein, W. 1999. *The Touchstone of Life*. New York: Oxford University Press.

Lycan, W. 1987. *Consciousness*. Cambridge, MA: MIT Press.

1996. *Consciousness and Experience*. Cambridge, MA: MIT Press.

2004. The superiority of HOP to HOT. In R. Gennaro (ed.), *Higher-Order Theories of Consciousness*. Amsterdam: John Benjamins, 115–136.

2020. Hearing As. In B. Brogaard and D. E. Gatzia (eds.), *The Epistemology of Non-Visual Perception*. New York: Oxford University Press, 118–145.

Lyons, W. 1986. *The Disappearance of Introspection*. Cambridge, MA: MIT Press.

Machamer, P., L. Darden, and C. Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67, 1–25.

Mack, A. and I. Rock. 1998. *Inattentional Blindness*. Cambridge, MA: MIT Press.

Maclachlan, D. L. C. 1989. *Philosophy of Perception*. Englewood Cliffs, NJ: Prentice Hall.

MacLeod, D. 2010. Into the neural maze. In J. Cohen and M. Matthen (eds.), *Color Ontology and Color Science*. Cambridge, MA: MIT Press, 151–178.

Marks, C. 1981. *Commissurotomy, Consciousness, and Unity of Mind*. Cambridge, MA: MIT Press.

Marr, D. 1982. *Vision*. San Francisco, CA: W.H. Freeman.

Martin, M. G. F. 1997. Self-observation. *European Journal of Philosophy* 5, 119–140.

Masrour, F. 2013. Phenomenal objectivity and phenomenal intentionality. In U. Kriegel (ed.), *Phenomenal Intentionality*. Oxford University Press, 116–136.

2015. The geometry of visual space and the nature of visual experience. *Philosophical Studies* 172, 1813–1832.

2016. Space perception, visual dissonance, and the fate of standard representationalism. *NOUS* 50, 565–593.

2020. The phenomenal unity of consciousness. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 208–229.

Matthen, M. 2005. *Seeing, Doing, and Knowing*. Oxford: Clarendon Press.

Mauk, M. and D. Buonomano. 2004. The neural basis of temporal processing. *Annual Review of Neuroscience* 27, 307–330.

McClelland, T. 2020. Self-representationalist theories of consciousness. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 458–481.

McDowell, J. 1996. *Mind and World*. Cambridge, MA: Harvard University Press.

McGovern, K. and B. J. Baars. 2007. Cognitive theories of consciousness. In P. Zelazo, M. Moscovitch, and E. Thompson (eds.), *The Cambridge Handbook of Consciousness*. Cambridge: Cambridge University Press, 177–205.

McGurk, H. and J. MacDonald. 1976. Hearing lips and seeing voices. *Nature* 264, 746–748.

Meinong, A. 1960. The theory of objects, I. Levi, D. B. Terrell, and R. Chisholm (trans.). In R. Chisholm (ed.), *Realism and the Background of Phenomenology*. Glencoe, IL: Free Press, 76–117.

Mel, B. W. 1986. A connectionist learning model for 3-dimensional mental rotation, zoom, and pan. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum, 562–571.

Mendelovici, A. 2018. *The Phenomenal Basis of Intentionality*. New York: Oxford University Press.

2020. How reliably misrepresenting olfactory experiences justify true beliefs. In B. Brogaard and D. Gatzia (eds.), *The Epistemology of Non-Visual Perception*. New York: Oxford University Press, 99–117.

Mendelovici, A. and D. Bourget. 2020. Consciousness and intentionality. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 560–885.

Mendola, J. 1997. *Human Thought*. Dordrecht: Kluwer.

2006. *Goodness and Justice*. Cambridge: Cambridge University Press.

2008. *Anti-Externalism*. Oxford: Oxford University Press.

2014. *Human Interests*. Oxford: Oxford University Press.

2021. *Experience and Possibility*. Oxford: Oxford University Press.

Merino-Rajme, C. 2014. A quantum of felt duration. *Analytic Philosophy* 55, 239–275.

Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Milner, D. and M. A. Goodale. 2008. Two visual systems reviewed. *Neuropsychologia* 46, 774–785.

Mitrani, L., S. Shekerdjiiski, and N. Yakimoff. 1986. Mechanisms and asymmetries in visual perception of simultaneity and temporal order. *Biological Cybernetics* 54, 159–165.

Mojet, J., E. P. Köster, and J. F. Prinz. 2005. Do tastants have a smell? *Chemical Senses* 30, 9–21.

Mölder, B. 2014. Constructing time: Dennett and Grush on temporal representation. In V. Arstila and D. Lloyd (eds.), *Subjective Time*. Cambridge, MA: MIT Press, 217–238.

Mole, C. 2020. Consciousness and attention. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 499–537.

Mollon, J. D. 2009. A neural basis for unique hues? *Current Biology* 19, R441–R442.

2012. The origins of modern color science. In S. Shevell (ed.), *The Science of Color*, second edition. Amsterdam: Elsevier, 1–39.

Montemayor, C. and H. H. Haladjian. 2015. *Consciousness, Attention, and Conscious Attention*. Cambridge, MA: MIT Press.

Moore, C. I. and R. Cao. 2008. The hemo-neural hypothesis: on the role of blood flow in information processing. *Journal of Neurophysiology* 99, 2035–2047.

Moore, P. B. 2012. How should we think about the ribosome? *Annual Review of Biophysics* 41, 1–19.

Morales, J. and H. Lau. 2020. The neural correlates of consciousness. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 233–260.

Murphy, A. C., S. Muldoon, D. Baker, et al. 2018. Structure, function, and control of the human musculoskeletal network. *PLoS Biology* 16, e22002811.

Nagel, T. 1974. What is it like to be a bat? *Philosophical Review* 83, 435–450.

2002. The psychophysical nexus. In T. Nagel (ed.), *Concealment and Exposure*. New York: Oxford University Press, 194–236.

Nash, L. K. 1974. *Elements of Statistical Thermodynamics*, second edition. New York: Addison-Wesley.

Neander, K. 2017. *A Mark of the Mental*. Cambridge, MA: MIT Press.

Newell, B. R. and D. R. Shanks. 2014. Unconscious influences on decision making: A critical review, with critical discussion. *Behavior and Brain Sciences* 37, 1–61.

Nida-Rümelin, M. 1993. Pseudonormal vision. *Philosophical Studies* 82, 145–157.

1998. On belief about experiences. *Philosophy and Phenomenological Research* 58, 51–73.

Noble, D. 2008. Genes and causation. *Philosophical Transactions of the Royal Society A* 366, 3001–3015.

2012. A theory of biological relativity: no privileged level of causation. *Interface Focus* 2, 55–64.

Noë, A. 2005. Real presence. *Philosophical Topics* 33, 235–264.

2006. *Action in Perception*. Cambridge, MA: MIT Press.

2012. *Varieties of Presence*. Cambridge, MA: Harvard University Press.

Northoff, G. 2023. *Neurowaves*. Montreal: McGill-Queens University Press.

Northoff, G. and W.-D. Heiss. 2015. Why is the distinction between neural predispositions, prerequisites, and correlates of level of consciousness clinically relevant? *Stroke* 46, 1147–1151.

Northoff, G. and F. Zilio. 2022. Temporo-spatial theory of consciousness (TTC). *Behavioural Brain Research* 424, 113788.

Nusbaum, N., D. Blitz, and E. Marder. 2017. Functional consequences of neuropeptide and small-molecule co-transmission. *Nature Reviews Neuroscience* 18, 389–403.

O'Callaghan, C. 2007. *Sounds*. Oxford: Oxford University Press.

2015. Speech perception. In M. Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press, 475–494.

2019. *A Multisensory Theory of Perception*. Oxford: Oxford University Press, 19–25.

Oizumi, M., L. Albantakis, and G. Tononi. 2014. From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. *PLoS Computational Biology* 10, e1003588.

O'Regan, J. K. 2011. *Why Red Doesn't Sound Like a Bell*. Oxford: Oxford University Press.

Orlandi, N. 2013. Embedded seeing: vision in the natural world. *NOUS* 47, 727–747.

2014. *The Innocent Eye: Why Vision Is Not a Cognitive Process*. Oxford: Oxford University Press.

O'Shea, M. 2005. *The Brain*. Oxford: Oxford University Press.

Palmer, S. E. 1999. *Vision Science*. Cambridge, MA: MIT Press.

Papineau, D. 2002. *Thinking About Consciousness*. New York: Oxford University Press.

2021. *The Metaphysics of Sensory Experience*. Oxford: Oxford University Press.

Parkes, L. M., J.-B. Marsman, D. Oxley, J. Goulermas, and S. Wuerger. 2009. Multivoxel fMRI analysis of color tuning in human primary visual cortex. *Journal of Vision* 9, 1–13.

Parsons, T. 1980. *Nonexistent Objects*. New Haven, CT: Yale University Press.

Pautz, A. 2006. Sensory awareness is not a wide physical relation. *NOUS* 40, 205–240.

2010. A simple view of consciousness. In G. Bealer and R. Koons (eds.), *The Waning of Materialism*. Oxford: Oxford University Press, 25–67.

2013. The real trouble for phenomenal externalists: new empirical evidence for a brain-based theory of consciousness. In R. Brown (ed.), *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*. Dordrecht: Springer, 237–298.

2020. Representationalism about consciousness. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 405–437.

Peacocke, C. 1983. *Sense and Content*. Oxford: Oxford University Press.

Pearson, H. 2006. What is a gene? *Nature* 441, 399–401.

Peirce, C. S. 1982. Lowell lectures. In M. H. Fisch (ed.), *Writings of Charles S. Peirce*. Bloomington, IN: Indiana University Press, 358–504.

Pelczar, M. 2010. Presentism, eternalism, and phenomenal change. *Synthese* 176, 275–290.

Penrose, R. 2005. *The Road to Reality*. New York: Knopf.

Perkins, M. 1983. *Sensing the World*. Indianapolis, IN: Hackett.

Phillips, I. 2016. Consciousness and criterion. *Philosophy and Phenomenological Research* 93, 419–451.

Piantanida, P. 2010. Letter. *Scientific American* 303, 6–8.

Pitt, D. 2004. The phenomenology of cognition or what is it like to think that P?. *Philosophy and Phenomenological Research* 69, 1–36.

Prinz, J. 2012. *The Conscious Brain*. New York: Oxford University Press.

Pylyshyn, P. 2001. Visual indexes, preconceptual objects, and situated vision. *Cognition* 80, 127–158.

Quiroga, R. Q., L. Reddy, G. Kreiman, C. Koch, and I. Fried. 2005. Invariant visual representation by single neurons in the human brain. *Nature* 435, 1102–1107.

Rassoulzadegan, M., V. Grandjean, P. Gounon, et al. 2006. RNA-mediated non-mendelian inheritance of an epigenetic change in the mouse. *Nature* 441, 469–474.

Recanati, F. 2012. *Mental Files*. Oxford: Oxford University Press.

Relkin, E. M. and J. R. Doucet. 1997. Is loudness simply proportional to the auditory nerve spike count? *Journal of the Acoustical Society of America* 101, 2735–2740.

Remez, R. E. and J. D. Trout. 2009. Philosophical messages in the medium of spoken language. In M. Nudds and C. O'Callaghan (eds.), *Sounds and Perception*. Oxford: Oxford University Press, 234–264.

Rescorla, M. 2015. Bayesian perceptual psychology. In M. Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press, 694–716.

2016. Bayesian sensorimotor psychology. *Mind & Language* 31, 3–36.

Robinson, H. 2016. *From the Knowledge Argument to Mental Substance*. Cambridge: Cambridge University Press.

Rock, I. 1966. *The Nature of Perceptual Adaptation*. New York: Basic Books.

Roelofs, L. 2019. *Combining Minds*. Oxford: Oxford University Press.

Rojahn, S. Y. 2013. What it's like to see again with an artificial retina. www.technologyreview.com/s/514081/can-artificial-retinas-restore-natural-sight

Röhl, M. B., B. Kollmeier, and S. Uppenkamp. 2011. Spectral loudness summation takes place in the primary auditory cortex. *Human Brain Mapping* 32, 1483–1496.

Rolls, E. T. 2007. The affective neuroscience of consciousness. In P. D. Zelazo, M. Moscovitch, and E. Thompson (eds.), *The Cambridge Handbook of Consciousness*. Cambridge: Cambridge University Press, 831–859.

Rosenbaum, D. A. 1991. *Human Motor Control.* San Diego, CA: Academic Press.

Rosenthal, D. M. 2005. *Consciousness and Mind.* Oxford: Oxford University Press.

Russell, B. 1905. On denoting. *Mind* 14, 479–493.

Sainsbury, R. M. and M. Tye. 2012. *Seven Puzzles of Thought.* Oxford: Oxford University Press.

Sakurai, T., T. Misaka, Y. Ueno, et al. 2010. The human bitter taste receptor, hTAS2R16, discriminates slight differences in the configuration of disaccharides. *Biochemical and Biophysics Research Communication* 402, 595–601.

Sarkar, S. 1996. Decoding "coding" – information and DNA. *BioScience* 46, 857–864.

Sasaki, T., K. Beppu, K. Tanaka, et al. 2012. Application of an optogenetic byway for perturbing neural activity via glial photostimulation. *Proceedings of the National Academy of Sciences of the United States of America* 109, 20720–20725.

Schechter, E. 2014. Partial unity of consciousness: a preliminary defense. In D. Bennett and C. Hill (eds.), *Sensory Integration and the Unity of Consciousness.* Cambridge, MA: MIT Press, 347–373.

Schnupp, J., A. King, K. Walter, and J. Bizley. 2010. The representation of the pitch of vowel sounds in ferret auditory cortex. In E. A. Lopez, A. R. Palmer, and R. Meddis (eds.), *Advances in Auditory Physiology, Psychophysics and Models.* New York: Springer, 407–416.

Schnupp, J., I. Nelken, and A. King. 2011. *Auditory Neuroscience.* Cambridge, MA: MIT Press.

Schrödinger, E. 1967. *What Is Life?* Cambridge: Cambridge University Press.

Scott, J. D. and T. Pawson. 2000. Cell communication: the inside story. *Scientific American* 313, 72–79.

Seager, W. 2016. *Theories of Consciousness*, second edition. London: Routledge.

Searle, J. 1991. Response. In E. Lepore and R. Van Gulick (eds.), *John Searle and His Critics.* Cambridge, MA: Blackwell, 181–192.

Shannon, C. 1948. A mathematical theory of communication. *Bell Systems Technical Journal* 27, 279–423 and 623–656.

Shepard, R. N. and S. Chipman. 1970. Second-order isomorphism of internal representations: shapes of states. *Cognitive Psychology* 1, 1–17.

Siegel, S. 2006. Which properties are represented in perception? In T. Gendler and J. Hawthorne (eds.), *Perceptual Experience.* New York: Oxford University Press, 481–503.

2010. *The Contents of Visual Experience.* New York: Oxford University Press.

Simons, D. and C. Chabris. 1999. Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception* 28, 1059–1074.

Sinha, P., J. Wulff, and R. Held. 2014. Establishing cross-modal mappings: empirical and computational investigations. In D. Bennett and C. Hill (eds.), *Sensory Integration and the Unity of Consciousness.* Cambridge, MA: MIT Press, 171–191.

Sjovall, B. 1967. *Psychology of Tension.* Norstedts: Svenska Bokforlaget.

Skillings, D. J. 2015. Mechanistic explanation of biological processes. *Philosophy of Science* 82, 1139–1151.

Sklar, L. 1993. *Physics and Chance*. Cambridge: Cambridge University Press.

Smith, B. 2015. The chemical senses. In M. Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*. New York: Oxford University Press, 314–352.

2020. Tasting flavors. In B. Brogaard and D. Gatzia (eds.), *The Epistemology of Non-Visual Perception*. New York: Oxford University Press, 29–52.

Smith, D. V., R. L. Van Buskirk, J. B. Travers, and S. L. Bieber. 1983. Coding of taste stimuli by hamster brain stem neurons. *Journal of Neurophysiology* 50, 541–558.

Sober, E. 1984. *The Nature of Selection*. Cambridge, MA: MIT Press.

1994. *Conceptual Issues in Evolutionary Biology*, second edition. Cambridge, MA: MIT Press.

Solomon, S. G. and P. Lennie. 2007. The machinery of color vision. *Nature, Reviews Neuro* 8, 276–286.

Soloway, P. D. 2006. Paramutable possibilities. *Nature* 441, 413–414.

Spence, C., M. Auvray, and B. Smith. 2014. Confusing tastes with flavours. In D. Stokes, M. Matthen, and S. Biggs (eds.), *Perception and Its Modalities*. Oxford: Oxford University Press, 247–274.

Spinney, L. 2011. You smell flowers, I smell stale urine. *Scientific American* 324, 26.

Spivey, M. 2007. *The Continuity of Mind*. Oxford: Oxford University Press.

Squire, L. R. 1987. *Memory and Brain*. New York: Oxford University Press.

Stalnaker, R. 1976. Possible worlds. *NOUS* 10, 65–75.

Stockman, A., D. MacLeod, and N. Johnson. 1993. Spectral sensitivities of the human cones. *Journal of the Optical Society of America A, Optics, Image Science, and Vision* 10, 2491–2521.

Stoughton, C. M. and B. R. Conway. 2008. Neural basis for unique hues. *Current Biology* 18, R698–R699.

Stratton, G. M. 1896. Some preliminary experiments on vision without inversion of the retinal image. *Psychological Review* 3, 611–617.

1897. Vision without inversion of the retinal image. *Psychological Review* 4, 341–360 and 463–481.

Strawson, G. 2006. Realistic monism: why physicalism entails panpsychism. *Journal of Consciousness Studies* 13, 3–31.

Strawson, P. F. 1966. *The Bounds of Sense*. London: Methuen.

Sundström, P. 2020. Visual experience. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness*. Oxford: Oxford University Press, 39–65.

Tailby, A., S. G. Solomon, and P. Lennie. 2008. Functional asymmetries in visual pathways carrying S-cone signals in macaque. *Journal of Neuroscience* 28, 4078–4087.

Taylor, J. G. 1962. *The Behavioral Basis of Perception*. New Haven, CT: Yale University Press.

Teller, D. 2003. Color: a vision scientist's perspective. *Behavioral and Brain Sciences* 26, 48–49.

Thompson, B. 2010. The spatial content of experience. *Philosophy and Phenomenological Research* 81, 146–184

Thompson, E. 2007. *Mind in Life*. Cambridge, MA: Harvard University Press.

Tjian, R. 1995. Molecular machines that control genes. *Scientific American* 308, 54–61.

Tononi, G. 2007. The information integration theory of consciousness. In M. Velmans and S. Schneider (eds.), *The Blackwell Companion to Consciousness*. Oxford: Blackwell, 287–299.

2008. Consciousness as integrated information. *Biological Bulletin* 215, 216–242.

Tononi, G. and C. Koch. 2015. Consciousness: here, there and everywhere? *Philosophical Transactions of the Royal Society B* 370, 20140167.

Trestman, M. 2013. The Cambrian explosion and the origins of embodied cognition. *Biological Theory* 8, 80–92.

2014. The modal breadth of consciousness. *Philosophical Psychology* 27, 843–861.

Trout, J. D. 2001. The biological basis of speech. *Psychological Review* 108, 523–549

2016. *Wondrous Truths*. New York: Oxford.

Turin, L. and T. Sanchez. 2008. *Perfumes, the A–Z Guide*. London: Profile.

Tye, M. 1995a. A representational theory of pains and their phenomenal character. *Philosophical Perspectives* 9, 223–239.

1995b. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.

2000. *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.

2003. *Consciousness and Persons*. Cambridge, MA: MIT Press.

2009. *Consciousness Revisited*. Cambridge, MA: MIT Press.

2017. *Tense Bees and Shell-Shocked Crabs*. New York: Oxford.

2021. *Vagueness and the Evolution of Consciousness*. Oxford: Oxford University Press.

Valberg, A. 2001. Unique hues: an old problem for a new generation. *Vision Research* 41, 1645–1657.

Van Cleve, J. 2014. Berkeley, Reid, and Sinha on Molyneux's question. In D. Bennett and C. Hill (eds.), *Sensory Integration and the Unity of Consciousness*. Cambridge, MA: MIT Press, 193–208.

Van Gelder, T. J. 1995. What might cognition be, if not computation? *Journal of Philosophy* 92, 345–381.

1998. The dynamical hypothesis in cognitive science, with commentaries. *Behavioral and Brain Sciences* 21, 615–665.

Van Gulick, R. 2007. What if phenomenal consciousness admits of degrees? *Behavioral and Brain Sciences* 30, 528–529.

Van Ness, H. C. 1969. *Understanding Thermodynamics*. New York: McGraw-Hill.

Varela, F. J., E. Thompson, and E. Rosch. 1991. *The Embodied Mind*. Cambridge, MA: MIT Press.

Wagner, M. 2006. *The Geometries of Visual Space*. New York: Psychology Press.

Walters, D. 1996. How are bitter and sweet related? *Trends in Food Science and Technology* 7, 399–403.

Weiskrantz, L. 1986. *Blindsight.* Oxford: Oxford University Press.

Weng, G., U. S. Bhalla, and R. Iyengar. 1999. Complexity in biological signaling systems. *Science* 284, 92–95.

Werblin, F. and B. Roska. 2007. The movies in our eyes. *Scientific American* 320, 72–78.

White, S. L. 1986. The curse of the qualia. *Synthese* 86, 333–368.

2010. The property dualism argument. In R. Koons and G. Bealer (eds.), *The Waning of Materialism.* Oxford: Oxford University Press, 89–113.

Wilkes, K. 1988. *Real People.* Oxford: Oxford University Press.

Wilson, D. A. and R. J. Stevenson. 2006. *Learning to Smell.* Baltimore, MD: Johns Hopkins University Press.

Wiltshire, J. 2014. *The Hidden Jane Austen.* Cambridge: Cambridge University Press.

Yoshioka, T., B. M. Dow, and R. G. Vautin. 1996. Neural mechanisms of color categorization in areas V1, V2, and V4 of macaque monkey cortex. *Behavioral Brain Research* 76, 51–70.

Young, B. D., A. Keller, and D. Rosenthal. 2014. Quality-space theory in olfaction. *Frontiers in Psychology* 5, 1.

Zwicher, E. and Fastl, H. 2006. *Psychoacoustics*, third edition. New York: Springer.

Index

https://doi.org/10.1017/9781009514774.010