



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Intentionality from Structure

A Non-Relational Account of the Constitution of Phenomenal Intentionality

Blaine Kenneally

PhD Philosophy
The University of Edinburgh
2018

Declaration

I declare that this thesis has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application for a degree. Except where stated otherwise by reference or acknowledgment, the work presented is entirely my own.

Signed

Date

Blaine Kenneally

Abstract

In recent decades, there has been a growth in the popularity of the view that there is a form of intentionality which is grounded in phenomenal consciousness. Some philosophers have argued that this form of intentionality is non-relational in nature. In this thesis, I consider what phenomenal conditions might ground a non-relational form of intentionality. I argue that this non-relational phenomenal intentionality is constituted by the structural organisation of phenomenal qualities.

In Part I, I consider the mistakes we are prone to, and should avoid, when it comes to theorising about phenomenally conscious experience. I show that, in the debate over the 'transparency' of experience, disputants have overlooked the crucial distinction between metaphysics and phenomenology, which we should take care to respect. I then show that disputants have also conflated two different sorts of 'seeming', and as such have made a mistake about what is 'given' to us in experience. I explain in detail why we are susceptible to mistakes about what is 'given', in order that we might take steps to avoid them.

With the avoidance of these errors in mind, in Part II I assess six suggestions as to what phenomenal conditions are constitutive of phenomenal intentionality. I identify, as a promising candidate, the suggestion that phenomenal intentionality is constituted by the structural organisation of more basic phenomenal elements. I then consider in more depth how mistakes about what is 'given' might be made in relation to phenomenal intentionality. I argue that the structural organisation account carries minimal risk of such errors, given that it commits only to phenomenological claims which are broadly agreed upon.

In Part III, I flesh out the structural organisation account, articulating my own more detailed theory as to how non-intentional elements of phenomenal character—phenomenal qualities—come to constitute phenomenal intentionality when structurally organised in the appropriate way.

Lay Summary

As you read these words, there is something it feels like for you to read them. There is a way it feels to see the black of the ink, a way it feels to see the white of the paper, and a way it feels to understand these words. These feelings are what philosophers call 'phenomenal consciousness' or 'phenomenality'. At the same time, your experience of reading these words is about these words, and what they mean. If you peel an orange, there will be a way it feels to you to see its orange colour, to touch its skin, and to smell its citrusy aroma, and your experience will also be about the orange you are peeling. This feature, of being about something, is what philosophers call 'intentionality'.

In recent decades, there has been a growth in the popularity of the view that our experiences can be about things just by virtue of how they feel—the view that there is a form of intentionality which is grounded in phenomenality. This form of intentionality has become known as 'phenomenal intentionality'.

Traditionally, it has been thought that, in order for one thing to be about another, there must exist a relationship between those two things. However, some philosophers have argued that an experience with phenomenal intentionality can be about something without bearing any relationship to it. In this thesis, I try to explain how this could be so. I try to explain how experiences could manage to be about things they bear no relationship to, just by virtue of how they feel.

I propose that the way experiences feel can be broken down into component parts. Individually, these parts are not about anything. But when these parts are combined together, and structurally organised in the right sort of way, they make up

experiences which are about things. In this sense, the component parts of the way experiences feel are like paints. On their own, paints are not about things, but when structurally organised on a canvass in the right sort of way, they can make up paintings which are about battles, or people, or landscapes.

Acknowledgements

It takes a village to raise a child, and something approaching a village to write a thesis. It is high time I thanked the cohort of family, friends, supervisors and funders who made this achievable.

First and foremost, I would like to thank Rosie, who knew better than I the challenges I would face in writing this thesis, and who married me in the midst of them nonetheless. Rosie's love and support has propelled me throughout, transforming an intimidating objective into a practical goal. It is no overstatement to say that this has been a joint endeavour.

I have benefitted enormously from the supervision of Alistair Isaac, whose pragmatic advice enabled me to stop worrying about plans and start writing prose. Alistair's patient encouragement helped me wrestle these words onto the page, and his perspicuous feedback has helped me to hone them. I am also grateful to Rob Rupert for incisive feedback and fruitful discussions, and to Jesper Kellestrup for support in the early stages and encouragement at difficult junctures.

I am grateful to my parents, Elaine and Shaun, and my sisters, Heidi and Fleur, for giving short shrift to my overblown worries, and for providing an unconditional bedrock of support throughout. That bedrock of support has been extended northwards by the Hydes. Alison, Andy, Anna and Phoebe have been a second family to me, and the warmth and humour of the Hyde household (not to mention Alison's cake) has helped lay many an anxiety to rest.

The momentum of this project benefitted enormously from time and conversation with friends. Although too numerous to mention, I am grateful to you

all. In particular, I am indebted to Peter 'Iolaus' Kukiewicz, who had the patience to vet my drafts with his acute grammatical eye, and who thoughtfully fielded my concerns in the final stages. I also owe special thanks to Suki Finn, who descended from the attic of Kyme Street to demystify postgraduate philosophy in the first instance, and who has been something of a philosophy life-coach at challenging junctures ever since. Thanks also goes to Darren Watkins, Marlowe Heywood-Thornes and Tom Kelsey for tolerating my long and rambling telephone conversations, to Becky Millar for cathartic conversations over coffee, and to the rest of the SGSAH philosophers' table—Claire Hamlett, Claire Field, Nathan Kirkwood and Joanna Leidenhag—for solidarity in the journey.

I was fortunate to be able to attend the 2016 'Matter, Mind and Consciousness' Summer University course at the Central European University in Budapest, organised by David Pitt. I am grateful to the lecturers and participants for contributing to a stimulating fortnight of philosophy, which galvanised the development of many of the ideas presented here. I am also grateful to the Scottish Graduate School for the Arts and Humanities, whose Student Development Fund generously covered most of the costs incurred, and to the University of Edinburgh School of Philosophy, Psychology and Language Sciences, whose Research Support Grant covered the remainder.

This project has been funded by the Arts and Humanities Research Council, through a studentship with the Scottish Graduate School for Arts and Humanities. I am immensely grateful for this financial support, and to SGSAH for providing additional opportunities along the way.

Finally, I would like to dedicate this thesis to the memory of my grandparents, Janet and Gordon Briggs, and my grandparent-in-law, Valerie Ayton. Their wise and mischievous smiles will not be forgotten.

Table of Contents

1. Introduction	17
1.1 Phenomenality and Intentionality	17
1.2 Separatism, and Three Types of Inseparatism.....	19
1.3 The Appeal of Relational Representationalism	25
1.4 Motivating a Non-Relational Phenomenal Intentionality View.....	28
1.4.1 The Relatum Problem	28
1.4.2 The Presence Problem	31
1.4.3 Reasons for Pursuing a Non-Relational Phenomenal Intentionality View	34
1.5 The Project of the Thesis.....	38
PART I	43
2. Distinguishing Between Phenomenology and Metaphysics	43
2.1 The Qualia Debate	44
2.1.1 Phenomenal Qualities	44
2.1.2 The Location Problem with Respect to Phenomenal Qualities	46
2.2 Claims About How Experience ‘Seems’ in the Dispute Over Transparency	50
2.3 The Relevance of What is Presented in Experience	55
2.3.1 Elucidating ‘Awareness’ and ‘Attention’	55
2.3.2 The Epistemic Dimension of the Qualia Debate	60
2.4 The Conflation of Phenomenology with Metaphysics.....	64
2.5 Conclusion.....	71
3. Avoiding Errors About What is ‘Given’ in Experience	73
3.1 Phenomenology as a Way of Establishing What Is ‘Given’	73
3.2 The ‘Given’ as What is Experientially Manifest	75
3.3 Mistakes Concerning What Is Given, arising from the Conflation of Two Kinds of ‘Seeming’ or ‘Appearance’	79
3.4 Distinguishing Between Appearance and Reality in Both Phenomenology and Metaphysics	91
3.4.1 Questions to Address	91
3.4.2 Dealigning the Phenomenology-Metaphysics and Appearance-Reality Distinctions....	95
3.4.3 Recognising Phenomenally Conscious Experience as a Sort of Appearance or Seeming	98
3.4.4 Errors due to Mistaken Alignment of the Phenomenology-Metaphysics and Appearance-Reality Distinctions	103
3.5 Conclusion.....	110
PART II	113
4. What is the Phenomenological Signature of Directedness?	113
4.1 Desiderata for an Account of the Phenomenological Signature of Directedness .	114
4.2 Phenomenality Itself	116
4.3 Sui Generis Phenomenal Directedness	125
4.4 ‘Taking’, Cognitive Intent, or Cognitive-Experiential Character	132
4.5 Transparency.....	136
4.6 Phenomenal Presence	140

4.7 A Further Problem with <i>Sui Generis</i> Candidates for the PSD	151
4.8 The Structural Organisation of Experience	156
4.9 Conclusion	167
5. Minimising the Risk of Error about how Phenomenal Intentionality is ‘Given’	169
5.1 The Grey-Zone of Seemings	170
5.2 The High-Level / Low-Level Dispute	172
5.3 Distinguishing What is Experientially Manifest from What Is Perceived	175
5.4 Phenomenal Intentionality in The High-Level / Low-Level Dispute	181
5.5 The Point of Implicit Agreement	185
5.5.1 Three Stalemate-Inducing Arguments	186
5.5.2 The Agreement Required for the Stalemates	190
5.6 Conclusion	197
PART III.....	199
Introduction to Part III.....	199
6. Learning from Depictive Artefacts	205
6.1 What is Depiction?	205
6.1.1 Pictorial Depiction	206
6.1.2 Depiction in General	211
6.2 Depiction Depends Upon Experienced Structure	215
6.2.1 Dependence Upon Experienced Structure in Pictorial Depiction	215
6.2.2 Dependence Upon Experienced Structure in Depiction in General	224
6.3 Structured Experience as an Explanation of Depiction’s Dependence on Experienced Structure	229
6.4 Conclusion	234
7. Depicting with Mental Paint.....	237
7.1 Impressions as Phenomenal Depictions.....	237
7.2 The Mental Paint Palette.....	241
7.2.1 Phenomenal Qualities as Exemplar Mental Paints	241
7.2.2 Quality Space Theory	244
7.2.3 Drawing Upon Quality Space Theory	251
7.3 Painting with Mental Paint.....	257
7.3.1 From Mental Paints to Impressions	257
7.3.2 The Spatiotemporal Form of Perceptual Experience.....	260
7.3.2.1 Temporal Form	260
7.3.2.2 Spatial Form.....	264
7.3.3 The Formation of Impressions	270
7.4 Conclusion	274
8. Beyond Perceptual Impressions	277
8.1 Accounting for Taking.....	277
8.2 Extending the Account Beyond Perception.....	285
8.2.1 Bodily, Emotional and Mood Experiences	285
8.2.2 Cognitive Experiences	290
8.3 Conclusion	296

9. Conclusion	297
9.1 Summary	297
9.2 Areas for Development and Further Inquiry	300
9.2.1 Incorporating Tracking Relations	300
9.2.2 A More Nuanced Origin Story for Intentionality in General	301
Bibliography	303

1. Introduction

1.1 Phenomenality and Intentionality

When we human beings wake up in the morning, there is something it feels like to be us. There is something it is like for us to hear the piercing sound of a digital alarm, and there is something it is like for us to open the curtains and see the street, bathed in morning light. There is something it is like for us to smell coffee, and something it is like to feel nervous about the day ahead. This fact, of there being ‘something it is like’ for us to be us, was famously identified by Nagel (1974) as what it means to have conscious experience.

Understood in this way, conscious experience is sometimes called ‘phenomenal consciousness’, to emphasise that the consciousness in question involves ‘phenomena’ or appearances—the *feeling* or *experiencing* of things. Block (1995), for example, distinguishes “phenomenal consciousness” from “access consciousness”. A mental state is *phenomenally* conscious if there is something it is like to be in that state, and *access* conscious if it is “availab[le] for use in reasoning and rationally guiding speech and action” (p. 227). Throughout this thesis, I will only be concerned with phenomenal consciousness,¹ and I will refer to it interchangeably as ‘experience’, ‘conscious experience’ and ‘phenomenally conscious experience’, as and when appropriate.

In this thesis, I am concerned with the relationship between two features or aspects of conscious experience: its *phenomenality* and its *intentionality*.

¹ I follow Montague (2016, p. 13) and Strawson (2008, p. 281), in this regard.

Phenomenality is an essential feature of conscious experience, as just defined. For to have phenomenality is precisely to carry some feeling for a subject. As such, 'phenomenality' will simply serve as a useful term for referring to that feeling-imparting characteristic of conscious experience, which is absolutely essential to its nature. As well as being an *essential* or *necessary* feature of conscious experience, phenomenality is also *sufficient* for conscious experience. If a mental state has phenomenality, then it carries some feeling for the subject, and hence involves conscious experience.

Intentionality is the characteristic of being *about* something. Sometimes, intentionality is equated with the very general phenomenon of *anything at all* being about something, often called 'aboutness'.² A photograph of Arran, a map of Skye, and the name 'Cumbrae' are all instances of intentionality in this broader sense, because they are all respectively *about* some particular Scottish island.³ Some philosophers reserve the term 'intentionality' for the more specific phenomenon of the mind or mental states being about things.⁴ For clarity, I will follow this more

² For example, Dennett and Haugeland (2004) and Montague (2010).

³ It seems obvious that photographs, maps and names manage to *be about* things in very different *ways*—it is at least partly by virtue of their different ways of being about things that we can distinguish the three from one-another, even when they are all about the same thing. 'Aboutness' should therefore be considered a general *kind* of phenomenon of which there are many more-specific instances; something can *be about* something by being *of* it, *representing* it, *depicting* it, *standing for* it, *indicating* it, *referring to* it, *expressing* it or *meaning* it, and these by no means exhaust the ways.

⁴ For example, Jacob (2014) writes that "[i]ntentionality is the power of minds to be about, to represent, or to stand for, things, properties and states of affairs." Similarly, Searle (1983, p.1) writes that "[i]ntentionality is that property of many mental states and events by which they are directed at or about or of objects and states of affairs in the world." The distinction between 'intentionality' and 'aboutness' is particularly important for Strawson (2008), who holds that "[t]here is a lot more aboutness in the

restricted usage, using 'intentionality' only ever to refer to the aboutness of the mind or of a mental state, whilst using 'aboutness' to refer to any instance of something being about something, including, sometimes, instances of intentionality. Although it is controversial as to whether intentionality is an essential or necessary feature of conscious experience, it is clear that some conscious experiences are about things. When I open my curtains and see the street, my conscious experience is about the street. As you read these words, your conscious experience is about these words.

Broadly speaking, this thesis is concerned with how the *phenomenality* and *intentionality* of conscious experience are related to one another. More specifically, it addresses the following question: if the intentionality of conscious experience is constituted by its phenomenality, then how does its phenomenality manage to constitute its intentionality? In this chapter, I will motivate and clarify this question, before laying out precisely how I will go about answering it. Roughly, my answer is that a conscious experience's phenomenality manages to constitute its intentionality by being structurally organised in the appropriate way.

1.2 Separatism, and Three Types of Inseparatism

As Kind (2010, pp. 902-903) and Montague (2016, pp. 21-23) attest, much of the twentieth century saw philosophers treating phenomenality and intentionality as entirely separate and independent aspects of the mind. On the one hand, it was considered possible to have mental states with intentionality, such as beliefs and

world ... than there is intentionality, for although all intentionality involves aboutness, not all aboutness involves intentionality."

thoughts *about* things, in the absence of any phenomenality.⁵ On the other hand, it was considered possible to have mental states with phenomenality, such as sensory experiences of colour and flavour, or bodily sensations of pain or itchiness, in the absence of any intentionality.⁶ And in so far as perceptual states, such as visually perceiving a street, were considered to have both intentionality (in being about the street) and phenomenality (in involving sensory experience of colour, texture, shape, etc.), those two aspects were considered to be distinct and only contingently associated.⁷ But in recent decades, this *separatism* about the phenomenal and the intentional has lost popularity, giving way to a now widely endorsed *inseparatism*, which maintains that phenomenality and intentionality are not entirely independent of one-another.⁸

⁵ For example, Searle (1994, p. 380) states that “[s]ome intentional states are conscious, some not. The belief that George Washington was the first president can be consciously entertained, but a person can have that belief while he or she is sound asleep. In such a case, the intentional state is unconscious.” Similarly, Rosenthal (2009, p. 246) states that “...thoughts and desires have intentional content independently of their being conscious”.

⁶ For example, Searle (1994, p. 380) states that “...a feeling of pain or a sudden sense of anxiety, where there is no object of the anxiety, are not intentional.” Similarly, Rosenthal (1994, p. 349) states that “sensations, such as pains and sense impressions, lack intentional content, and have instead qualitative properties of various sorts.”

⁷ Kind (2010, p. 903), Smithies (2013, p. 746) and Bordini (2017, p. 1108) attest to this. The view that there is only a *contingent* relationship between the intentionality and phenomenality of such states is called “the ‘medium conception’” by McGinn (1988, p. 236), and is summarised well by Zahavi (2013, p. 83): “According to this view, the relation between consciousness and intentionality is like the relation between a medium of representation and the message it conveys ... consciousness is nothing but a (rather mysterious) medium in which something relatively mundane, namely intentionality, is contingently embedded”.

⁸ I borrow the term “separatism” from Horgan and Tienson (2002, p. 520). Kriegel (2013, p. 5) characterises “inseparatism” more strongly than I do, as the view that phenomenality and intentionality are “inseparably intertwined”. I employ a weaker characterisation in recognition of the fact that there *sometimes* being dependence

Although the motivations for the various kind of inseparatism are many and varied, reflection on one's own experiences provides some degree of motivation. For example, Horgan and Tienson (2002, p. 521) identify some of the phenomenal qualities typically involved in perceiving an apple: the red colour we see, the tactual smoothness we feel, and the sweetness we taste. They point out that these phenomenal qualities each present to us a different quality *of* the apple, and together present to us a single object—the apple—in which those qualities are unified. In other words, the experience is *about* a red, smooth, sweet apple purely because of its phenomenal character — it *says* certain things about the world just by *feeling* a certain way. And this extends to our ongoing perceptual experiences of the entire world around us: the phenomenal qualities we experience present to us the qualities and objects of the world. Given that the intentionality of perceptual experiences appears to be determined by and dependent upon phenomenality in this way, such experiences clearly lend support to an inseparatist view.

Relatedly, Siewert (1998, p. 12 & pp. 188-194) points out that conscious experiences, purely in virtue of their phenomenal character, are often assessable for accuracy. That is, the phenomenal character of an experience will often entail accuracy conditions—conditions which would need to be fulfilled in order for the experience to be an accurate one. The phenomenal character of the experience

relationships between phenomenality and intentionality does not preclude them from *ever* being independent. Hence, one can consistently believe that phenomenality and intentionality are *not entirely separate*, and hence deny separatism, without believing that they are *inseparably intertwined*. Kriegel's stronger characterisation misses this.

described in the previous paragraph, for example, entails, at the very least, the condition that there must be something red, smooth and sweet, in order for the experience to be accurate. As such, phenomenal character can effectively stipulate a way the world should be according to it, and in this sense is *about* the world being that way, and thus has intentionality just in virtue of its phenomenal character.

McGinn (1988) summarises this ostensible interconnectedness of phenomenality and intentionality point very vividly, describing perceptual experience as “Janus-faced”:

“...perceptual experiences are Janus-faced: they point outward to the external world but they also present a subjective face to their subject: they are of something other than the subject and they are like something for the subject. But these two faces do not wear different expressions; for what the experience is like is a function of what it is of, and what it is of is a function of what it is like.”

(McGinn, 1998, p. 231)

In sum, we have good reason for upholding some form of inseparatism about the phenomenal and the intentional. Given that I am concerned with the relationship between phenomenality and intentionality as characteristics of *conscious experience*, I am going to set aside concerns about intentionality outside of conscious experience.⁹ With these concerns set aside, we can distinguish three types of

⁹ As such, I am not going to consider broader questions about the source or origin of *all* intentionality, the extent to which intentionality permeates the world, what constitutes derived and underived intentionality, and how any derived intentionality manages to be derived. See Kriegel (2013) for an excellent survey of the positions one might adopt on these issues.

inseparatism about phenomenality and intentionality as they feature in conscious experience.

The first type of inseparatism, known as *representationalism*, takes phenomenality to be reducible to intentionality.¹⁰ That is, it takes intentionality to be more fundamental, or more basic than, phenomenality, such that phenomenality *depends upon* intentionality, but not vice versa. This is often expressed in terms of the ‘phenomenal character’ of an experience—the particular character of its phenomenality—being dependent upon its ‘representational content’, or what the experience is about. For example, if one were to consciously experience looking at and biting into a lime, the phenomenal character of the experience would include *greenness* and *bitterness*, and the representationalist would explain this in terms of the experiential mental state representing the properties of being green and being bitter.¹¹

¹⁰ From here onwards, I employ the taxonomy of positions used by Bordini. Although this taxonomy makes sense for my present purposes, it is worth noting that not everyone taxonomises views in this way. The distinction Bordini and I draw between *representationalism* and *egalitarianism* is recognised by Chalmers (2004, p. 178) as the distinction between “reductive representationalism” and “non-reductive representationalism”.

¹¹ This is a somewhat crude characterisation, but it nonetheless captures the thrust of the positions held by Harman (1990), Dretske (1995), Lycan (1996) and Tye (1995, 2000). Note that I have set aside a distinction drawn by Chalmers (2004), Bourget and Mendelovici (2014) and Mendelovici (2014) between *pure* and *impure* representationalism. *Pure* representationalism takes phenomenal character to reduce to representational content *alone*, whilst *impure* representationalism takes phenomenal character to reduce to representational content *together with some other features*, such as representing in a certain manner. Delving into the intricacies of this distinction will not be necessary for my purposes. Note also that representationalism is sometimes called ‘intentionalism’, as in Mendelovici (2014).

The second type of inseparatism is what Bordini (2017, p. 1109) calls *egalitarianism*. On this sort of view, upheld by Chalmers (2004) and Pautz (2008, 2013), intentionality and phenomenality are mutually dependent, but neither reduces to the other.

The third type of inseparatism is the *phenomenal intentionality view* or PIV, according to which intentionality is reducible to phenomenality. PIV takes phenomenality to be more fundamental, or more basic than, intentionality, such that intentionality *depends upon* phenomenality, but not vice versa. Again, this view can be expressed in terms of the ‘representational content’ of an experience being dependent upon its ‘phenomenal character’. According to PIV, if one were to consciously experience looking at and biting into a lime, the experience would *be about* or would *represent* the properties of being green and being bitter, and this would be explained by the fact that the experience’s phenomenal character involved *greenness* and *bitterness*.¹²

We can summarise the positions as follows. For the representationalist, how experiences feel depends upon what they are about. For the PIV theorist, what experiences are about depends on how they feel. And for the egalitarian, what experiences are about, and the way they feel, are entirely on a par—they are simply different sides of the same coin.

¹² Although this is, again, a somewhat crude characterisation, it nonetheless captures the thrust of the positions held by Siewert (1998), Horgan and Tienson (2002) and Loar (2003/2017).

The question this thesis addresses relates to PIV. PIV is of interest to us, because it looks like the most promising form of inseparatism to endorse if one also wishes to endorse a non-relational view of intentionality. But what constitutes a non-relational view of intentionality? Why would one wish to endorse such a view? And why does PIV look like the most promising inseparatist way of doing so? In order to answer these questions, it will be useful to consider the appeal of a form of representationalism which takes a relational view of intentionality. With a relational form of representationalism on the table, it will be easier to motivate a non-relational form of PIV.

1.3 The Appeal of Relational Representationalism

As Davide Bordini (2017, pp. 1109-10) points out, representationalism has become a popular position amongst philosophers, plausibly because many hold that intentionality either *already has* or *will be* naturalised, and so hope to naturalise phenomenality by reducing it to intentionality. Let us unpack this motivation in more detail.

To 'naturalise' some phenomenon is to account for it in a manner consistent with the ontology of the natural sciences—that is, entirely in terms of scientifically respectable entities. Those who wish to uphold a scientifically respectable worldview, whilst also upholding the reality of some phenomenon, will therefore seek to naturalise that phenomenon.

Since the late 1970s, much of the work on intentionality has been part of a consolidated effort to naturalise it, which Uriah Kriegel (2013, p. 2) calls “the

Naturalist-Externalist Research Program, or NERP". Given that intentionality doesn't appear to feature amongst the basic physical components of the universe, it stands to reason that it must be *reducible to* or *made-up-of* those components.¹³ In the words of Mendelovici and Bourget (2014, p. 327), "intentionality is real but not fundamental, so it must ultimately be reducible to some combination of fundamental physical ingredients." NERP has, therefore, been engaged in a project of *reductive* naturalization. And as Montague (2016, p. 17) attests, there has been a consensus "that the way to reductively naturalize intentionality is to identify a natural relation that holds between states of the brain and states of the environment when and only when the former are about the latter." It is this outward-looking approach, of determining what brain states are about by looking at their relationship to states of the external environment, which makes NERP an 'externalist' research program. Although the precise nature of the relevant relation between the brain and its environment has been hotly disputed, it has been broadly agreed to be a 'tracking relation' of some sort,¹⁴ where the 'tracking' in question is "a matter of detecting, carrying information about, or otherwise corresponding with states or items in the environment" (Mendelovici and Bourget, 2014, p. 326).¹⁵

¹³ This point is made with brilliant clarity by Fodor (1987, p.97).

¹⁴ The term "tracking relation" is used by Kriegel (2013, p. 3), Mendelovici and Bourget (2014, pp. 326-327), and Montague (2016, p. 18).

¹⁵ More specifically, different tracking theories provide different ways of cashing out the same common schema: *a vehicle V represents or 'tracks' a property P just in case V bears relation R to P*. (A vehicle is what 'carries' the representational content, such as a brain state.) Different tracking theories provide different accounts of what R consists in. Stampe (1977), for example, suggests that the relationship is one of *P causing V under the right conditions*. Dretske (1981), on the other hand, holds that the relationship is one of V being *nomically dependent* on P, meaning that it is a law

If one thinks that NERP has successfully naturalised intentionality, or is well on the way towards doing so, then one may be motivated towards representationalism by the prospect of naturalising phenomenality. In other words, if one thinks that intentionality can be naturalised, then one may be tempted to try to reduce phenomenality to intentionality in order to secure naturalisation for phenomenality also, effectively killing two birds with one stone. The overall picture would be this: the phenomenal character of an experience reduces to the representational content of a mental state, which in turn reduces to tracking relations between brain states and states of the environment. One could, of course, be a representationalist, maintaining that phenomenality reduces to intentionality, without maintaining that intentionality reduces to tracking, in which case one would not be motivated to adopt representationalism in this particular form. In order to distinguish this particular form of representationalism from others, we can follow Bourget and Mendelovici (2014) in labelling it “tracking representationalism”.

Tracking representationalism constitutes a form of *relationism* about intentionality. As Bourget and Mendelovici (2017, §4.6) characterise it, “[r]elationism about intentionality is the view that intentionality is a relation to distinctly existing entities that serve as contents.” Tracking representationalism commits to relationism about intentionality in taking the content of a brain state to be whatever distinctly

of nature that V is not tokened unless P is instantiated. This account is augmented in Dretske (1988) with a teleological component, wherein V represents P if it is *supposed to* nomically depend upon P, in the sense that that nomic dependence is being utilised by an organism to guide its behaviour, ultimately as a means toward survival. Another account with a teleological element is due to Millikan (1984), whose core suggestion is roughly that V represents P only if it is the *evolved function* of V to do so.

existing entity that state bears a tracking relation to. As such, tracking representationalism's commitment to relationism about intentionality is an essential concomitant of its naturalistic appeal.

However, as we will now see, there are serious problems with relationism about intentionality. These problems motivate a non-relational view of intentionality, which PIV looks well-placed to provide.

1.4 Motivating a Non-Relational Phenomenal Intentionality View

Mendelovici (2018, §9.2, pp. 200-207) identifies two problems with a relational view of intentionality, which I will call *the relatum problem* and *the presence problem*.¹⁶ In line with the scope of my concerns, I will consider these problems only insofar as they apply to the intentionality of conscious experience. I will first explain these two problems, and explain how a non-relational view avoids them. Given this reason for preferring a non-relational view, I will then explain why PIV is perfectly placed to provide a non-relational account of the intentionality of conscious experience.

1.4.1 The Relatum Problem

A relational view of intentionality takes a conscious experience's intentionality to consist in a relation which the experiential mental state bears to

¹⁶ Mendelovici (2018, §9.2, pp. 200-207) calls these "the inventory problem" and "the real problem", respectively. Although these titles make sense in the context of Mendelovici's more in-depth discussion of these problems, '*the relatum problem*' and '*the presence problem*' make more sense for my purposes.

distinctly existing entities. The relatum problem is the problem of specifying what sort of distinctly existing entity that experiential mental state bears the relevant relation to—it is the problem of specifying what sort of relatum sits at the other end of the intentionality-constituting relation. The solution would be straightforward if our conscious experiences were limited to veridical perception, for we could maintain that our experiential states are related precisely to those concretely existing objects, properties and states-of-affairs which we perceive, or have perceptual experiences *about*. However, our intentional conscious experiences clearly stretch beyond veridical perception, such that we can have conscious experiences about things which do not concretely exist. We can imagine or dream about non-existent entities such as kelpies and gorgons, and we can have hallucinatory or illusory experiences about objects which aren't really there, properties which aren't really instantiated, and states-of-affairs which do not really obtain. This is problematic, because the relata of any relationship must all exist. So, if intentionality consists in a relation to distinctly existing entities, but in many cases cannot consist in a relationship to the entities which the experience is ostensibly about, then we need to explain what distinctly existing entities our experiential mental states are in fact related to.

A tempting answer is that our experiences are about abstract entities of some kind, such as uninstantiated properties, Platonic universals, sets of possible worlds, or propositions. However, this answer is deeply problematic, because we are looking for relata which definitely *exist*, and it is not at all clear that these abstract entities fulfil this basic requirement. Mendelovici (2018, p. 201) suggests that what our experiential mental state is related to might be clusters of independently existing

properties. The idea is that, although gorgons do not in fact exist, women and snakes do, and as such one might visualise a gorgon by virtue of relations between that experiential mental state and the respective properties of being a woman and being a snake, amongst others. Although it would be problematic, for the reason just mentioned, to hold that the various properties related to are Platonic universals, Mendelovici (Ibid.) suggests that the properties “might be taken to be tropes, which are particular ways that individual objects are[, or] Aristotelian universals, which are abstract properties that can be instantiated in multiple items and exist only in their instances.” But although these relata are not so problematic in the existential sense, they remain problematic in other ways.

Tropes are supposed to be particular instances of properties. So the idea would be that, in hallucinating, visualizing or dreaming about a cube in one’s hand, which isn’t really there, my experiential mental state is in part related to some particular instance of the property of being a cube. However, that instance cannot be the instance *as it appears to one*, of a cube in one’s hand, precisely because there is no such instance. Therefore, the instance one is related to must be some instance *other* than the instance as it appears. We need to specify a *real* instance of the property of being a cube, which one’s mental state is related to. But there are plausibly many instances of the property of being a cube. How do we specify which particular instance is the relevant relatum of the intentional relation? We seem to have merely pushed the relatum problem a step back.

Aristotelian universals may seem to solve the trope-specific relatum problem. An Aristotelian universal is, like a Platonic universal, an abstract property which can

be instantiated in multiple instances. However, unlike platonic universals, which exist independently of their instantiations, Aristotelian universals exist only in their instantiations. As such, if we take the relatum in the cube case to be the Aristotelian universal of being a cube, we do not encounter the problem of specifying *which* instance of being a cube is the relevant relatum—the relatum is instead the Aristotelian universal consisting in all of its multiple instantiations. However, it is extremely counterintuitive to think that, in hallucinating a cube, one has an experiential mental state which is related to *all instances* of the property of being a cube. Although Aristotelian universals may *technically* solve the relatum problem by specifying a relatum, they are far from offering a satisfying and convincing solution. As such, the relatum problem can still be considered a problem for relational views of intentionality in general.

A non-relational view avoids the relatum problem precisely because it does not take intentionality to consist in a relation. Rather, according to a non-relational view, whatever makes a conscious experience about what it is about must be some aspect of the conscious experience itself—some intrinsic feature of the conscious experience.

1.4.2 The Presence Problem

Setting aside the issue of exactly what sort of a relatum could possibly sit at the end of an intentionality-constituting relation, there remains a more fundamental problem with the capabilities of such a relation. The presence problem is the problem

of there being a disparity between what a mere relation is capable of, and what the intentionality of conscious experience manages to achieve.

Our intentional conscious experiences—our experiential mental states—make the things they are about *present* to us. When I open my curtains and have a conscious, visual experience *of* or *about* the street, for example, there is a clear sense in which the street is *present* or *there* for me, by virtue of my experiential state.¹⁷ And as Johnston (2007) explains, this is not something such states simply *happen* to do. Rather, it is an essential feature of those states:

“The connection between a state that makes something present and what it makes present is not an accidental feature of that state, it enters into the essence of that state, it partly defines what it is to be that state.”

(Johnston, 2007, p. 234)

However, it is difficult to see how any mere relation between a mental state and some distinctly existing entity could establish a connection of the sort Johnston describes, which enters into the essence of the mental state and partly defines its nature. Mendelovici (2018) puts this point particularly vividly, stating that:

¹⁷ Johnston (2007, p. 233) provides the following example of ‘presence’: “When one sees one's dogs running in the front yard, the whole content of the perceptual experience is of the dogs and their running being present in a certain way, a way that discloses something of the nature of the dogs and their running. THERE the dogs are, immediately available as objects of attention and demonstration, and as topics of one's further thought and talk.”

“[i]t’s hard to see how *any* relation ... can grab hold of items that exist distinctly and sometimes independently of us and make them available to us, or allow us to reach out and make some kind of epistemic contact with them.”

(Mendelovici, 2018, pp. 205-206, original emphasis)

One might be tempted to insist that the relation is an extremely special sort of relation, capable of precisely those feats which Johnston and Mendelovici judge any ordinary relation to be incapable of. However, this is an extremely *ad hoc* manoeuvre, implausible in the absence of any independent reasons for thinking that such a special relation could obtain.

It is more plausible that the connection between a state and what it makes present (i.e. its content) is so essential to the state’s nature simply because this so called ‘connection’ is really an intrinsic feature of the mental state. In other words, it is more plausible that the state makes present what it does just by virtue of its own nature. To adopt this view is, of course, to adopt a non-relational view of intentionality, according to which whatever makes a conscious experience about what it is about must be some aspect of the conscious experience itself—some intrinsic feature of the conscious experience.

1.4.3 Reasons for Pursuing a Non-Relational Phenomenal Intentionality View

Given that a non-relational view avoids the problems of the relational view, and given that the views seem equally plausible in other respects,¹⁸ we have good reason to prefer a non-relational view. PIV is perfectly placed to provide a non-relational account of the intentionality of conscious experience, because it explains intentionality in terms of phenomenality, and, as we will now see, we have strong intuitive grounds for taking phenomenality to be non-relational.

Phenomenality intuitively seems non-relational, in the sense that the phenomenal character of our conscious experience, at any given moment, seems to be an intrinsic property of us. By intrinsic, I mean that the instantiation of that property does not depend upon anything external to or distinct from us, but only upon what is internal to and constitutive of us.

To explain by analogy, experienced phenomenal character seems much more like the non-relational property of mass than the relational property of weight. Our weight changes depending upon the strength of gravitational field we are in, such that we would weigh less on Mars than we do here on Earth, given that Mars has a weaker gravitational field. By contrast, our mass depends only upon how much matter we, as an organism, are comprised of. So, whilst our weight is a property which depends upon something external to and distinct from us, our mass is a property which depends just upon what is internal to and constitutive of us as an

¹⁸ Mendelovici (2018, §9.3, pp. 206-29) considers five alleged virtues of a relational view over a non-relational view, and shows, for each alleged virtue, that a non-relational view in fact fares no worse than a relational view.

organism. We can change our weight just by changing our location (albeit somewhat drastically), but to change our mass, we need to lose or gain fat or muscle.

Just like mass, the phenomenal character we experience seems to be an intrinsic property of us, regardless of any relations we might bear to anything external to or distinct from us. The idea that experienced phenomenal character is a property of *us*, which *we* instantiate, is captured well by Nida-Rümelin (2016):

“Experiences are special kinds of events. Your blue experience as you look at the sea is an example. What does it consist in? The obvious answer is: it consists in your instantiation of a certain property. You are visually presented with blueness. The experience at issue consists in the fact that *you* have *that property* at the moment at issue.”

(Nida-Rümelin, 2016, p. 165, original emphasis)

The strong intuition that such properties of subjects are *intrinsic* or *non-relational* is elicited by the fact that we can conceive of subjects instantiating such properties in situations wherein they bear vastly different relationships to their environment, and situations wherein they lack any such relationships whatsoever. The first such situation is that of a brain in a vat.¹⁹ Imagine that a subject’s brain is, unbeknownst to them, relocated from their body and into a vat, and hooked up to a machine which provides exactly the same electrical inputs, and adjusts them on the basis on the brain’s outputs, such that the brain is in exactly the same situation, with regard to proximal inputs and outputs, as it would have been on an ordinary day

¹⁹ The brain-in-a-vat thought experiment is used to elicit such intuitions by Loar (2003/2017), Horgan and Tienson (2002), Horgan, Tienson and Graham (2004) and Kriegel (2014), to name a few.

situated in the body.²⁰ Ostensibly, the subject's phenomenally conscious experience will be exactly the same as it would have been on an ordinary day. The second such situation is that of a disembodied mind.²¹ Imagine that the subject's mind is instead relocated to an entirely different universe, where it simply exists alone in empty space or in the complete absence of space, but nonetheless undergoes the same phenomenally conscious experiences.

The conceivability of these scenarios draws out the intuition that so long as factors intrinsic to the subject—internal to their brain or mind—are kept the same, then the subject's phenomenally conscious experience will remain the same. The supreme differences in the relationships between the subject and their environment, even if we confine 'the subject' to their brain or mind, would not seem to make any difference to the subject's phenomenally conscious experience. Therefore, phenomenality seems to be an intrinsic, non-relational property of the subject. In sum, we have strong intuitive grounds for taking phenomenality to be non-relational.

Independently of this positive reason for thinking that phenomenality is non-relational, we also have a negative reason—a reason for doubting that phenomenality is relational. For a relational view of phenomenality suffers from the same problems as a relational view of intentionality. If the phenomenal experience

²⁰ This is my own preferred scenario for conveying the brain-in-a-vat thought experiment. Kriegel (2014, p. 161) and Horgan, Tienson and Graham (2004, p. 297) have their own preferred variations—appealing to possible worlds and brain-synthesising aliens, respectively—but the differences are merely cosmetic.

²¹ Kriegel (2014, p. 164) describes this sort of scenario under the moniker of a "disembodied soul" or "space soul". Horgan and Tienson (2002) refer to this sort of scenario as a "disembodied Cartesian mind" or "Cartesian duplicate".

of visual *greenness* is a relation to some distinctly existing thing, then one again encounters *the relatum problem*, of specifying what exactly the distinctly existing relatum is. As with intentionality, the same problems arise with taking the relevant relata to be tropes, Aristotelian universals, or abstracta such as uninstantiated properties, platonic universals, possible worlds and propositions.²² One also encounters *the presence problem* for the second time, because it is essential to the very nature of phenomenal character that it *presents itself* to the subject. As with intentionality, it is extremely difficult to see how any mere relation between the subject, or their mental state, and some distinctly existing entity, could establish a connection which is essential to phenomenal character's presentational nature.

In sum, we have strong intuitive grounds for thinking that phenomenality is non-relational, and independent reason to doubt that it is relational. Given this, and given that we have reason for taking a non-relational view of intentionality, it makes sense that the inseparatist position we adopt is that of PIV, for it is perfectly placed to provide a non-relational view of intentionality by reducing intentionality to non-relational phenomenality.

Although egalitarianism might also seem like a viable option at this point, some egalitarian views, such as those of Pautz (2008, 2013), treat intentionality as relational. As such, given the problems with relationism, I think it is preferable to pursue PIV—to pursue an account according to which intentionality definitely

²² Papineau (2014, p. 8) highlights the particular problem with taking phenomenality to consist in a relation to abstracta, stating that: "My conscious sensory feelings are concrete, here-and-now, replete with causes and effects. How can their metaphysical nature essentially involve relations to entities that lie outside space and time?"

terminates in something non-relational, namely phenomenality, rather than an account where a non-relational termination is less well guaranteed.

1.5 The Project of the Thesis

In this chapter, we have established the motivation for pursuing non-relational PIV as an account of the relationship between the intentionality and phenomenality of conscious experience. That is, we have established the motivation for pursuing a form of inseparatism which takes a non-relational view of intentionality, by reducing intentionality to non-relational phenomenality. As such, we are now in a position to understand the motivation for undertaking the specific project of this thesis—a project which is neatly summarised by Masrour (2013):

"[The idea that] there is a kind of intentionality that is constituted by phenomenality ... implies that the conditions that a state needs to satisfy in order to count as a state with phenomenal intentionality are phenomenal conditions. One task for the phenomenal intentionalist is thus to provide an account of the phenomenal conditions that constitute phenomenal intentionality. We can call this the constitution problem."

(Masrour, 2013, p.117)

In this thesis, I am going to consider what phenomenal conditions might constitute the non-relational phenomenal intentionality of conscious experience. I will argue that this non-relational phenomenal intentionality is constituted by the structural organisation of phenomenal qualities.

Before we inquire as to what phenomenal conditions are constitutive of phenomenal intentionality, Part I considers the mistakes we are prone to, and should

avoid, when it comes to theorising about conscious experience. In Chapter 2, I highlight the importance of distinguishing phenomenology from metaphysics. I do so by showing how, in the debate over the 'transparency' of experience, disputants have overlooked this crucial distinction, and as such have given phenomenological claims undue metaphysical weight. In Chapter 3, I show that disputants have also conflated two different sorts of 'seeming', and as such have made mistakes about what is 'given' to us in experience. I explain in detail why we are susceptible to mistakes about what is 'given', in order that we might take steps to avoid them.

Part II considers some of the options as to what phenomenal conditions are constitutive of phenomenal intentionality. In Chapter 4, I assess six suggestions as to what constitutes the 'phenomenological signature of directedness'. I identify, as a promising candidate, the suggestion that phenomenal intentionality is constituted by the structural organisation of more basic phenomenal elements. In Chapter 5, I assess this suggestion in light of the lessons of Part I, regarding the mistakes we are liable to make in what we take to be 'given' in conscious experience. I consider how these mistakes might be made specifically in relation to how we take phenomenal intentionality to be 'given', and assess whether the structural organisation account manages to avoid these mistakes. I argue that the structural organisation account carries minimal risk of error concerning what is 'given' in experience, because it only commits to claims about what is 'given' which are broadly agreed upon, and is neutral regarding more contentious claims.

Part III builds upon the core structural organisation account, fleshing it out in greater detail. I articulate an account of how non-intentional elements of

phenomenal character—phenomenal qualities—come to constitute phenomenal intentionality when structurally organised in the appropriate way. I identify *impressions* and *taking* as the two aspects of phenomenal intentionality which need to be accounted for, and then account for both in terms of the structural organisation of phenomenal qualities. I argue that phenomenal qualities act as ‘mental paints’. Although mental paints are not in themselves about anything, they can, like paints on a canvas, be structurally arranged so as to depict things. I argue that *impressions* are structural arrangements of mental paints which manage to depict things, and that *taking* consists in the higher-order structural organisation of *impressions*. I do so as follows.

In Chapter 6, I elucidate the notion of depiction through consideration of depictive artefacts, and thereby lay the groundwork for treating impressions as phenomenal depictions. In Chapter 7, I clarify the sense in which impressions are phenomenal depictions, before showing it to be plausible that the phenomenal character of experience has the capacity to instantiate such depictions. I show that the brain can plausibly support a mental paint palette of the sort required to supply the basic phenomenal elements of all possible experiences, and then explain how these phenomenal building blocks are brought to instantiate the depictive structural arrangements which constitute *impressions*. In Chapter 8, I account for *taking* in terms of the higher-order structural organisation of *impressions* themselves. Then, having so far accounted for *impressions* and *taking* primarily in relation to perceptual experience, I show that the account can feasibly be extended beyond perceptual experience, to account for the constitution of phenomenal intentionality in all types

of phenomenally intentional experience. Finally, in Chapter 9, I conclude with a summary of the thesis, and identify some promising areas for further inquiry.

PART I

2. Distinguishing Between Phenomenology and Metaphysics

The overall aim of this thesis is to work towards an account of what phenomenal conditions constitute the non-relational phenomenal intentionality of conscious experience. As a first step, this chapter and the next will consider the mistakes we are prone to, and should avoid, when it comes to theorising about conscious experience. In this chapter, I will highlight the importance of distinguishing between phenomenology and metaphysics, when we make claims about phenomenally conscious experience. I will do so by critiquing a particular strand of the debate over the existence of qualia.

Qualia realism is a philosophical view according to which the phenomenal character of our experience is constituted by mind-dependent properties of experience itself. It is to these phenomenal-character-constituting properties of experience that the name 'qualia' refers, and it is the qualia realist's insistence on the genuine existence of such properties that earns the view its name.

A great deal of recent work in the debate over qualia has involved claims relating to a particular trait, feature or characteristic which experience is alleged to possess, namely the characteristic of being 'transparent' or 'diaphanous'. Purported observations of this characteristic intuitively militate against qualia realism, and are frequently deployed in arguments against it. Equally, qualia realists have attempted to defend their position by purporting to observe that experience is *not* transparent. These purported observations are expressed as claims about how things 'seem' or

‘appear’ in relation to experience. In this chapter, we will see that these claims are far less significant than the disputants take them to be. For these are *phenomenological* claims concerning how experience seems from a first-person perspective, and as such have little bearing on the *metaphysical* issue of whether qualia exist. This will show us the importance of distinguishing between phenomenology and metaphysics, when it comes to claims about phenomenally conscious experience.

I will first introduce the qualia debate, explaining the point of contention and describing the various positions held. Second, I will identify the paradigmatic claims about how things ‘seem’ or ‘appear’ in experience, deployed for and against qualia realism. Third, I will clarify precisely how these claims are supposed to bear upon the truth or falsity of qualia realism. Fourth, I will show how the use of these claims overlooks the crucial distinction between metaphysics and phenomenology, and hence is misguided.

2.1 The Qualia Debate

2.1.1 Phenomenal Qualities

In order to understand the matter of contention at the core of the qualia debate, one must first grasp what is meant by ‘phenomenal character’ and ‘phenomenal qualities’. If you consider, for example, your current experience of reading these words, the ‘phenomenal character’ of your experience is *what the experience is like*, for you, as you undergo it. All experiences have *some* phenomenal character, in the same way that we might say that all people have *some* sort of

character or personality. And just as ‘character’ or ‘personality’ does not refer to a way in which all people commonly comport themselves, ‘phenomenal character’ does not refer to what all experiences, commonly, are like to undergo. Rather, it is supposed to refer to what a *particular* experience is like to undergo, where what it is like to undergo *that* experience might differ from what it is like to undergo *other* experiences, just as one person’s character might differ from another’s. Nida-Rümelin (2008) provides a number of examples of the ways in which experiences might differ in terms of their phenomenal character:

“Smelling a rose is phenomenally different from smelling basil. Seeing something as slightly reddish blue is phenomenally different from seeing it as pure blue. Feeling one’s own leg without touching it and with closed eyes as being bent is phenomenally different from feeling it as being straight. Looking into the face of a person without recognizing her as one’s childhood friend is phenomenally different from looking into the same face under otherwise identical conditions after the recognition. Thinking of someone with love is phenomenally different from thinking of someone with emotional indifference...”

(Nida-Rümelin, 2008, pp. 309-10)

Whilst the singular ‘phenomenal quality’ can be used synonymously with ‘phenomenal character’, the plural ‘phenomenal *qualities*’ allows one to be more articulate in describing and analysing what it is like to undergo experiences. Let us take, for example, the experience of walking around a colourful garden, surrounded by birdsong. The overall phenomenal character or quality of the experience might be said to be *made up of* or *constituted by* a number of different phenomenal *qualities*—each with their own distinctive phenomenal quality, character or feel—such as the various colours in one’s visual field and the various sounds heard from the birds. It is

by no means incontestably obvious that the phenomenal character of an experience can be broken-down or analysed into such components, but as Coates and Coleman (2015, pp. 3-4) attest, “it seems natural to distinguish elements within [experiences] on account of their felt differences ... and the term [phenomenal qualities] is routinely employed to pick them out.”

2.1.2 The Location Problem with Respect to Phenomenal Qualities

At the centre of the qualia debate lies what Jackson (1998) calls a “location problem”, a type of problem summarised by Kriegel (2013, p. 4) as “the problem of finding a place for some phenomenon (typically familiar from the manifest image) in a description of the world cast entirely in some privileged vocabulary.” In the qualia debate, it is the phenomenal qualities just described that constitute ‘the phenomenon typically familiar from the manifest image of the world’. And to find a place for these qualities in ‘a description of the world cast entirely in some privileged vocabulary’ would be to find a place for them in some *ontological framework*.

‘Ontology’ is the study of *what there is*, or what exists, and *what it is like*. An ontological framework *privileges* a particular vocabulary in the sense that it aims to provide a comprehensive account “of what there is and what it is like ... *in terms of a limited number of more or less basic notions*” (Jackson, 1998, p. 5, emphasis added). In other words, ontological frameworks attempt to accommodate *everything* that exists *exclusively* in terms of a set of prioritised, basic ingredients. Therefore, any manifest phenomenon *not* in the list of basic ingredients must either be composed of or grounded in those ingredients, or not really exist. If we adopt materialism, for

example, where the basic ingredient is physical matter, we can only maintain the existence of ghosts if we accept that they are composed of or grounded in matter. Otherwise, we must deny their existence. Another ontological framework would be naturalism, according to which the basic ingredients are the objects, properties and events postulated by the natural sciences. Yet another would be dualism, according to which the basic ingredients are of two fundamentally different kinds: the physical and the mental. To *locate* some phenomenon within a given framework is to accommodate it amongst that framework's catalogue of existing things, either as one of the basic ingredients, or as something that can be built out of or grounded in them. On the other hand, to *eliminate* it is to banish it from the catalogue of existing things, maintaining that it is neither a basic ingredient nor a composition, nor something grounded in the basic ingredients.

The qualia debate is a dispute over how to solve the location problem for phenomenal qualities. All parties share the common goal of solving this location problem, but differ dramatically in the solutions they offer. *Qualia realism*, *representationalism* and *naïve realism* each propose a different way of locating phenomenal qualities within our ontology, whilst *eliminativism* maintains that they have no place in our ontology.

Naïve realism holds that experience puts us in direct relation with actual mind-independent properties of the world. Thus, the phenomenal qualities we experience when we perceive the world are properties of the actual world, which would exist regardless of the existence or state of our minds. For example, the phenomenal greenness involved in a visual perception of some grass would be

considered an actual property of the real grass out in the world, which perception places us in direct relation to.²³

By contrast, *representationalism* takes experience to *represent* actual mind-independent properties of the world. Thus, the phenomenal qualities we experience when we perceive the world are nothing over and above the ‘content’ of a representation. So, a representationalist would take the phenomenal greenness of the grass to be nothing more than the colour property the grass is represented as having.²⁴ A representationalist’s precise solution to the location problem depends upon which form of the view they uphold. *Internalist* representationalists think that everything required for the existence of phenomenal character, and hence phenomenal qualities, resides within the brain.²⁵ By contrast, *externalist* representationalists think that aspects of the wider world, such as properties in the world and one’s relation to them, are also required.²⁶ Hence, internalist representationalists fall into the broader category of *phenomenal internalists*, who take phenomenal qualities to be ‘in the head’, whilst externalist representationalists

²³ See, for example, Martin (2002) and Kennedy (2009).

²⁴ The view I here call ‘representationalism’ is referred to as “strong representationalism” by Kind (2003, p. 227; 2010, p. 904) and “strong intentionalism” by Crane (2000), in order to distinguish it from “weak” representationalism or intentionalism. Strong representationalists maintain that a mental state’s phenomenal character is nothing over and above its representational content, whilst weak representationalists merely hold that mental states with phenomenal character always have some representational content.

²⁵ Rey (1998) advocates internalist representationalism, but calls it “narrow” representationalism.

²⁶ Tye (1995, 2000, 2014), Dretske (1995) and Lycan (1996) all advocate externalist representationalism.

fall alongside naïve realists into the broader category of *phenomenal externalists*, who take phenomenal qualities to be ‘outside of the head’.

In distinction to both of these positions, *qualia realism* takes phenomenal qualities to be something other than actual properties out in the world, and something over and above the content of a representation. Qualia realists take phenomenal qualities to be properties of experience itself, and call such properties ‘qualia’ (with each individual property called a ‘quale’). On this view, the phenomenal greenness is not a property of the grass, nor is it the colour property the grass is represented as having in the content of a representation; it is instead a property of the experience itself.²⁷

Finally, *eliminativism* affords phenomenal qualities or qualia no place whatsoever in our ontology, on the grounds that our conception of them is somehow misguided or incoherent. For an eliminativist, phenomenal greenness is not something which actually exists, though it may illusorily appear to.²⁸

It is important to recognise that these proposed solutions to the location problem are accounts of the *metaphysics* of phenomenal qualities—they attempt to describe the true nature or reality of phenomenal qualities, in a way which either accommodates them within our ontology, or shows that they do not really exist. I will focus mostly upon the three *realist* positions on phenomenal qualities—those which try to accommodate phenomenal qualities within our ontology—for it is primarily these which engage in the transparency dispute.

²⁷ See, for example, Kind (2001, 2008), Hatfield (2007).

²⁸ See, for example, Dennett (1988).

2.2 Claims About How Experience 'Seems' in the Dispute Over Transparency

A great deal of recent work in the qualia debate has involved claims relating to a particular trait, feature or characteristic which experience is alleged to possess, namely the characteristic of being 'transparent' or 'diaphanous'. As Amy Kind (2010, p. 905) attests, the suggestion that experience possesses this characteristic dates back at least to the early modern period, as exemplified by the following excerpt from Thomas Reid:

"We are so accustomed to use the sensation [of hardness] as a sign, and to pass immediately to the hardness signified, that, as far as appears, it was never made an object of thought, either by the vulgar or by philosophers... There is no sensation more distinct, or more frequent; yet it is never attended to, but passes through the mind instantaneously, and serves only to introduce that quality in bodies, which, by a law of our constitution, it suggests."

(Reid, 1764/1970, pp. 61-62; quoted in Kind, 2010, p. 905)

But in recent work, the following passage from G. E. Moore (1903) is often cited as the paradigmatic articulation of the claim that experience is transparent:

"...the moment we try to fix our attention upon consciousness and to see *what*, distinctly, it is, it seems to vanish : it seems as if we had before us a mere emptiness. When we try to introspect on the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous."

(Moore, 1903, p. 450, original emphasis)

And the paradigmatic contemporary elaborations on the point are due to Gilbert Harman and Michael Tye:

“When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experience. And that is true of you too. ... When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree...”

(Harman, 1990, p. 39)

“Focus your attention on a square that has been painted blue. Intuitively, you are directly aware of blueness and squareness as out there in the world away from you, as features of any external surface. Now shift your gaze inward and try to become aware of your experience itself, inside you, apart from its objects. Try to focus your attention on some intrinsic feature of the experience which distinguishes it from other experiences, something other than what it is an experience *of*. The task seems impossible: one’s awareness always seems to slip through the experience to blueness and squareness, as instantiated together in an external object. In turning one’s mind inward to attend to the experience, one seems to end up concentrating on what is outside again, on external features or properties.”

(Tye, 1995, p. 30, original emphasis)

Abstracting from the commonalities amongst these excerpts,²⁹ the paradigmatic transparency claim amounts to the following assertion, which I will call *Seeming Transparency*, or ST:

²⁹ The paradigmatic status of these excerpts as expressions of the claim that experience is transparent is evident from their ubiquitous citation in the literature discussing transparency. See, for example, Kennedy, 2009; Kind, 2003, 2008, 2010; Nida-Rümelin, 2007, 2008; Rey, 1998; Siewert, 2004; Stoljar, 2004; and Tye, 2014.

ST Experience does not seem to present any properties of its own, but rather seems to exclusively present properties of things that the experience is of.³⁰

For concision, I will use the verb ‘present’ as an abbreviation for ‘provide one with awareness of’ and ‘enable one to attend to’. In the literature, ST is deployed in arguments against qualia realism and in favour of either representationalism, as in Harman (1990) and Tye (1995, 2000, 2014), or naïve realism, as in Kennedy (2009).

On the other side of the debate, qualia realists such as Ned Block and Amy Kind put forward a claim different from ST, which nonetheless relates to the transparency of experience. In contrast to what Kind (2008) calls the “mundane” examples which Tye and Harman use to support their claims, which exemplify relatively ordinary visual experiences—at least for those of us with a tendency to ogle trees and blue squares—Block and Kind assemble a collection of more “exotic”

³⁰ I take it to be obvious from Tye and Harman’s excerpts that ‘intrinsic features of experience’ refers to features or properties of experience itself, or features or properties which are experience’s own. I prefer to use the latter two phrasings in order to avoid unnecessarily committing qualia realists to the problematic notion of ‘intrinsicity’. One might object that the term ‘intrinsic’ is required to distinguish the qualia realist’s alleged non-representational properties *of experience itself* from what Tye (2014, p. 48) calls the “representational properties” of experience, which are nonetheless properties *of experience itself*, but not ‘intrinsic’ ones. I maintain that this distinction is already captured: the so-called ‘intrinsic’ properties are properties of experience itself which are themselves presented in experience, whereas the so called non-‘intrinsic’ representational properties are properties of experience which are *not* themselves presented in experience, but which represent those properties which *are* presented. In other words, so called ‘intrinsic’ properties are self-presenting—they make *themselves* available to attention or awareness—whereas so called non-‘intrinsic’ properties are other-presenting—they make something *other than themselves* available to attention or awareness.

experiences selected precisely because they are supposed to illustrate the relatively rare occasions on which experience is *less transparent* and *more opaque*; when experience seems to present properties of its own.

The visual examples Kind (2008, pp. 288-289) marshals are blurred vision, phosphene experiences, and afterimages, the latter two of which were formerly highlighted by Block (1996, pp. 31-32 & pp. 34-35). It is claimed that all three kinds of experience seem to present properties which belong to *the experiences themselves* rather than to something the experience *represents* or *relates us to*. First, someone with poor or unfocused vision would experience blurriness when looking at a newspaper headline. Kind claims that many have the strong intuition that attending to the blurriness is different from attending to the headline, and suggests that the blurriness feels like an aspect of one's experience, rather than an aspect of the headline represented to us, or to which we are related. Second, applying pressure to the retina can produce colour sensations known as phosphene experiences. Kind claims that, in attending to these, we do not seem to be attending to coloured objects or expanses which are represented to us or to which we are related, but rather seem to be attending to the experience itself. Third, there are afterimages which linger in one's vision after, for example, a bright camera-flash. Again, it is claimed that these do not seem to be actual things we are related to, nor to be objects or expanses represented to us. As Block (1996) attests, afterimages "don't look as if they are really objects or as if they are really red. They look... illusory" (p. 32, original ellipsis). In each example scenario, it is claimed that experience seems to provide us with

awareness of, and enables us to attend to, properties *of its own*, rather than properties that it *represents or relates us to*.

Block and Kind augment their battery of visual-experience examples with examples of more common non-visual experiences which lie “somewhere between the exotic ... and the mundane” (Kind, 2008, pp. 290). Block (1996, pp. 31-34) highlights pains and orgasms, as does Kind (2008, pp. 290-93), who also adds emotions and moods to the list. The force of these examples is supposed to come from the fact that, although these experiences often seem to present actual properties out in the world, or represented properties, they *also* seem to present qualitative properties of the experience itself, over and above the former properties. Block (1996, p. 33), for example, highlights the fact that we can be aware of and attend to the “hurting” of pain and the “phenomenally impressive” sensation of orgasm, neither of which, he claims, seems entirely reducible to an actual or represented bodily occurrence. Similarly, Kind (2008, pp. 292-93) claims that although emotions and moods occasionally seem to present actual or represented bodily occurrences—such as queasiness in the stomach in the case of nervousness, or an increased pulse rate in the case of anger—it seems far from the case that these are the *only* sorts of properties that emotions and moods present.

In sum, Block and Kind’s examples are deployed in support of another paradigmatic claim relating to the transparency of experience, albeit one borne of an opposing paradigm. I will call this *Seeming Opacity*, or SO:

SO Experience seems to present properties of its own, not merely properties of things that the experience is of.

2.3 The Relevance of What is Presented in Experience

As established in §2.1, qualia realism is an account of the *metaphysics* of phenomenal qualities, which maintains that phenomenal qualities are properties of experience itself. And as established in the previous subsection, advocates and deniers of qualia make claims about whether or not experience seems to present properties of its own, where ‘to present’ is ‘to provide one with awareness of’ or ‘to enable one to attend to’. In order to see more clearly how these claims are supposed to bear on the truth or falsity of qualia realism, we need to clarify what is meant by ‘awareness’ and ‘attention’, and recognise the epistemic dimension of the qualia debate.

2.3.1 Elucidating ‘Awareness’ and ‘Attention’

In relation to the transparency of experience, philosophers tend to employ ‘awareness’ and ‘attention’ as intuitive, pre-theoretic terms which enable them to articulate claims such as ST and SO. This is evident in the passages quoted in §2.2. Moore claims that “...the moment we *try to fix our attention* upon consciousness and to see *what*, distinctly, it is, it seems to vanish” (1903, p. 450, first emphasis added). Similarly, Harman (1990, p. 39) implores the reader: “try to turn your attention to intrinsic features of your visual experience”, and predicts that they will be unable to do so. Tye (1995, p. 30) also prompts the reader in this way: “try to focus your

attention on some intrinsic feature of the experience”, before asserting that “[t]he task seems impossible”. The implication is that ‘attend’ is a success verb; an action which one can try and fail to do, and which one can only be said to be properly doing if they are succeeding at it. Moreover, it is suggested that attention can be ‘fixed upon’ or ‘focused on’ particular aspects or features of an experience, and it is implied, at least by Moore, that focusing our attention upon something can put us in a better epistemic position with respect to that thing; a better position from which to gain knowledge about that thing; ‘to see *what*, distinctly, it is’. These implied characteristics paint a picture on which attention is conceived to be much like a mental spotlight, which one can wilfully—or at least *sometimes* wilfully, to *some extent* of wilfulness—direct onto the world and onto one’s own mind. One’s attention has a focus in the same way that a spotlight highlights a particular area, and attending to something places one in a better epistemic situation with respect to it, just as spotlighting something can help one to know more about it. Furthermore, attending and spotlighting are both actions which we can try and fail to do for a number of reasons: just as one may fail to spotlight the dancer who is moving too fast, one may fail to pay attention to instructions spoken too quickly.

‘Awareness’ is sometimes simply used as a synonym for ‘attention’. In his excerpt, for example, Tye switches back and forth between ‘attention’ and ‘awareness’ with no discernible difference in meaning; though Tye writes that “one’s awareness always seems to slip through”, and prompts the reader to “try to become aware”, he could have written that ‘one’s *attention* always seems to slip through’, and prompted the reader to ‘try to *attend*’, without any discernible difference in the

meaning of the passage. But 'awareness' is often reserved to refer to something one can have *without* thereby having attention, and so it is worth recognising the distinction. Block (2003), for example, puts forward an example in which the two notions intuitively come apart:

"...attention and awareness are distinct, ... One can be aware of what one is not attending to. For example, one might be involved in an intense conversation while a jackhammer outside causes one to raise one's voice without ever noticing or attending to the noise until someone comments on it—at which time one realizes that one was aware of it all along."

(Block, 2003, p. 171)

Given this distinction, it looks like 'awareness' is supposed to refer to a kind of wide-ranging, general and unfocused acquaintance with things, which is bestowed upon us when we undergo experiences. By contrast, 'attention' is the mental spotlight which is selectively shone upon a particular thing or group of things in one's awareness, in order to improve one's acquaintance with them, or place one in a better epistemic position with respect to them. This conception of awareness and attention can be broken down into two components. First, awareness is taken to be a *prerequisite* for attention; one's awareness of something is taken to be a *necessary requirement* for being able to attend to that thing, such that one's awareness of something is at least partly responsible for *enabling* one to attend to it. Second, one's attention is taken to have narrower scope than one's awareness, such that one will necessarily be aware of more than they are attending to.

Moore (1903) appears to have implicitly upheld the first component, asserting that “to be aware of a sensation of blue is ... to be aware of an awareness of blue” (p. 449), and subsequently claiming that this awareness of blue can itself “be distinguished if we look *attentively* enough, and if we know that there is something to look for” (p. 450, emphasis added). It is not unreasonable to infer that, on Moore’s account, one’s awareness of an awareness contributes towards one’s ability to attend to it. Moore’s advocacy of the second component is slightly more overt. In the case of an awareness of an awareness of blue, Moore thinks that one *can* attend to one’s awareness of blue if one looks attentively enough and knows that there is something to look for, but otherwise just ends up attending to the blue. The blue to which one otherwise ends up attending to is only a *part* of what one is aware of overall, and hence it is implied that attention has a narrower scope than awareness.

This two-component conception of attention and awareness also seems to have been adopted by Tye in his more recent work. Here he explicitly accepts Block’s aforementioned distinction, and advocates the first component in describing awareness as a ‘basis’ for attention:

“...attention and awareness are indeed distinct. Still, if one is aware of something, in standard visual cases, one can attend to the relevant thing directly on the basis of that awareness even if one does not do so. One can switch one’s attention/mental focus to the relevant thing directly on the basis of that awareness”

(Tye, 2014, p. 49, original emphasis)

The second component is implied by Tye’s assertion that one can ‘switch’ one’s attention to something on the basis of one’s awareness of it. To say that one’s

attention has been 'switched' implies that one has ceased to attend to one thing and is now attending to another. If, as Tye suggests, awareness of something affords one the ability to attend to it, regardless of whether or not one actually does so, then one's awareness includes *both* what is presently attended to *and* everything that one is not presently attending to, but has the ability to attend to by 'switching' one's attention. This suggests that the scope of one's attention is narrower than the scope of one's awareness, such that attention must switch between one thing or another rather than accommodating both at the same time. Though his stance on the first component is unclear, Harman might also be said to implicitly endorse the second component along similar lines to Tye, in describing attention as something which one can 'turn' from one thing to another.

There seems to be a significant degree of consensus, then, that awareness *is a prerequisite for, and has a wider scope than,* attention. Awareness is taken to be the most general kind of acquaintance one has with things in experience, which serves as a necessary basis for focusing in on a particular thing, aspect or feature by attending to it, which in turn improves one's acquaintance with that feature, or puts one in a better epistemic position with respect to it. In this vein, Brian O'Shaughnessy (1985, p. 202) writes that "it is the essential function, and hence an essential causal property, of the attention that it make[s] possible a knowing of its object". If we combine this knowledge-enabling conception of attention with the conception of awareness as a prerequisite for attention, then we can conceive of the joint operation of awareness and attention as a knowledge-enabling capacity; a means of epistemic access to the world, and potentially also to one's own mind.

It is also worth recognising that this conception of attention and awareness, which is employed in relation to the transparency or opacity of experience, is concerned with what Dennett (1969, p. 93) calls the “personal” level of explanation, in contrast to the “subpersonal” level of explanation. Whilst the former is the domain of explanations given in terms of “people and their sensations and activities”, the latter is the domain of explanations given in terms of “brains and events in the nervous system”. Those who make claims relating to the transparency of experience are concerned with ‘attention’ or ‘awareness’ only at a level at which that ‘attention’ or ‘awareness’ may be ascribed to a *person* or *subject*. Furthermore, they are concerned with that person or subject primarily as a *conscious being*, as opposed to a mechanistically-conceived living organism. Thus, their conception of attention is limited to the attention *of the conscious subject*, and does not include just *any* filtering of information in the brain, which might not manifest as the subject’s conscious attention. Likewise, their conception of awareness is limited to the awareness *of the conscious subject*, and does not include just *any* kind of sensitivity-to something, or disposition-to-react-to something, which might not manifest as conscious awareness. In sum, the use of these terms in the transparency dispute is limited to the domain of a subject’s conscious experience.

2.3.2 The Epistemic Dimension of the Qualia Debate

Having established precisely what is meant by ‘awareness’ and ‘attention’ in the transparency dispute, I will now consider why what we are aware of or able to attend to in experience is relevant to the truth or falsity of qualia realism. I will show

that qualia realism, representationalism and naïve realism all entail commitments about what experience provides one with awareness of or enables one to attend to. Because these commitments are concerned with our *knowledge-enabling* capacities of awareness and attention, they are *epistemological* commitments, and constitute what Kind (2001, especially pp. 149-151) calls the “epistemic dimension” of the qualia debate. Whilst Kind highlights the epistemological commitment of qualia realism alone, I will here extend her reasoning to representationalism and naïve realism.

The epistemological commitments in question follow only from *realist* accounts of the metaphysics of phenomenal qualities. That is, they follow only from those positions which propose ways of locating phenomenal qualities within our ontology, and hence hold that phenomenal qualities exist. The commitments do not arise for those who would rather *eliminate* phenomenal qualities from our ontology, and who therefore deny their existence. To see why this is so, it is helpful to recognise how the commitments stem from the very notion of phenomenal qualities.

Recall that, in §2.1.1, it was established that phenomenal qualities are *what an experience is like for a subject to undergo*. And in being defined as something an experience is like *for the subject*—something the subject *has the ability to know*, or *has epistemic access to*—these qualities are necessarily something that the experiencing subject is aware of and able to attend to. For example, if we take a subject’s visual experience of green grass to involve the phenomenal quality of greenness, then that greenness must necessarily be something that the subject is aware of and can attend to. A less obvious epistemic feature of phenomenal qualities is that the subject’s awareness of and attention to them must be *direct* or *non-*

inferential. This means that the subject is taken to be aware of the greenness, and able to attend to the greenness, *without having to infer* that there is greenness. In other words, the greenness is taken to be something that is immediately presented to the subject; the subject does not have to ‘work out’ that there is greenness, but is rather directly and immediately confronted with greenness. Furthermore, as Martine Nida-Rümelin (2008, p. 314) and Tye (2014, pp. 41-45) point out, ‘being aware of’ and ‘attending to’ a thing are *de re* descriptions; they describe a relationship between the subject and some particular thing.³¹ And it is of course necessary, for such a relationship to obtain, that both of the relata actually exist. Therefore, to hold that the subject of an experience *is aware of* and *can attend to* ‘what their experience is like’ is to commit to the *existence* of ‘what their experience is like’, and hence to commit to the existence of phenomenal qualities. All of the above is contained within the definition of ‘phenomenal qualities’, and hence a commitment to the existence of these qualities entails a commitment to the subject having direct epistemic access to them via awareness and attention.

The epistemological commitment to a subject’s *direct epistemic access* to phenomenal qualities is entailed by a metaphysical commitment to the existence of those qualities, and hence is shared by all of the metaphysical accounts which opt to

³¹ This should not be taken to contradict the view that phenomenality is a non-relational or intrinsic property of a subject, which was motivated in the previous chapter (§1.4.3). For even if a subject’s phenomenality is a non-relational or intrinsic property of theirs, they could nonetheless bear a relationship of ‘being aware of’ or ‘attending to’ to that non-relational, intrinsic property. In the same way, my mass is a non-relational, intrinsic property of me, but I can nonetheless know my mass, and therefore bear the relationship of ‘knowing’ to my mass.

locate phenomenal qualities within our ontology rather than to *eliminate* them. But although qualia realism, representationalism and naïve realism agree that we should locate phenomenal qualities within our ontology *somehow*, they disagree about *precisely how* we should locate them; although they agree that phenomenal qualities exist, they disagree about their specific metaphysical nature. And because they disagree about what phenomenal qualities essentially *are*, they disagree about *what it is that an experiencing subject has direct epistemic access to*. In other words, different accounts of what it is that phenomenal qualities metaphysically consist in will yield different commitments as to what it is that is directly epistemically accessible for the subject. Thus, the qualia realist is committed to the subject having direct epistemic access, via awareness and attention, to *the properties of experience itself*. Alternatively, the representationalist is committed to the subject having such access to *the properties that their mind represents*, whilst the naïve realist is committed to them having such access to *actual properties out in the world*. In contrast to all of these positions, the eliminativist is committed to the view that anything which the subject has access to, or is aware of or able to attend to, is not a phenomenal quality.

Having uncovered this epistemic dimension of the qualia debate, we can see that the qualia realist is committed to holding that experience presents—or provides one with awareness of, or enables one to attend to—properties of its own. The question of what experience presents is therefore of utmost relevance to the truth or falsity of qualia realism. For the notion of qualia is precisely the notion of

properties of experience itself which are presented in experience, and so qualia realism is true if and only if experience presents properties of its own.

2.4 The Conflation of Phenomenology with Metaphysics

We have seen that advocates and deniers of qualia each make claims—SO and ST, respectively—about what seems to be presented in experience. And we have seen why *what is presented* in experience is relevant to the truth or falsity of qualia realism. But we have yet to see how claims about what *seems* to be presented in experience feature in complete arguments for and against the existence of qualia. Below, I have reconstructed the arguments for and against qualia realism in which SO and ST respectively feature. I have aimed to capture the simplest valid arguments one can make for and against qualia realism whilst remaining faithful to the literature.

ST features as follows in the argument against qualia realism:

The Argument from Seeming Transparency (Against Qualia Realism)

- P1** Qualia exist if and only if experience presents properties of its own.
- P2/ST** Experience does not seem to present any properties of its own, but rather seems to exclusively present properties of things that the experience is of.

P3 If P2/ST, then experience *does not* present any properties of its own,
but only properties of things that the experience is of.

C Qualia do not exist; qualia realism is false.

This argument can be found in Tye’s ten-step appeal to transparency on behalf of representationalism (2000, pp. 46-51). The argument is spread out across all steps apart from 5 and 6, which are less concerned with rejecting a qualia-realist-type view and more concerned with promoting representationalism. Evidence of a commitment to P1 is peppered throughout, in Tye’s repeated denial of the existence of “qualit[ies] of your experience to which you have direct access” (p. 48) and of “qualities [of experiences] of which the subjects of ... experiences are directly aware...” (p. 49). Alongside this, step 1 asserts P2/ST, step 2 asserts P3 and C—though the assertion of P3 is somewhat obscure³²—and steps 3-4 and 7-10 reassert the consequent of P3, as well as C.

The same argument can be found in Tye (2014), but with P3 glossed-over. In this argument, P1 is explicitly stated in Tye’s definition of qualia (p. 41), and P2/ST is

³² It may not be immediately clear how Tye’s Step 2 asserts P3. The step states that “[t]o suppose that the qualities of which perceivers are directly aware in undergoing ordinary, everyday visual experiences are really qualities of the experience would be to convict such experiences of massive error. That is just not credible” (Tye, 2000, p. 46). In asserting that this supposition would convict experiences of massive error, Tye implies P2/ST: that experiences do not *seem* to us to present properties of their own (and hence would commit massive error if they *were* in fact presenting properties of their own). And in claiming that such massive error is not credible, Tye implies that experiences must be how they seem, rather than how they do not seem. Hence, he effectively asserts P3.

asserted in the paradigmatic excerpts from Moore, Harman and Tye, which Tye quotes at the beginning of his paper (pp. 39-40). Tye's commitment to P3, however, is brushed under the carpet. He covertly assumes the truth of P3 without visibly committing to it. Instead of explicitly stating the conditional, Tye simply affirms its consequent: that experience *does not* provide one with awareness of or enable one to attend to properties of experience itself.³³

SO plays a corresponding role in an argument of the same general form, but this time *for* qualia realism:

The Argument from Seeming Opacity (For Qualia Realism)

P1 Qualia exist if and only if experience presents properties of its own.

P2*/SO Experience seems to present properties of its own, not merely properties of things that the experience is of.

P3* If P2*/SO, then experience *does* present properties of its own, not merely properties of things that the experience is of.

C* Qualia exist; qualia realism is true.

This argument can be found, for example, in Kind (2008). P1 is assumed for the purpose of Kind's paper, firstly on the basis that one "might plausibly suppose that any qualia worthy of the name must be introspectible", and secondly on the basis that, even if qualia aren't *necessarily* introspectible, our ability to introspect upon

³³ See premises 10, 12, 15 and 16 in Tye (2014, pp. 42-43 & p. 45).

them constitutes the qualia realist's main reason for believing that they exist (p. 287), and as such might as well be considered a necessary and sufficient condition for their existence. The bulk of the paper is concerned with establishing P2*/SO by appealing to the examples of more *opaque*, 'exotic' experiences described in the previous subsection (pp. 288-293). As in Tye (2014), P3 is very much glossed-over; Kind simply ceases to write in terms of what it 'seems' we can attend to (2008, pp. 289-293) and begins to write just in terms of what we can attend to (pp. 293-296), hence assuming P3. Finally, it is concluded, on the basis that we can attend to phenomenal-quality-constituting properties of experience itself, that qualia realism is true (p. 296).

P3 and P3* here play a crucial role in building upon SO and ST to make valid arguments for and against qualia realism. These claims are integral, because they move each argument from a claim about how things *merely seem* to a claim about how things *actually are*. And yet they are glossed-over in the sense that neither Tye nor Kind explicitly express them in their entirety. Instead, the premise is either implied obscurely,³⁴ or bypassed completely in unscrupulous jumps straight from assertions of ST and SO to assertions of the consequents of P3 and P3*—from assertions that experiences *seem* a certain way, to assertions that experiences *are* that way, without any premise linking those assertions. The fact that these integral premises are obscured or bypassed should give us serious cause for concern.

Laura Gow (2016) exposes the extent to which philosophers have moved unscrupulously between claims about how things *seem* to claims about how things

³⁴ See fn. 32, above.

are specifically with regard to the transparency of experience. Gow shows (pp. 723-729) that philosophers have drifted carelessly between the claim that experience *seems* transparent (ST) and the claim that experience *is* transparent (the consequent of P3), to the point that they have effectively conflated the two claims.³⁵ Because the former claim is one about how things *seem to us*, Gow points out that it “should, in the first instance, be understood as a claim about the *phenomenology* of perceptual experience” (p. 725, original emphasis)—as a claim about what is consciously experienced from the first person perspective.³⁶ And because the latter claim is a claim about how experience *actually* is, it should be understood as a claim about the *metaphysics* of experience—a claim about the real nature of things. So according to Gow, arguments appealing to transparency have conflated *phenomenological transparency* with *metaphysical transparency*. We can here build upon Gow’s point, and add that arguments appealing to opacity have conflated *phenomenological*

³⁵ Although it does not affect the way I here draw upon Gow, it is worth noting that Gow and I employ different characterisations of transparency. Whereas I characterise transparency in terms of *experience only providing one with awareness of or enabling one to attend to properties of things that the experience is of*, Gow characterises it in terms of *all the properties we are aware of being externally located*. I think that my characterisation is preferable, given that whether the properties we are aware of are located internally or externally seems orthogonal to the issue of whether the experience is transparent or not. For example, one might cite real-time perception of surgery on one’s own brain as an example of a transparent experience in which one is aware of internally located properties. This seems like a conceptually coherent case in which transparency and externality come apart, which Gow’s characterisation cannot accommodate. But there does not appear to be any conceptually coherent example of a transparent experience which provides awareness of or enables one to attend to something more than the properties of things that the experience is of. Therefore, my characterisation of transparency more faithfully captures the notion.

³⁶ As David Woodruff Smith (2018, Introduction) states, “[p]henomenology is the study of structures of consciousness as experienced from the first-person point of view”.

opacity with *metaphysical opacity*. And so broadly, on both sides of the debate, we may consider there to have been a conflation between phenomenological claims and metaphysical claims, and hence a blurring of the boundary between phenomenology and metaphysics.

This conflation between phenomenology and metaphysics is problematic, because phenomenological claims do not simply translate into metaphysical claims, and do not straightforwardly entail corresponding metaphysical claims in the way that premises P3 and P3* assume that they do. Phenomenology is concerned with how things seem or appear to us.³⁷ And it might be tempting to think that phenomenology *does* correspond to metaphysics, given that how things seem to us is, in our everyday lives, a largely reliable guide to how things are. Most of the time, paths which seem to us to be clear of obstacles *are* clear of obstacles, mugs which seem to be full *are* full, and traffic lights which seem to have turned red *have* turned red. We therefore have an entirely justifiable default attitude of taking things to be as they seem, which we adopt unless there are good reasons for not doing so. Of course, we occasionally *do* have good reason for not taking things to be as they seem. If a magician appears to make a card suddenly cease to exist or teleport to an unseen location, the fact that this would flout the laws of physics gives us a good reason for doubting that this actually happened. But most of the time, we have no such reason to doubt that things are as they seem. This is what might motivate one to assume that phenomenology corresponds to metaphysics when it comes to the claims ST and

³⁷ As David Woodruff Smith (2018, §3) states “in its root meaning ... phenomenology is the study of *phenomena*: literally, appearances as opposed to reality.”

SO in the qualia debate. But to assume this would ultimately be misguided, because the very fact that we are engaging in a debate concerning *metaphysics* gives us good reason to abandon our default stance of simply taking things to be as they seem. For the modus operandi of metaphysics is precisely to look beyond mere appearances in an attempt to establish the real nature of things—to follow the escapee of Plato’s cave, coming to recognise the images on the wall as mere shadows of real objects.

In other words, when we engage in metaphysics, and ask questions about the real nature of things, one of our aims is precisely to probe beyond how things seem to us in our own experience, and hence to exceed the bounds of phenomenology. This is not only an aim of metaphysics, but of science as well. A block of metal seems to us to be made up of homogenous, space-occupying stuff, and yet the scientific community’s accepted atomic theory of matter tells us that most of the block is empty space. In exceeding the bounds of phenomenology, one must of course employ a methodology which does more than simply attend to how things seem or appear to us. Just as scientists do not rely purely upon how things seem to us, but instead hypothesise and experiment in order to establish a scientific picture of reality, philosophers engaged in metaphysics do not rely purely upon how things seem to us, but instead engage in reasoned argumentation in order to establish a clearer picture of the real nature of things. It is therefore misguided to rely heavily upon phenomenological claims about how things seem in order to establish a metaphysical conclusion, as advocates and deniers of qualia do with SO and ST in the debate over the metaphysics of phenomenal qualities.

In summary, although advocates and deniers of qualia attempt to argue for their respective views using claims about what seems to be presented in experience, these claims cannot establish or refute qualia realism unless combined with the assumption that the phenomenology of experience corresponds with its metaphysics, which is unwarranted in a debate concerning the metaphysics of phenomenal qualities. To avoid such unwarranted assumptions ourselves, we should ensure that our stance on phenomenal consciousness respects the crucial distinction between phenomenology and metaphysics. More specifically, we should take care not to base our metaphysical view of phenomenal consciousness solely upon phenomenological observations of how things seem to us in phenomenally conscious experience.

2.5 Conclusion

In this chapter, I have sought to highlight the importance of distinguishing between phenomenology and metaphysics, when we make claims about phenomenally conscious experience. I have done so by critiquing a particular strand of the debate over the existence of qualia, concerning the ‘transparency’ of experience. We saw that disputants purport to observe that experience either *is* or *is not* transparent, expressing these purported observations as claims about how things ‘seem’ or ‘appear’ in relation to experience. After having clarified how these claims are supposed to bear on the truth or falsity of qualia realism, we then saw that the arguments in which they feature erroneously conflate phenomenological claims with metaphysical claims. To avoid such errors ourselves, we should ensure that our claims

about phenomenal consciousness respect the distinction between phenomenology and metaphysics.

3. Avoiding Errors About What is 'Given' in Experience

In the previous chapter, we saw that there is a particular respect in which we *shouldn't* let phenomenology influence our thinking about phenomenal consciousness. We saw that we should take care not to base our metaphysical view of phenomenal consciousness *solely* upon phenomenological observations of how things seem to us in phenomenally conscious experience. In this chapter, I wish to turn my attention to the way in which we *should* let phenomenology influence our thinking about phenomenal consciousness, albeit whilst taking considerable precautions. I will first briefly explain why it makes sense to turn to phenomenology in order to establish what is 'given' to us in phenomenally conscious experience. Second, I will clarify precisely what I mean by the 'given'. Third, I will show that some philosophers have made mistaken phenomenological observations of what is given, due to a conflation between two distinct notions of 'seeming' or 'appearance'. Fourth, I will explain why we are so prone to error in what we take to be given to us in experience, and establish some rules to help us avoid making mistakes.

3.1 Phenomenology as a Way of Establishing What Is 'Given'

As we saw in the introduction, phenomenally conscious experiences feel like something to the subject who undergoes them, and can be *of* or *about* things. As such, a satisfactory conception of phenomenally conscious experience should have some idea of its capabilities and limitations in this regard. It should lay out what phenomenally conscious experience is capable of providing to its subject, both in terms of its *phenomenality*, or how it can feel, and in terms of its *intentionality*, or

what it can be *of* or *about*. In short, it should lay out what is 'given' to us in phenomenally conscious experience. And to establish what is 'given' in phenomenally conscious experience, it makes perfect sense to turn to phenomenology. For phenomenology is precisely the study of how things 'seem' or 'appear' to us in our own conscious experience, and a complete characterisation of how things seem or appear to us would encompass everything our experiences are of or about, and all aspects of what it is like for us to undergo those experiences.

Furthermore, although we shouldn't trust that how things phenomenologically seem or appear to us corresponds with how things metaphysically are, engaging in phenomenology is nonetheless a sensible first step towards establishing the real nature or metaphysics of phenomenal consciousness. For in establishing what is 'given' in phenomenally conscious experience, phenomenology establishes precisely what it is that a metaphysical account of phenomenal consciousness must explain, and can thereby guide us in developing a satisfactory metaphysical account.

There is, then, a legitimate role for phenomenology to play in influencing how we think about phenomenal consciousness. It can help us to establish what is 'given' to us in phenomenally conscious experience, and this in turn can guide our stance on the metaphysics of phenomenal consciousness. Even so, phenomenology is not an *infallible* means of establishing what is given. Before we see why, I must clarify precisely what I mean by what is 'given' to us in phenomenally conscious experience.

3.2 The 'Given' as What is Experientially Manifest

In defining the notion of what is 'given' in experience, I take my lead from Michelle Montague, in her 2016 book *The Given*. For Montague, everything that is given in an experience constitutes the 'content' of that experience. Although I will not employ the term 'content' in the same way, I will adopt Montague's definition of the content of an experience as my starting point for a definition of what is given in experience. According to Montague, the content of an experience is:

"everything one *experiences* in having the experience ... everything that is given to one, experientially, in the having of the experience, everything one is aware of, experientially, in the having of the experience."

(Montague, 2016, p. 30, original emphasis)

There is a lot going on in this definition, which it would be useful to tighten-up into something more clear and concise. I think this definition can be boiled down to three components: the intuitive idea of the given, plus two important qualifications. Intuitively, to think of something being 'given' in experience is to think of something being provided to the subject. As such, 'the given' intuitively refers to everything that a subject *gets* from their experience. Montague captures this intuitive idea by talking of "everything one *experiences*", "everything that is given to one", and "everything one is aware of".

Montague's definition regiments this intuitive idea with two important qualifications, which I wish to here make more explicit. First, the given is not simply *everything one experiences*, but only everything provided "in the having of the

experience.” This excludes from ‘the given’ certain things which a subject can colloquially be said to *experience*, but which do not plausibly feature amongst the things provided to a subject in the having of any particular experience. For example, a subject can be said to experience ‘a wonderful skiing holiday’ or ‘a stimulating fortnight of philosophy’. But it is implausible that such lengthy events are provided to the subject, in the having of an experience, in the same way that the presence of a mug of tea, the sounding of an alarm, or a certain shade of blue, are plausibly provided to the subject in the having of an experience. In other words, our colloquial classification of ‘things which can be experienced’ is much broader than the more technical classification with which we are here concerned, of things which are provided to the subject in the having of an experience.

The difference between the two classifications can be rendered more explicitly as follows. For something to be provided to the subject “in the having of an experience”, that thing must be, for the subject, a constituent feature of an episode of experience.³⁸ This accommodates our examples well, for it seems right to say that, if an experience provides to the subject the presence of a mug of tea, then the presence of a mug of tea must feature, for the subject, as a constituent part of an experiential episode. By contrast, it is implausible that ‘a wonderful skiing holiday’ or ‘a stimulating fortnight of philosophy’ could feature, for the subject, as a constituent

³⁸ Engaging fully with the technicalities of what constitutes an ‘episode’ of experience is beyond my present remit, but I maintain that such episodes are at least separated by interruptions in phenomenal consciousness, such as periods of unconscious sleep. (See §7.3.2.1—in particular fn. 126—on the temporal form of experience, for more precise considerations relating to this matter.)

part of a single experiential episode. Rather, to say that a subject 'experiences' such an event is to say that a subject undergoes a *collection* of experiences relating to the event, and not merely a *single* experiential episode in which the event features, for the subject, as a constituent part. So, we can summarise Montague's first qualification as follows: the 'given' is everything provided to the subject in the having of an experience, in the sense that it features, for the subject, as a constituent part of an experiential episode.

It is tempting, for the sake of simplicity, to consider the 'given' to be everything involved in the occurrence of a single episode of experience. But there are many factors plausibly 'involved' in the occurrence of a single episode of experience which are not provided to the subject, or do not themselves feature, for the subject, in their experience. In perceptually experiencing a table, for example, such factors might include one's brain activity, and the reflected light by which one sees the table. Such factors could play a significant role in the occurrence of an episode of experience, without ever actually featuring, for the subject, in their experience, and hence without ever being provided to the subject. Even if we tighten this classification, to only those factors involved in the occurrence of an experience which are 'given' to the subject, some of these may be considered to be 'given' to the subject in a sense different from the technical sense with which we are concerned. For example, an experience involving a certain pattern of brain activity might be said to 'give' that pattern of brain activity to the subject, in the sense that the subject's brain instantiates that pattern of activity when the experience occurs, even if the pattern of activity never features, for the subject, within their experience.

Montague's second qualification prevents such cases from being admitted into our technical category of 'the given'. Montague specifies that the given must be "*experienced*", or provided to the subject "*experientially*", and thereby excludes anything involved in the occurrence of an experience which is not actually provided to the subject, or does not actually feature, for them, in their experience.

In combination, these two qualifications distinguish our technical sense of 'the given' from yet another sense in which something may be 'given' in experience—the colloquial sense in which an experience might be said to 'give' one confidence, or 'give' one nightmares. In most cases, this departs from our technical sense of 'the given' because the experience referred to is in fact a collection of experiences, rather than a single experiential episode. For example, the 'experience' of an internship lasting a matter of months might be said to 'give' one confidence or nightmares, depending on its success. As such, whatever is supposedly 'given' is not something provided to the subject within a single experiential episode, and so is excluded by Montague's first qualification. But even if the experience referred to *is* a single experiential episode, such as that of seeing a tarantula, the nightmares that it 'gives' are a later event supposedly *caused by* that experience, rather than something provided to the subject contemporaneously in the having of the experience. Similarly, an experiential episode of telling a joke in front of an audience might 'give' someone confidence in the colloquial sense, by *causing* them to be more confident in future, even if they do not experience any confidence of their own during the telling of the joke. So, things 'given' in this sense are not necessarily provided to the subject *experientially*, in that they are not *experienced* contemporaneously, as part of the

experience which supposedly 'gives' them. As such, they are still excluded from our technical classification of 'the given', this time by Montague's second qualification.

The intuitive idea of the 'given', along with Montague's two qualifications, can be captured succinctly as:

Everything which is experientially manifest to the subject.

As such, 'the given' does not include everything which one might colloquially be said to experience, does not include everything involved in the occurrence of a single episode of experience, and does not include everything which experience might be said to 'give' a subject in the more colloquial, causal sense. Rather, it only includes everything which manifests itself to the subject as a constituent feature of one or another of the experiential episodes which they undergo. With this more refined definition of 'the given' in hand, I will now show how some philosophers have made inaccurate phenomenological observations of what is given to us in phenomenally conscious experience.

3.3 Mistakes Concerning What Is Given, arising from the Conflation of Two Kinds of 'Seeming' or 'Appearance'

In the previous chapter (§2.4), we saw that our view of the metaphysics of phenomenal consciousness should not be based purely upon how things seem or appear to us, from our own first-person-perspective. So, we should not assume that experience simply *is* how it *seems* to be, according to claims such as ST and SO. But

we have also established that phenomenology is a legitimate means of establishing what is 'given' to us in phenomenally conscious experience. So, a natural question to ask, at this point, is whether ST or SO qualify as accurate phenomenological observations. In other words: do either ST or SO report how things genuinely seem to us in phenomenally conscious experience, and thereby capture something that is genuinely given to us in phenomenally conscious experience?

So far, I have treated ST and SO as phenomenological claims, because they make claims about how things seem, as opposed to metaphysical claims about how things actually are. But it is now important to recognise that both ST and SO contain, nested within their phenomenological claims about how things seem, a metaphysical claim about how things are. In other words, they are both claims about how things *seem to actually be*. This is precisely why it is so easy for disputants such as Tye and Kind to move erroneously from such claims about how things *seem* to claims about how things *actually are*; one need only drop the qualification that things merely *seem* this way, and one is left with the claim that things *actually are* this way. So, ST and SO each assert that a certain metaphysical state-of-affairs, regarding what we are aware of or able attend to in experience, *seems* to obtain. We have seen that it is a mistake to trust that things *are* as ST and SO claim that they *seem* to be, but the question I am now concerned with is whether either of the seemings in question are *experientially manifest*, and hence given to us in phenomenally conscious experience.

It is evident that some of those who uphold ST and SO genuinely take the seemings captured in those claims to be manifest to us in our experience. For they claim that the metaphysical transparency and opacity of experience, respectively, are

evident simply from reflecting upon experience. Take, for example, Tye's retort to the following point made by Ned Block:

"...it is not easy (perhaps not possible) to answer on the basis of introspection alone the highly theoretical question of whether in so doing I am aware of intrinsic properties of my experience."

(Block, 1990, p. 73)

"It seems pretty easy to me. Besides, what's so highly theoretical here?"

(Tye, 2014, p. 49)

Consider also the following bold statement from Kind:

"Why believe in qualia? Because our every experience reveals their existence ... [philosophers who advocate ST] deny the phenomenological data."

(Kind, 2008, p. 285)

Ironically, Tye and Kind's own published work is testament to the absurdity of the views they express here. For if it were the case that either seeming transparency or seeming opacity were manifest to us in experience, Tye and Kind would not need to expend so much effort contriving examples to show that experience seems either transparent or opaque. Rather, expending such effort would be entirely unnecessary, much like trying to convince those with colour-vision that things seem to be coloured, or trying to persuade professional tasters that things seem to have tastes. "We already know that," they might say, "it's manifest to us in our experience." Evidently,

this is not the case when it comes to the seeming transparency or seeming opacity of experience. ST and SO are not simply given to us in experience.

Bernard Molyneux (2009) has different reasons for thinking that ST and SO are not experientially manifest. Molyneux argues that experience is effectively *silent* on the matter of its own metaphysics, by illustrating that we could make changes to the metaphysics of perceptual experience which would not affect what is given in experience, and would not affect the accuracy of the experience. For this he uses the example of Bob, who is fitted with tiny televisual contact lenses (pp. 123-127). These lenses have miniscule cameras on their outer surface, which feed live images to a display on their inner surface, such that they “perform a useless, ‘total-of-nothing’ operation on the incoming light signals” (p. 123). There is a strong sense here in which the metaphysics of the perceptual experience have changed in comparison with Bob’s ordinary perceptual experience, since his visual experience is now presenting properties of the televisual lens displays, which it otherwise wouldn’t have been.³⁹ However, since the lenses perform a ‘total-of-nothing’ operation on the incoming light signals, no aspect of this metaphysical change is manifest to Bob in his experience. Furthermore, the fact that the new metaphysical situation is not experientially manifest does not render the experience *inaccurate*, for the ‘total-of-nothing’ operation of the lenses ensures that Bob’s perceptual experience remains

³⁹ Molyneux (2009, p. 123) claims that “Bob’s experience of the world, in the lens example, is indirect.” Although I think that the metaphysics of perceptual experience, in the lens example, differs from the metaphysics of ordinary perceptual experience, I do not think that this is correctly characterised as a change in the directness of the perceptual experience.

sensitive to the world in exactly the same way that it is sensitive to the world in ordinary *accurate* perception. And given that Bob's *accurate* perceptual experiences are completely insensitive to whether or not he is wearing televisual lenses, his perceptual experience cannot be judged to be *inaccurate* for being insensitive to the presence of the lenses when they are there. It therefore seems that experience is entirely neutral on the metaphysical matter of whether or not televisual lenses mediate Bob's perception of the world. And this makes it plausible that no state-of-affairs whatsoever, concerning the metaphysics of experience, is manifest to us in our experience. In other words, it is plausible that experience is completely silent on matters of its own metaphysics, including the matter of whether it is transparent or opaque.

Christopher Frey (2013, p. 73) also maintains that experience is silent about its own metaphysics, on the basis that phenomenological reflection alone does not fully determine the metaphysical status of the "sensuous elements" which "one appreciates in experience",⁴⁰ nor the metaphysical status of "experiential episodes themselves." For, as Frey explains, concerning the sensuous elements of experience,

⁴⁰ Judging by the examples he gives, Frey's "sensuous elements" are equivalent to what I am calling 'phenomenal qualities': "the way the redness of a Red Delicious apple looks when one sees it, the way middle C sounds when one hears it being played on a Bösendorfer piano, and the way a pain feels when one experiences a pin pierce one's finger" (Frey, 2013, fn. 2).

“...even if we always appreciate an experience’s sensuous elements as being or qualifying physical objects in a publicly accessible environment, it may turn out that they are, as a matter of fact, instantiated by amalgams of sense-data or are intrinsic features of one’s experiential state.”

(Frey, 2013, p. 73)

In other words, even if a phenomenal quality such as blueness always *seems* to be a property of the surfaces of objects out in the world, it may nonetheless be the case that blueness is *actually* an intrinsic property of one’s own experience. For as Gow (2016) points out, as we saw in §2.4, *phenomenological transparency* does not entail *metaphysical transparency*. How things phenomenologically *seem* does not fix how they metaphysically *are*. Similarly, with the experiential episodes themselves, a diverse array of candidate metaphysical accounts remain on the table even once one has engaged in phenomenological reflection:

“Perhaps one has an experience with a particular phenomenal character in virtue of being in a representational state with a special kind of content or a special functionally specified role. Perhaps one has this experience in virtue of standing in some primitive relation of acquaintance or direct awareness to an appropriate class of entities. Perhaps one has this experience by virtue of the divine dispensation of an omnipotent god. Whatever the source, it is invisible. The means by which a scene becomes experientially present to one is not itself phenomenally appreciable.”

(Frey, 2013, p. 73)

The fact that one’s phenomenological reflections are compatible with such a diverse array of metaphysical accounts certainly suggests that experience does not disclose its own metaphysical nature, but is instead neutral or silent on the matter.

Not only is the silence of experience plausible, but we also lack any positive reason for thinking that experience is, so to speak, vocal on matters relating to its own metaphysics. Gary Hatfield (2007, p.149) points out that “there is no reason to suppose that the senses have evolved in order to reveal the principles of their own operation”, and the same may be said for phenomenally conscious experience—there is no reason to suppose that it works in such a way that it must have something to say regarding its own metaphysics.⁴¹

Thus, ST and SO are not simply given to us in experience, and Kind and Tye are grossly mistaken in purporting to phenomenologically observe them amongst what is given, as part of the ‘phenomenological data’. So, how is it that Kind and Tye have made this error? I believe that the error is due to a conflation between two distinct kinds of ‘seeming’ or ‘appearance’. Only one of these is genuinely given to us in experience, but because both are expressed in terms of how things ‘seem’ or ‘appear’, we are liable to mistakenly treat both as phenomenologically observed aspects of how things ‘seem’ or ‘appear’ to us in experience, and hence as aspects of what is given to us in experience. So, if ST and SO exemplify the sort of seeming *not*

⁴¹ It is worth noting that experience’s silence or neutrality with regard to its own metaphysics does not preclude phenomenological reflection from being a sensible first step towards establishing the metaphysics of phenomenal consciousness, as I argued in §3.1. For in establishing what is ‘given’ in phenomenally conscious experience, phenomenology can establish what it is that a metaphysical account of phenomenal consciousness must explain. So although it cannot determine precisely which metaphysical account is correct, it does specify what criteria a satisfactory metaphysical account must meet, if it is to explain everything that is ‘given’ in phenomenally conscious experience. Therefore, phenomenology still has some role to play in guiding us towards a satisfactory metaphysical account of phenomenal consciousness.

given to us in experience, what examples can we give of the sort of seeming which *is* given?

Examples of seemings which patently *are* given to us in experience can be found in paradigmatic examples of what is supposed to constitute phenomenal consciousness. Consider, for example, the experience of seeing the blueness of a clear summer sky, or feeling the warmth of bathwater. These can be equally well characterised as instances of the sky ‘appearing’ blue, or the bathwater ‘seeming’ warm. And it seems that all experiences involving phenomenal qualities can be characterised as instances of something ‘seeming’ or ‘appearing’ to have those qualities. In fact, it seems that phenomenal qualities can be faithfully characterised as *ways* something can ‘seem’ or ‘appear’ to us; something can ‘seem’ or ‘appear’ blue, warm, sweet, loud, round, smooth, etc.⁴² It almost goes without saying that such seemings are experientially manifest, since it is not clear what the notion of a

⁴² One might contend that not all phenomenal qualities can be faithfully characterised in this way. For example, it seems that bodily sensations such as aches and pains should be counted as phenomenal qualities. However, to say that your own broken ankle ‘seems’ or ‘appears’ painful to you, or that your upset stomach ‘seems’ or ‘appears’ achy to you, would seem to mischaracterise the phenomenon. It seems rather more appropriate to say that your broken ankle simply *is* painful, and that your stomach simply *is* achy. But although our way of characterising these particular phenomenal qualities does not translate smoothly into a characterisation in terms of *ways things* (i.e. the broken ankle or the stomach) seem, they are not strong counterexamples to the view that all phenomenal qualities are themselves ‘seemings’ or ‘appearances’. For we can nonetheless recognise that, when we consciously experience a broken ankle or an upset stomach, the pain or the ache ‘appears’ to us or ‘seems’ to us to be there, in the same way that blueness and warmth ‘appear’ to us or ‘seem’ to be there when we consciously experience a clear summer’s day.

phenomenal quality could refer to, if not to a way in which something can manifestly seem or appear to us in experience.

Now that we have some examples of the two kinds of seeming at issue, we can begin to elucidate their differences. It is worth heeding Block, in the excerpt above, when he points out that the question concerning experience's transparency or opacity is a "highly theoretical" one. For ST and SO are indeed highly theoretical, in the sense that they each assert that a certain metaphysical state-of-affairs seems to obtain. ST effectively asserts that it *seems* that qualia realism is false, and SO effectively asserts that it *seems* that qualia realism is true. By contrast, phenomenal qualities are far less theory-laden. When something seems or appears, 'blue', 'square' or 'cold' to us, for example, what is manifest to us is not some metaphysical state-of-affairs, but rather the phenomenal qualities themselves, of blueness, squareness or coldness. In fact, it seems ludicrous that any metaphysical state-of-affairs could ever manifest itself within our experience in this way. For the truth or falsity of qualia realism would need to appear before us similarly to blueness or coldness, as something like qualia-realism-is-true-ness or qualia-realism-is-false-ness. But this grossly mischaracterises what occurs when a metaphysical state-of-affairs such as qualia realism seems to one to be true or false. Rather than being something that is simply 'given' to us in experience, the seeming truth of a metaphysical state-of-affairs is more appropriately characterised as an *inclination to believe* that the relevant state-of-affairs obtains. And although one might insist that some aspect of this inclination is manifest to us in experience, such as the experience of a 'eureka moment' when one suddenly grasps the solution to some problem, this

is not equivalent to the manifestation, in experience, of the very metaphysical state-of-affairs which supposedly seems to be true.

Let us call experientially manifest seemings 'e-seemings'. An archetypal e-seeming is the visual greenness involved in looking at a lime. In this case, something *seems green* just because the phenomenal quality of *greenness* is manifest to the subject, in their experience. And let us call non-experientially manifest seemings, or inclinations to believe, 'i-seemings'. An archetypal i-seeming is the inclination to believe that qualia realism is true. In this case, qualia *seem* to exist in the sense that one is inclined to believe in a metaphysical state-of-affairs in which qualia exist, but that metaphysical state-of-affairs would not be manifest to the subject, in their experience.⁴³

⁴³ Andrew Cullison (2010, §II, pp. 264-66) argues that 'seemings' are not equivalent to 'inclinations to believe', as I am here taking i-seemings to be. However, I think that all three of the reasons he gives for his view can be overturned.

First (pp. 264-65), Cullison presents a counterexample, wherein one's inclination to believe comes apart from what seems to be true. Sam is inclined to believe in his wife's new found religious beliefs because he is worried she might divorce him otherwise, even though what his wife believes seems false to him. We can overcome this counterexample simply by adding the qualification that seemings are equivalent only to a *specific sort* of inclination to believe, namely the inclination one has to believe in something when it coheres with or follows from what one already believes or understands to be true.

Second (p. 265), Cullison observes that something's seeming to be true is sometimes part of the explanation as to why we are inclined to believe it. Therefore, he claims, the seeming is explanatorily prior to the inclination to believe, and cannot itself consist in the inclination to believe. But we can deflect this objection as follows. We can hold that, when one explains that they have an inclination to believe in something *because* it seems to be true, they are merely qualifying that their inclination is of the specific sort wherein one is inclined to believe something which coheres with or follows from what one already believes or understands to be true. As such, we may talk about the seeming-to-be-true *as if* it is explanatorily prior to the inclination to believe, when in fact the former simply qualifies the latter. Qualifications presented in the form of explanations are common, as in the case of

This distinction, between e-seemings and i-seemings, somewhat proves its worth in helping to explain what is actually going on in the dispute concerning transparency and opacity. In asserting ST, Tye is really expressing that, on the basis of e-seemings such as the squareness and blueness involved in seeing a blue square, it i-seems to him that experience is transparent, and that qualia realism is false. And in asserting SO, Kind is really expressing that, on the basis of e-seemings such as the blurriness of blurred vision, it i-seems to her that experience is opaque, and that qualia realism is true. In other words, the dispute is over what metaphysical account of experience we should be inclined to accept on the basis of the e-seemings given to us in experience. And although Kind clearly takes examples of more 'exotic' e-seemings to be more effective in inclining one towards opacity and qualia realism, and Tye clearly takes more 'mundane' e-seemings to be more effective in inclining one towards transparency and away from qualia realism, their efforts to

sentences such as 'the road is bumpy because it is potholed', or 'the paint is wet because it hasn't yet dried'.

Third (pp. 265-66), Cullison presents another counterexample, wherein something comes to seem to be true even though one has always been inclined to believe it. He claims that, as a result of his father having certain beliefs about government, and his inclination to follow his father, he grew up with a strong inclination to believe the same. Nonetheless, as an adult, Cullison claims that what he believes now seems to him to be true in a way it did not when he was a child. Yet again, we can overcome this counterexample by appeal to our qualified, more specific sort of inclination to believe. Even if Cullison's childhood inclination to believe was of the correct sort, in that he was inclined to believe what his father believed because it cohered with or followed from his existing belief or understanding that his father's testimony was reliable, it is plausible that the life experience he has since acquired has furnished him with a wider range of beliefs and a deeper understanding of what is true, which the beliefs in question now more tightly cohere with, or follow from with even greater certainty, such that what he believes now seems to be true in a way which it did not when he was a child.

accommodate each other's examples into their respective metaphysical accounts show that they do not in fact dispute one another's accounts of the e-seemings given to us in experience,⁴⁴ and so do not disagree over the 'phenomenological data'.⁴⁵ However, because we tend to express both e-seemings and i-seemings in terms of how things 'seem' or 'appear' to us, the two can easily become conflated. As a result, Kind and Tye mistakenly treat their i-seemings as e-seemings, and as such make erroneous, conflicting claims about what is given to us in experience.

To sum up, we have seen how mistakes are easily made about what is given to us in experience. It makes sense to engage in phenomenology in order to establish what is given in experience, precisely because phenomenology is the study of how things 'seem' or 'appear' to us in experience. However, there are two distinct kinds of 'seeming' or 'appearance'—i-seemings and e-seemings—only the latter of which is actually given to us in experience. As such, when we engage in phenomenology, and consider how things seem or appear to us in experience, we are liable to fail to discriminate between these two kinds of seeming, and mistakenly count *both* as being given to us in experience.

⁴⁴ For example, Kind (2008, pp. 295-296) makes efforts to show that qualia realism is compatible with the examples of "mundane" experiences which proponents of transparency appeal to, and Tye (2000) dedicates Chapter 4 to explaining why examples of the sort Kind and Block raise, which he calls "oddities", are in fact compatible with experience consisting wholly of representational content, and hence presenting represented properties only.

⁴⁵ Frey (2013, p. 75) also recognises that there is in fact a great deal of agreement over the 'phenomenological data', amongst competing metaphysical accounts. As he puts it: "Competing philosophical accounts of perception—for example, representationalism, naïve realism, projectivism and sense-datum theories—issue identical phenomenological 'predictions.'"

The lesson here, then, is that we must exercise caution about what we take to be given to us in experience. More specifically, we must ensure that we are not conflating the e-seemings which are given to us in experience with the i-seemings which are not. But how precisely are we to guard against this error? And are there any other errors we are prone to, in what we take to be given to us in experience? In order to address these questions thoroughly, we will need to return to the distinction drawn between metaphysics and phenomenology in §2.4, and reassess it in light of our distinction between i-seemings and e-seemings. This will be the concern of the following section.

3.4 Distinguishing Between Appearance and Reality in Both Phenomenology and Metaphysics

3.4.1 Questions to Address

We have seen that phenomenology, as the study of how things ‘seem’ or ‘appear’ to us in our own conscious experience, is prone to error about what is given to us in experience. This is because the notion of how things ‘seem’ or ‘appear’ to us is ambiguous between two different kinds of ‘seeming’—between e-seemings which *are* given to us in experience, and i-seemings which are *not*.

However, we have so far characterised each type of seeming using a single archetype. We characterised e-seemings using the archetype of a phenomenal quality, which is patently manifest in our experience, and we characterised i-seemings using the archetype of an inclination to believe in some metaphysical state-of-affairs, in which case that state-of-affairs is patently not manifest in our

experience. But there is plausibly a significant grey-zone between these two archetypes, of seemings whose nature, as either an e-seeming or an i-seeming, is not immediately clear. Therefore, if we are to avoid errors in what we classify as e-seemings, and thereby take to be given in experience, we require a more comprehensive policy regarding what we admit into the category of e-seemings, and what we exclude as an i-seeming.

But we may be getting ahead of ourselves, for one may still harbour doubts that there is any need for concern when it comes to establishing what is given in experience. For even if one accepts that we must be somewhat careful not to conflate these two kinds of seeming, one may think that the likelihood of doing so is relatively low, and that there is therefore little need for caution. "After all," one might say, "what is given to us in experience is precisely what is experientially manifest to us, and if it is experientially manifest to us, then we can't go far wrong in characterising what is given to us in experience, because what is given will be right there before us!" In other words, it looks like there is little room for error in characterising what is given in experience, precisely because what is given is simply given to us. As M.G.F. Martin (1998) points out, this view isn't uncommon in philosophical circles:

"A common assumption in most philosophical discussions of appearances and experience is that, when one does engage in just such reflection, the character of how things appear to one is just obvious to me."

(M.G.F. Martin, 1998, p. 157)

There is, of course, a sense in which this assumption is entirely warranted. For the given is what is *experientially manifest* to us, and it is not at all clear what this could mean, if not that the given is in some way directly accessible to us, directly available to us, right in front of us, obvious to us, or there for us to see. But if this is so, then it seems that the potential for error regarding what is given in experience should be minimal, if not negligible. Therefore, if I am to make a sufficiently convincing case that we in fact have a significant proclivity for error regarding what is given, I need to square this with the fact that the given is experientially manifest, and thereby directly there for us to see.

We therefore have two questions to address:

1. How can it be that what is 'given' is what is experientially manifest to us, and yet we are so liable to make mistakes about it?
2. What should our policy be, regarding what we admit into the category of e-seemings, and what we exclude as an i-seeming?

To answer the first question, and begin to answer the second, we need to reassess our distinction between phenomenology and metaphysics. In the previous chapter (§2.4), the distinction between phenomenology and metaphysics was presented as the distinction between how things *seem* and how they *actually are*—as the distinction between *appearance* and *reality*. We characterised phenomenology as the domain of how things *seem* or *appear* in first-person, conscious experience, and metaphysics as the domain of how things *really are*. In other words, we characterised

the distinction between *phenomenology and metaphysics* as being in complete alignment with the distinction between *appearance and reality*. However, as we will now see, our distinction between i-seemings and e-seemings reveals the erroneous simplicity of this characterisation. The appearance-reality distinction does not in fact align with the phenomenology-metaphysics distinction. Rather, the distinctions are orthogonal, such that we can distinguish between appearance and reality within the domain of phenomenology, as well as within the domain of metaphysics.

Recognising this more complex state of affairs will enable us to answer the first question which needs addressing. It will enable us to see why we are so liable to make mistakes about what is experientially manifest, in two ways. First, it will reveal a more substantive reason as to *why* one would be inclined to conflate i-seemings with e-seemings, besides the fact that they are both widely conceived of under the same banner of 'seeming' or 'appearance'. Second, it will enable us to see that there is a *further* mistake which obscures the given, and which occurs at the intersection of i-seemings and e-seemings, but which is not due to conflation of the two.

Recognising the complexities of the situation will also help us begin to answer the second question which needs addressing. For it will involve unpacking yet further ambiguities in seeming and appearance talk, thereby providing us with a more fine-grained taxonomy of seeming and appearance claims. This more fine-grained taxonomy of seeming and appearance claims will not itself establish a definitive policy on what we admit into the category of e-seemings, but it will enable us to articulate two rules which will help us to avoid making errors about what is given.

So, let us begin by seeing how our distinction between i-seemings and e-seemings demands the dealignment of the phenomenology-metaphysics distinction and the appearance-reality distinction.

3.4.2 Dealigning the Phenomenology-Metaphysics and Appearance-Reality Distinctions

The very notion of e-seemings, of ‘seemings’ or ‘appearances’ which are experientially manifest to us, demands a distinction between appearance and reality within the domain of phenomenology. For e-seemings constitute phenomenological *reality*. They are what is *genuinely* given to us in first-person, phenomenally conscious experience. They are what our first-person, phenomenally conscious experience is *actually* comprised of. So, if e-seemings are phenomenological reality, what constitutes phenomenological appearance?

Well, for any given aspect of reality, one might try to observe that reality. But observation is not an infallible means of access to reality. There is the possibility that what one observes to be the case is not in fact the case. Nonetheless, whatever one sincerely observes to be the case is what *seems* or *appears* to one to be the case, regardless of whether the observation turns out to be accurate. An observation reports an *appearance*, regardless of whether it reports the *reality*. A phenomenological *appearance*, then, is an observation of what is given in one’s own experience, which may or may not accurately capture the *reality* of what is given in experience. In other words, a phenomenological appearance is how things *seem* to e-seem, which may or not accurately capture how things *actually* e-seem.

Patently, the sort of *seemings* which constitute phenomenological appearances cannot themselves be e-seemings. As established in §3.3, e-seemings are not theory-laden. When something e-seems 'blue', 'square' or 'cold' to us, what is manifest to us is just the blueness, squareness or coldness, and not the obtaining of some metaphysical state-of-affairs. To put this more precisely, e-seemings are not *highly conceptual*. That is, although e-seeming blueness might be taken to involve a concept of blueness, and perhaps also a more general concept of colour, it does not involve anything like the sophisticated array of concepts required for it to i-seem to one that qualia realism is true. When a metaphysical state-of-affairs i-seems to be true, the i-seeming is 'highly theoretical' in the sense that it is highly conceptual, requiring the subject to possess a sophisticated array of advanced concepts. By contrast, e-seemings are *not* highly conceptual, requiring a minimal number of rudimentary concepts, if any at all. As such, an observation as to how things e-seem cannot itself be an e-seeming. The very act of stepping back from e-seemings, to observe how things e-seem, takes one away from the business of simply experiencing things such as 'blueness', 'squareness' and 'coldness', and puts one in the business of making highly conceptual observations of states-of-affairs. In other words, it takes one away from simply *experiencing* what is experientially manifest, and puts one in the business of *conceptualising* what is experientially manifest.

If things *seeming* to e-seem some way is too conceptual to be an e-seeming, then it looks like the seeming in question could be an i-seeming. This makes a great deal of sense. For we know that the seeming in question is an observation concerning the state-of-affairs of what is experientially manifest to one, and it would be fair to

count a sincere observation of some state-of-affairs as an *inclination to believe* in that state-of-affairs. As such, things *seeming* to e-seem some way is a case of things *i-seeming* to e-seem some way. Phenomenological appearances are, therefore, how things i-seem to e-seem—they are *inclinations to believe* that something or other is ‘given’ in one’s experience, or is experientially manifest to one. In sum, phenomenological *reality* is how things e-seem to one, and phenomenological *appearance* is how things i-seem to e-seem to one. This is the appearance-reality distinction within phenomenology.

The demand for an appearance-reality distinction within metaphysics is much easier to illustrate. Our archetypal i-seemings were inclinations to believe in a certain metaphysical state-of-affairs, such as qualia realism seeming true or seeming false. As ways in which the metaphysical state-of-affairs *seems* to one, such i-seemings constitute metaphysical *appearance*. In this case, what constitutes metaphysical *reality* is not difficult to discern. Metaphysical reality is precisely the metaphysical state-of-affairs which *actually does* obtain. So, the appearance-reality distinction within metaphysics is simply the distinction between how things i-seem to be, and how things actually are.

The appearance-reality distinction is so much easier to spell out for metaphysics than for phenomenology precisely because metaphysical reality is not itself a kind of appearance. Rather, metaphysical reality is rock-bottom, fundamental reality—it is the ultimate true nature of things. The appearance-reality distinction within phenomenology is so much more complex, and thereby conducive to errors concerning what is really given in phenomenally conscious experience, precisely

because phenomenally conscious experience is *itself* a sort of appearance or seeming—it is comprised of those experientially manifest e-seemings which are genuinely given to us in experience.

As we will shortly see, errors occur when one recognises that phenomenally conscious experience is itself a sort of appearance or seeming, but fails to recognise that the appearance-reality distinction can be drawn within both metaphysics and phenomenology. First, it is worth acknowledging that there is a strong precedent for taking phenomenally conscious experience to be a sort of appearance or seeming.

3.4.3 Recognising Phenomenally Conscious Experience as a Sort of Appearance or Seeming

The recognition of phenomenally conscious experience as itself a sort of appearance or seeming is widespread and longstanding. It can be found in Thomas Reid's *Essays on the Intellectual Powers of Man*:

“When I smell a rose ... the agreeable odour I feel, considered by itself, without relation to any external object, is merely a sensation. ... This sensation can be nothing else than it is felt to be. Its very essence consists in its being felt; and when it is not felt, it is not. There is no difference between the sensation and the feeling of it; they are one and the same thing.”⁴⁶

(Reid, 1785, pp. 226-27)

⁴⁶ Demian Whiting (2016, p. 47) instead uses the following excerpt as an illustration of Reid's view: “When I am pained, I cannot say that the pain I feel is one thing, and that my feeling it is another thing. They are one and the same thing, and cannot be disjoined even in imagination. Pain, when it is not felt, has no existence.” Although this excerpt more cogently expresses the view, it is not taken from Reid's original 1785 edition of *The Intellectual Powers of Man*, but appears only in James Walker's abridged 1850 edition (p. 133), and subsequent editions.

The phenomenally conscious experience in question here is the sensation of smelling the agreeable odour of the rose. And Reid recognises that this experience consists precisely in the sensation's 'being felt'. That is, the experience consists in the odour being *felt*, being *apparent* to one, or *seeming* to be there. If there is no feeling of an odour, no appearance of an odour, and no seeming occurrence of an odour, then there is no phenomenally conscious experience of an odour, because such an experience consists precisely in the feeling, appearance or seeming occurrence of an odour. The phenomenally conscious experience *is itself* a sort of seeming or appearance.

More recently, this has been recognised by Saul Kripke (1980), Thomas Nagel (1974) and John Searle (1992). All three articulate the point by highlighting an asymmetry between the phenomenon of conscious experience, and other natural phenomena such as heat, colour and sound. Whilst Nagel and Searle both characterise the asymmetry in terms of reductive explanation, Kripke frames it in terms of the way our natural kind terms work.

Nagel (1974, pp. 444-45) and Searle (1992, pp. 118-24) point out that experience cannot be reductively explained using the same strategy which works for most natural phenomena. The strategy is as follows. For some phenomenon (e.g. heat, the colour red, loud sounds), first set aside the appearances that it produces for us (e.g. hotness, redness, loudness), and then redefine or redescribe it without any reference to the appearance it produces, in terms of scientifically respectable entities or occurrences (e.g. molecular motion, certain spectral reflectance

distributions, high amplitude sound waves).⁴⁷ In sum, the strategy is to set aside the *appearance* of the phenomenon and prioritise its *reality*. But as Nagel and Searle point out, we cannot employ the same strategy to reductively explain conscious experience, because conscious experience *consists in* the very appearances—of hotness, redness, and loudness—which would ordinarily be set aside. As such, there is no ‘appearance’ we can set aside without thereby setting aside the phenomenon of conscious experience itself. In the words of Nagel and Searle themselves,

“Experience itself ... does not seem to fit the pattern. The idea of moving from appearance to reality makes no sense here.”

(Nagel, 1974, pp. 444-45)

“...we can’t make that sort of appearance-reality distinction for consciousness because consciousness consists in the appearances themselves. *Where appearance is concerned we cannot make the appearance-reality distinction because the appearance is the reality.*”

(Searle, 1992, pp. 121-22, original emphasis)

Kripke (1980, pp. 144-55) highlights the same asymmetry, but frames it in terms of the way natural kind terms function. For Kripke, the asymmetry is not that phenomena such as heat can be *redefined* or *redescribed* in a way which conscious experience cannot be. Rather, it is that terms such as ‘heat’ relate to conscious experience differently to the way in which terms such as ‘pain’ do.

⁴⁷ In Nagel’s own terms, “...the process of reduction is a move in the direction of greater objectivity ... accomplished by reducing our dependence on individual or species-specific points of view toward the object of investigation. We describe it not in terms of the impressions it makes on our senses, but in terms of its more general effects and of properties detectable by means other than the human senses” (1974, p. 444).

On Kripke's view, the term 'heat' needn't be redefined or redescribed, because it already rigidly designates 'molecular motion'. In other words, that phenomenon which we call 'heat' is necessarily identical to molecular motion. By contrast, it is entirely contingent that the phenomenon of heat produces in us the appearance of hotness. For sentient beings like us might never have existed, or might have been wired such that the phenomenon of heat produced the appearance of blueness or sweetness. Nonetheless, according to Kripke, the appearance of hotness fixes molecular motion as the referent of the term 'heat'. For it is only because molecular motion causes the appearances of hotness that we are sensitive to its occurrence, and thereby able to ascribe the name 'heat' to it. In sum, 'heat' rigidly designates molecular motion, and the appearance of hotness simply fixes molecular motion as the referent of 'heat'.

Kripke points out that there is an asymmetry in the way that the terms 'pain' and 'heat' work, because the appearance of painfulness doesn't merely fix the referent of the term 'pain' in the same way that the appearance of hotness does for 'heat'. Rather, painfulness *constitutes* pain. The appearance of painfulness is precisely what the term 'pain' rigidly designates. As such, where there is an element of contingency between the phenomenon of heat and the appearance of hotness, there is no analogous contingency between the phenomenon of pain and the appearance of painfulness.⁴⁸ There is no appearance which pain contingently produces, precisely because it already *is* an appearance. In Kripke's own words:

⁴⁸ Kripke (1980) makes this point as part of an argument against brain-identity theory. Having established that "identity is not a relation which can hold contingently

“...in the case of mental phenomena there is no ‘appearance’ beyond the mental phenomenon itself.”

(Kripke, 1980, p. 154)

Although Kripke expresses the point in different terms, he evidently recognises what Reid, Nagel and Searle also recognise: that phenomenally conscious experiences, such as the experience of pain, are themselves a sort of *appearance* or *seeming*.

Far from being repudiated or dismissed as trivial, the point continues to be emphasised in contemporary philosophy. For example, Terry Horgan reiterates the point in his introduction to a recent paper concerning introspection on phenomenal consciousness:

“...in the case of phenomenal consciousness there is no gap between appearance and reality, because the appearance just is the reality...”

(Horgan, 2012, p. 406)

In sum, there is a strong and continuing precedent for recognising phenomenally conscious experience to be a sort of *appearance* or *seeming*.

between objects” (p. 154), Kripke tries to show that there is an apparent element of contingency between the mental state of ‘pain’ and the supposedly identical physical state of ‘C-fibre stimulation’ which cannot be explained away. Any apparent contingency between ‘heat’ and ‘molecular motion’, he maintains, can be explained away in terms of the contingency between the phenomenon of heat and the appearance of hotness, leaving the necessary identity of ‘heat’ and ‘molecular motion’ intact (pp. 149-51). But, Kripke argues, the apparent contingency between ‘pain’ and ‘C-fibre stimulation’ cannot be explained away in the same way, precisely because pain doesn’t contingently cause an appearance of painfulness in the same way that heat contingently produces the appearance of hotness (pp. 151-54). Rather, pain *necessarily consists in* the appearance of painfulness. As such, this avenue for explaining away the apparent element of contingency between ‘pain’ and ‘C-fibre stimulation’ is not available to identity theorists.

3.4.4 Errors due to Mistaken Alignment of the Phenomenology-Metaphysics and Appearance-Reality Distinctions

Let us now see why we are so prone to error about what is given, even though it is experientially manifest to us. There are two errors, both of which occur when one correctly recognises experience as itself a sort of seeming or appearance, but mistakenly aligns the phenomenology-metaphysics distinction with the appearance-reality distinction.

We have just seen that there is a strong precedent for recognising experience as a sort of seeming or appearance, but one may doubt that there is any significant tendency to align the phenomenology-metaphysics distinction with the appearance-reality distinction, and hence doubt that this can make any significant contribution to errors concerning what is given in experience. The fact that there *is* such a tendency, and a significant one at that, is evident in some of the articles cited in §2.4, in support of my own characterisation of the phenomenology-metaphysics distinction in terms of the appearance-reality distinction. Smith (2018, §3), for example, writes that “in its root meaning ... phenomenology is the study of phenomena: literally, appearances as opposed to reality.” Also, Gow (2016) explicitly presents “phenomenological transparency” as a claim about how things “seem”, and “metaphysical transparency” as a claim about how things “are in fact”, thereby aligning phenomenology with appearance and metaphysics with reality.

There is, therefore, a genuine tendency to mistakenly align the phenomenology-metaphysics distinction with the appearance-reality distinction. When this mistaken alignment is paired with the recognition that experience itself is

a sort of seeming or appearance, it results in two different errors. The first error is the one we recognised in §3.3, of conflating e-seemings that *are* given in experience with i-seemings that *aren't* given. This error occurs as follows.

If one mistakenly aligns appearance with phenomenology, then one will consign all appearances or seemings to the domain of phenomenology. As such, when things *appear* or *seem* to be a certain way, metaphysically, one will consider that appearance or seeming to be a phenomenological observation as to how things seem or appear to us, in our first-person conscious experience. This approach, paired with the view that conscious experience itself consists in how things seem or appear to us, results in one taking the way things seem to be, metaphysically, to be a constitutive part of conscious experience. In other words, it results in one taking metaphysical states-of-affairs to be given to us in experience.

In this way, the mistaken alignment of appearance with phenomenology inclines one towards the error of conflating e-seemings that *are* given in experience with i-seemings that *aren't* given. By aligning appearance entirely with phenomenology, one neglects the fact that there can be seemings or appearances which are not phenomenological observations, such as i-seemings concerning metaphysical states of affairs. Furthermore, one takes the fact that such i-seemings are a sort of seeming or appearance as indication that they belong within the domain of phenomenology. And given that no appearance-reality distinction is recognised *within* the domain of phenomenology, all seemings or appearances consigned to that domain are taken to constitute the very seeming or appearance which conscious experience itself consists in. One thereby ends up taking metaphysical appearances

to constitute a part of phenomenological reality. In other words, i-seemings about metaphysics are conflated with the e-seemings which comprise phenomenally conscious experience.

We hereby see a more substantive reason as to *why* one would be inclined to conflate i-seemings with e-seemings, besides the fact that they are both widely conceived of under the same banner of 'seeming' or 'appearance'. We are liable to conflate them due to a tendency to mistakenly align the phenomenology-metaphysics distinction with the appearance-reality distinction, combined with the recognition that conscious experience is itself a sort of seeming or appearance. Alignment of the distinctions consigns *all* seemings or appearances to the phenomenological domain, and all seemings consigned to that domain are taken to constitute conscious experience, or be given in experience. As such, one ends up with an overinflated view of what is given in experience.

The second error we are prone to is that of treating introspection as an infallible guide to what is given in conscious experience. This error arises as follows. If one mistakenly aligns the appearance-reality and phenomenology-metaphysics distinctions, then there is no distinction available to be drawn *within* phenomenology, between phenomenological appearance and phenomenological reality. As such, there is no way of distinguishing between observations of how things e-seem from e-seemings themselves. That is, how things i-seem to e-seem is not distinguished from how things e-seem. This is not to say that the i-seemings about how things e-seem are conflated with the e-seemings themselves, but rather that the gap between how things i-seem to e-seem, and how things genuinely e-seem, is

overlooked. In overlooking this gap, one neglects the potential for an i-seeming about how things e-seem to be mistaken about how things genuinely e-seem. One neglects, for example, the possibility that they are inclined to believe that something e-seems green to them, despite nothing actually e-seeming green to them. In other words, one neglects the possibility that one's observations of how things e-seem might fail to accurately capture how things actually e-seem. One thereby commits the error of taking introspection to be an infallible guide to what is given in conscious experience.

If one mistakenly aligns the appearance-reality and phenomenology-metaphysics distinctions, then the recognition that conscious experience is itself a sort of appearance or seeming is likely to entrench the error of treating introspection as an infallible guide to what is given in experience. For if conscious experience is itself a sort of seeming or appearance, such that, as Horgan (2012, p. 406) claims, "the appearance just is the reality", then it looks as if one needn't worry about how things seem or appear to one being misleading in any way. In other words, if there is no gap or distinction between appearance and reality when it comes to conscious experience, then it looks as if one should be able to take one's introspective observations of one's own conscious experience at face value. It looks as if how one's own conscious experience seems or appears to one should be the final word on how it really is, with no further room for doubt. This line of reasoning is captured well by Keith Frankish:

"It is incoherent to doubt that experiences are as they seem, since experience reports are already reports of how things seem. I may come to doubt my initial claim that there is a green patch in front of me and retreat to the more cautious claim that there seems to be a

green patch, but I cannot coherently retreat from that claim to the claim that there seems to seem to be a green patch. The first claim expresses all the epistemic caution that is necessary or possible.”

(Frankish, 2016, p. 33)

This line of reasoning is, however, mistaken. One *can* coherently retreat from the claim that there seems to be a green patch to the claim that there seems to seem to be a green patch. Our distinction between i-seemings and e-seemings reveals how this is so. One can coherently retreat from the claim that there e-seems to be a green patch, to the claim that there i-seems to e-seem to be a green patch. That is, one can coherently retreat from the claim that a green patch is given in one’s experience, to the claim that one is inclined to believe that a green patch is given in one’s experience. As such, the claim that there seems to be a green patch evidently *does not* express ‘all the epistemic caution that is necessary or possible’. Not only is it possible to exercise caution over whether or not there actually *is* a green patch, it is also possible to exercise caution over whether or not there actually *seems* to be a green patch. And not only is this further caution possible, it also seems necessary.

As mentioned earlier (§3.4.2), the act of stepping back from e-seemings, to observe how things e-seem, takes one away from the business of simply experiencing things such as ‘blueness’, ‘squareness’ and ‘coldness’, and puts one in the business of making highly conceptual observations of states-of-affairs. It takes one away from simply *experiencing* what is experientially manifest, and puts one in the business of *conceptualising* what is experientially manifest. And it is this process of *conceptualising* what is experientially manifest—of capturing or characterising what

is experientially manifest using the concepts in one's possession—which opens up significant room for error about what is given in experience. For we have no reason to suppose that we are naturally apt at accurately characterising what is given to us in experience. Therefore, we cannot simply trust our introspective observations to deliver infallible reports about what is really given in conscious experience. Rather, we must be aware that our introspective observations might be inaccurate.

So, in sum, mistakenly aligning the appearance-reality and phenomenology-metaphysics distinctions will result in one overlooking the distinction between appearance and reality *within* the domain of phenomenology. One will thereby fail to recognise that what *appears* to be given in experience (how things i-seem to e-seem) might not accurately capture the *reality* of what is given in experience (how things e-seem). And if one fails to recognise this potential for error, the claim that “in the case of phenomenal consciousness ... the appearance just is the reality” (Horgan, 2012, p. 406) only serves to entrench one's negligence.

There is, of course, one specific sense in which it is true that, in the case of phenomenally conscious experience, the appearance just is the reality. It is true in the sense that phenomenally conscious experience is comprised of e-seemings, or made up of everything which is experientially manifest to one in undergoing phenomenally conscious experience. But one is only able to interpret the claim in this nuanced, qualified way if one already recognises that the appearance-reality and phenomenology-metaphysics distinctions are orthogonal rather than aligned. If one fails to recognise this, then one is likely to interpret the claim that “the appearance just is the reality” as the claim that there is *no* distinction or gap *whatsoever* between

appearance and reality when it comes to phenomenally conscious experience, and this will only entrench one's negligence to the potential error of introspective observations.⁴⁹

In conclusion, we are prone to two errors which obscure what is given to us, or what is experientially manifest, in experience. Both errors occur due to a tendency to mistakenly align the phenomenology-metaphysics distinction with the appearance-reality distinction, combined with the recognition that conscious experience is itself a sort of seeming or appearance. The first error we are prone to is that of conflating i-seemings which are *not* given to us in experience with e-seemings which *are* given. This can result in i-seemings about metaphysical states-of-affairs being counted amongst what is given in experience, thereby overinflating the category of what is given in experience. The second error we are prone to is that of treating introspection as an infallible guide to what is given in experience. This can result in erroneous introspective observations of what is given being unreflectively

⁴⁹ If one is in any doubt that the claim is likely to be interpreted in this way, one need only look to Horgan (2012) to see how close he gets to claiming that introspection is infallible: "...in the case of phenomenal consciousness there is no gap between appearance and reality, because the appearance just is the reality: *how the phenomenal character seems, to the agent, is how it is*" (p. 406, emphasis added). Furthermore, Searle (1995) clearly makes efforts to cancel the implication that introspection is infallible, thereby acknowledging that the implication is there: "If it seems to me exactly as if I am having conscious experiences, then I am having conscious experiences. This is not an epistemic point. I might make various sorts of mistakes about my experiences, for example if I suffered from phantom limb pains. But whether reliably reported or not, the experience of feeling the pain is identical with the pain in a way that the experience of seeing a sunset is not identical with a sunset."

accepted as the truth, thereby giving one a distorted impression of what is given in experience.

3.5 Conclusion

Let us return to the questions posed earlier, and review the answers we now have.

1. How can it be that what is 'given' is what is experientially manifest to us, and yet we are so liable to make mistakes about it?

Everything which is 'given' or experientially manifest to us is indeed directly available to us, or obvious to us, in the sense that it consists in those *seemings* or *appearances* we are privy to when things seem or appear 'blue', 'loud' or 'cold', for example. In other words, what is given is precisely the 'blueness', 'loudness' or 'coldness' which one experiences when things seem or appear 'blue', 'loud' or 'cold' to one. As such, everything which is experientially manifest is, by its nature, a sort of seeming or appearance. To be privy to such seemings or appearances is precisely what it is to undergo a phenomenally conscious experience. It is in this sense that phenomenally conscious experience is itself a sort of seeming or appearance.

However, it is crucial to recognise that these seemings or appearances which are experientially manifest to us, and which phenomenally conscious experience consists in, are of a specific sort. Namely, they are what I have been calling e-seemings. We are liable to make mistakes about what is given in experience precisely because we are prone to two sorts of error when it comes to discerning e-seemings.

First, we are liable to conflate i-seemings which are *not* given to us in experience with e-seemings which *are* given. For example, we might mistakenly take metaphysical states-of-affairs to be given to us in experience. Second, we are inclined to treat introspection as an infallible guide to what is given in experience, when it can in fact deliver inaccurate observations. Therefore, we might blindly accept erroneous introspective observations as the truth, and end up with a distorted impression of what is given in experience.

In sum, what is 'given' or experientially manifest to us is directly available or obvious to us in precisely the way that e-seeming 'blueness' is directly available or obvious to us when it occurs. Nonetheless, we are liable to make mistakes about what is 'given' precisely because we are prone to two sorts of error when it comes to discerning e-seemings.

2. What should our policy be, regarding what we admit into the category of e-seemings, and what we exclude as an i-seeming?

In order to understand our proclivity for error regarding what is given, we first had to recognise that the distinction between phenomenology and metaphysics does not align with the distinction between appearance and reality, but is in fact orthogonal to it. Recognising this involved unpacking further ambiguities in seeming and appearance talk. Further to the distinction between i-seemings and e-seemings, we distinguished between two sorts of i-seeming. First, on top of the e-seemings which constitute phenomenological reality, there are the i-seemings we have *about* e-seemings when we make introspective observations about e-seemings, and these

i-seemings constitute phenomenological appearance. Second, there are the i-seemings we have about metaphysical states-of-affairs, and these i-seemings constitute metaphysical appearance. This more fine-grained taxonomy of seeming and appearance claims does not itself establish a definitive policy on what we admit into the category of e-seemings, but it does enable us to articulate two rules which will help us to avoid the two aforementioned errors concerning what is given:

1. Not all appearances or seemings are what is 'given' or experientially manifest in phenomenally conscious experience. I-seemings are distinct from e-seemings, and only e-seemings are experientially manifest.
2. Our introspective observations of what is 'given' are not infallible. An introspective observation is how things i-seem to e-seem, and this may not accurately capture how things really e-seem.

With these two rules in mind, we can now proceed to consider what phenomenal conditions might be constitutive of the non-relational phenomenal intentionality of conscious experience.

PART II

4. What is the Phenomenological Signature of Directedness?

'Phenomenal intentionality' is the notion of a kind of intentionality which is exhaustively constituted by phenomenality or phenomenal character. As Farid Masrour points out,

"This idea implies that the conditions that a state needs to satisfy in order to count as a state with phenomenal intentionality are phenomenal conditions. One task for the phenomenal intentionalist is thus to provide an account of the phenomenal conditions that constitute phenomenal intentionality. We can call this the constitution problem."

(Masrour, 2013, p. 117)

In other words, if one wishes to uphold that there is phenomenal intentionality—a form of intentionality which is constituted by phenomenality alone—then one should be able to overcome Masrour's "constitution problem" by explaining what phenomenal conditions are constitutive of phenomenal intentionality. Uriah Kriegel refers to these conditions as the "phenomenological signature of intentionality" (2011, p. 156) or "the phenomenological signature of directedness" (2013, p. 17). I will adopt the latter label, hereafter abbreviating it to PSD.

Across two publications (2011 & 2013), Kriegel surveys six different suggestions as to what the PSD consists in, but does not embark on a critical evaluation of them. In this chapter, I will do just that. I will argue that the most promising candidate for the PSD is the *structural organisation* of our actual phenomenal experiences, because it is less problematic and has greater explanatory

power than the other candidates. I will first establish some essential desiderata for an account of the PSD, and then consider the candidates one-by-one.

4.1 Desiderata for an Account of the Phenomenological Signature of Directedness

In order to be taken seriously, an account of the PSD must:

- A. Specify only experientially manifest phenomenal conditions.
- B. Be informative, or reduce the mysteriousness of phenomenal intentionality.

Desideratum A is essential for two reasons. First, in requiring an account of the PSD to specify only *phenomenal* conditions, it limits the scope of our inquiry to the phenomenal domain. Phenomenal intentionality is supposed to be a kind of intentionality which is constituted only and entirely by phenomenal character, where the phenomenal character of an experience is what-it-is-like for the subject to undergo that experience. Therefore, the conditions necessary and sufficient for phenomenal intentionality must be specifiable solely in terms of what-it-is-like to undergo an experience instantiating phenomenal intentionality, from the subject's own first-person perspective. One could not include, for example, the condition that the experience must be realised by a certain brain state, because although this specifies a way the experience must be, it does not specify what the experience must be like *for the subject to undergo, from their own first-person perspective*. Second, in

requiring an account of the PSD to specify only *experientially manifest* phenomenal conditions, Desideratum A focuses our inquiry on phenomenological *reality* as opposed to mere phenomenological *appearance*. In other words, it prompts us to consider whether a candidate set of phenomenal conditions are genuinely given in experience, or whether we are merely inclined to believe, falsely, that they are given. In the terms of the previous chapter, it prompts us to consider whether the phenomenal conditions are genuine e-seemings, or merely i-seeming e-seemings. In this chapter, I will not delve too deeply into these particular considerations. I will consider the desideratum sufficiently fulfilled so long as it is ostensibly plausible that the candidate phenomenal conditions are genuinely experientially manifest, and there is no obvious reason for suspecting that we are making one of the mistakes identified in Part I. The candidate we identify as the most promising will be put through more rigorous scrutiny, in this regard, in the next chapter.

Desideratum B is necessary for ensuring that an account of the PSD is sufficiently motivated by its own informativeness. Patently, anything worth calling an ‘account’ of something must offer some insight into the deeper nature of that thing, rather than leaving it equally as mysterious as it was prior to the account being articulated. For example, Theresa May’s pronouncement that “Brexit means Brexit” does not qualify as an informative account of what ‘Brexit’ means, for it leaves the meaning of ‘Brexit’ equally as mysterious as it ever was. Of course, informativeness admits of degree, so it is also worth bearing the desideratum in mind to distinguish more informative accounts from less informative ones.

4.2 Phenomenality Itself

Horgan and Tienson (2002) argue that the phenomenal character and intentional content of conscious intentional states are so deeply interconnected that they are best regarded as inseparable aspects of such states. In their own pithy terms, their claim is that “phenomenal consciousness is intentional through and through” (p. 530). Chalmers (2004, p. 179) also expresses this view, stating that “intentional content appears to be part and parcel of phenomenology: it is part of the essential nature of phenomenology that it is directed outward at a world.” According to this view, then, the PSD simply consists in phenomenality itself, for any conscious intentional state possesses its intentionality simply by virtue of having some phenomenal character or other.

On this view, Desideratum A is trivially fulfilled—we need not worry whether the conditions said to constitute the PSD all fall within the experientially manifest phenomenal domain, because the proposed condition of *being phenomenal* is, by definition, fulfilled exclusively within the experientially manifest phenomenal domain. If the view holds up under examination, it also has the potential to be informative, and hence to fulfil Desideratum B. The extent to which it is informative to recognise phenomenality as being necessary and sufficient for phenomenal intentionality depends upon what more there is, or isn't, to say about phenomenality itself. The nature of phenomenality itself is, of course, a stubborn mystery. If one takes phenomenality to supervene on the physical, then one faces the so-called ‘hard problem of consciousness’, and is tasked with explaining why certain physical states are accompanied by certain phenomenally conscious states, and why they are

accompanied by phenomenal consciousness at all. On the other hand, if one takes phenomenality to *not* supervene on the physical, one faces the equally formidable task of explaining how this can be so. The suggestion is at least minimally informative because it indicates that *two* mysteries—the nature of phenomenal intentionality, and the nature of phenomenal consciousness—are in fact *one*, just as it would be informative to realise that wherever the car keys are will be wherever I left my coat, given that my car keys were in my coat pocket. However, just as this realisation won't help me locate my keys unless I can find my coat, the recognition that phenomenality itself is the PSD won't be any more informative unless the nature of phenomenality is more successfully spelled out. So, in sum, phenomenality itself is a minimally informative candidate for the PSD, and has the *potential* to be more informative, depending upon what more can or cannot be said about the nature of phenomenality. This is enough to warrant taking phenomenality itself into consideration as a candidate for the PSD.

Let us now consider the plausibility of the view that any conscious intentional state possesses its intentionality simply by having some phenomenal character or other. The natural critical response is to try to produce a counterexample: an example of a phenomenally conscious state which is not intentional. And the most promising examples are precisely those qualitative experiences which paradigmatically set phenomenally conscious states, such as conscious perceiving or imagining, apart from merely intentional ones, such as unconscious believing or desiring; experiences of, for example, redness, warmth, sourness and pain. Horgan and Tienson appear to anticipate this line of objection, erecting two lines of pre-

emptive defence. Horgan and Tienson make their case using the example of an experience of redness, but their points are intended to apply to any qualitative experience. I will continue in the same vein.

As a first line of defence, they make clear that the primary target of their argument is not qualitative experiences of redness *per se*, but “experiences of red *as we actually have them*” (2002, p. 521, emphasis added). Rather than experiencing pure redness in isolation, Horgan and Tienson claim that red is always seen first as a property of objects located in space, and is experienced as part of a complete three-dimensional scene. Furthermore, they claim that the red is experienced as unified in space with other properties of the red object, such as the sweetness, heaviness, firmness and smoothness of a red apple. In sum, they claim that our actual experiences of red do not occur in isolation, but rather as parts of broader experience-complexes which, together as a whole, present us with a three-dimensional world populated with objects. These are intended to be reasonable, uncontroversial claims which draw attention to the ways in which qualitative experiences play a role in depicting or representing things to us, and hence play a role in determining the intentional content of our conscious states.

However, even if it is the case that experiences of red *as we actually have them* are always parts of experience-complexes which carry intentional content, it doesn't follow that we *couldn't* have experiences of red which are not a part of such experience-complexes, and which do not carry any intentional content. In other words, it doesn't follow that non-intentional experiences of red are *impossible*. The role that redness experiences play in determining intentional content might be

entirely contingent, rather than necessary. Therefore, Horgan and Tienson's claims about our *actual* experiences do not rule out the *possibility* of phenomenally conscious states which do not carry intentional content. And if such states are possible, then phenomenal consciousness is not intentional through and through, and the PSD cannot be phenomenality itself.

We can therefore push through Horgan and Tienson's first line of defence by insisting on the possibility of those experiences which they imply to be absent from our actual lives, and which are the most plausible examples of non-intentional phenomenal consciousness: namely, experiences of a single phenomenal quality, such as redness, in complete isolation. Indeed, a strong case can be made for the possibility of such experiences. Christopher Frey (2013) highlights two ways in which it is eminently plausible that such experiences might occur, given our existing empirical knowledge.

"Spatially Punctiliar Experiences" (p. 75) are one way. Frey points out that, given that damage to one's occipital cortex can result in a 'shrinking' of the visual field, it is in principle possible for one's occipital cortex to be damaged in such a way that one "would undergo visual experiences in which a single phenomenal point, *and nothing more*, is appreciable" (p. 76, original emphasis).⁵⁰ Frey reasonably maintains that such experiences, consisting only of a single phenomenal point, would not be appreciated as being spatially divisible in the way an ordinary visual field might be

⁵⁰ Frey (2013, pp. 75-76) clarifies that the area of the visual field lost "is not replaced with darkness." As such, the experiences which are in principle possible are "not experiences as of a point of light in a sea of darkness," but rather are experiences which are exhaustively constituted by a single phenomenal point.

considered to be, and would therefore not present the subject with any apparent spatial extension (Ibid., fn. 14). Nonetheless, the single phenomenal point must consist in some visual phenomenal quality or other, such as that of redness. Spatially punctiliar experiences thus provide an example of a single phenomenal quality experienced in isolation, rather than as part of an experience-complex. Furthermore, their possibility is eminently plausible, given that it is the upshot of reasonable inferences from empirical facts concerning the effect of brain damage on the visual field.

Ganzfeld experiences are another way in which an experience of a single phenomenal quality may plausibly occur in isolation. As Frey states, “a ganzfeld is a visual field that is completely permeated by a constant and homogeneous sensuous colour” (p. 75). The possibility of such experiences is eminently plausible, not only because it is easy to conceive of an experience of a particular shade of colour expanding so as to fill one’s entire visual field, but also because people have actually reported having such experiences. In these experiences, the constancy of a homogeneous shade means that no kind of spatially extended object or scene is presented, sometimes to the extent that the expanse of colour is not even experienced as a surface located in front of the subject.⁵¹ Given that people actually report having such experiences, ganzfelds don’t only strengthen the case for the possibility of non-intentional experiences of a single phenomenal quality in isolation, but also counter Horgan and Tienson’s implication that such experiences are absent

⁵¹ For an investigation which explores the breakdown of ordinary spatial experience in ganzfelds, see Cohen (1957).

from our actual lives. So ganzfelds also undermine Horgan and Tienson's first line of defence from a different direction: by providing evidence that Horgan and Tienson's claims about actual experience are false.

We therefore have a strong case against Horgan and Tienson's first line of defence. Frey's two examples provide us with a strong case for the possibility of non-intentional experiences of a single phenomenal quality in isolation, and ganzfelds in particular provide evidence that Horgan and Tienson are simply relying upon a false claim about our actual experiences. But now we hit up against the second line of defence:

“...even considered in isolation from any total visual-experiential state, the what-it-is-like of experiencing red is already intentional, because it involves red *as the intentional object of one's experience*. Again, redness is not experienced as an introspectible property of one's own experiential state, but rather as a property of visually presented objects.”

(Horgan and Tienson, 2002, p. 521, original emphasis)

In the first sentence, Horgan and Tienson effectively claim that, even if an experience of red in isolation were possible, it would not be devoid of intentionality because it would involve red as the intentional object of one's experience. The second sentence then spells out what it means for red to be an intentional object of experience: namely that it is experienced transparently—not as a property of one's own experience, but as a property of a visually presented object.

However, this conception of what it means to be an intentional object appears to undermine Horgan and Tienson's own claim to be considering an

experience of red *in isolation* from any total visual-experiential state. For if red is to be experienced *as a property of a visually presented object*, then presumably that experience must feature a visually presented object. But although it is clear that an experience-complex, as described above, could feature a 'visually presented object' by virtue of multiple phenomenal qualities being experienced as unified in space, it is not at all clear that an isolated experience of red could achieve the same. An isolated experience of red appears incapable of featuring a 'visually presented object', precisely because it lacks the experiential resources of the experience-complex. So insofar as Horgan and Tienson are insisting that an isolated experience of red involves red as the intentional object of one's experience, where this requires the experience to feature a visually presented object, they are failing to recognise the full implications of the experience being a completely isolated one.

One can, however, spell out what it means for red to be an intentional object of one's experience in a way which does not commit one to the experience featuring a visually presented object. One could say that red is the intentional object in the very basic sense that it would *represent* or *be about* the property of being red. And one can avoid postulating 'visually presented objects' by employing the conception of intentionality presented by Siewert (1998, pp. 12 & pp. 188-194), wherein a feature is intentional if it is a feature in virtue of which its possessor is assessable for accuracy. Applying this to experience, an experience is intentional if the subject of that experience is assessable for accuracy just in virtue of having that experience. In other words, an experience is intentional if it entails a set of conditions which must be fulfilled for the subject of the experience to be experiencing accurately. This

conception provides us with a clearer sense in which an isolated experience of pure redness might be considered to *represent* or *be about* the property of being red; it might entail that the subject of the experience is only experiencing accurately if the property of being red is actually instantiated.

However, we are presently concerned with the question of whether or not phenomenal consciousness is itself *essentially* intentional. On the conception just outlined, an experience is essentially intentional if the subject is assessable for accuracy *just by virtue of having the experience*. But it does not appear that a subject *is* assessable for accuracy *solely* by virtue of having an isolated experience of pure redness. Rather, it seems that the experience only entails accuracy conditions if it is assumed to have the function of representing or indicating the instantiation of the property of being red.

We take our own redness experiences to have this function because of the way in which they contribute to our successes and failures; when redness experiences correspond to the instantiation of the property of being red, and hence to the presence of red objects, we are able to successfully navigate the world and avoid eating poisonous red berries. But when redness experiences do not so correspond, we might be liable to bump into red objects, or accidentally eat poisonous red berries. Given that we succeed when our own redness experiences correspond to the presence of red objects, it makes sense to treat those experiences as having the function of representing or indicating the presence of red objects, or the instantiation of the property of being red.

But independently of these considerations, there is no reason to assign redness experiences in general any function to indicate or represent. And without being taken to have such a function, redness experiences in and of themselves do not entail accuracy conditions, and are therefore not *essentially* intentional. This becomes apparent if we consider an instance of an isolated experience of pure redness, and strip it of all context which would incline us to assign some representational function or other. Let us imagine a conscious being whose entire life experience consists in a momentary, isolated experience of pure redness, either as a ganzfeld or spatially punctiliar experience. Consider this momentary episode of consciousness in abstraction from any physical substrate, and hence in abstraction from any phylogenetic or ontogenetic history, or any mechanical role in a physical system. Without any of these avenues to plunder for a hint of the representational function the experience might serve, the experience itself is not assessable for accuracy or inaccuracy; it is simply a phenomenally conscious state of a particular qualitative nature, which in and of itself does not represent or indicate anything at all.⁵²

Horgan and Tienson could, of course, insist that such experiences are *not possible*. It is certainly true that I have not *proven* that phenomenally conscious states devoid of intentionality are possible; I have merely invited the reader to conceive of

⁵² One might be tempted to respond that it is the experience's potential to serve a representational function which makes it intentional, but given that literally anything harbours the potential to represent something or other, this position would make intentionality indefensibly ubiquitous. (For example, my coffee cup has the potential to represent mars, for I can, if I choose, use it as a prop standing-in for mars when demonstrating the arrangement of the solar system.)

such a state, and this is by no means a good guide to genuine possibility. But it is also true that Horgan and Tienson have not *proven* that that non-intentional phenomenal states are *not* possible. Therefore, although we can't completely rule out the possibility that the PSD is phenomenality itself, we can conclude that phenomenality itself isn't a particularly convincing candidate for the PSD, for it is not at all clear that phenomenality entails intentionality.

4.3 Sui Generis Phenomenal Directedness

A different candidate for the PSD is proposed by Brian Loar (2003/2017).⁵³ Loar argues that, in our own experiences, we can recognise a phenomenal feature of 'directedness' which is alone responsible for furnishing our experiences with phenomenal intentionality.⁵⁴ This 'directedness' is not explained in terms of any simpler or more basic phenomenal features, and in this sense is taken to be *sui generis*: a unique phenomenal feature in its own right. Let us see whether it fulfils our desiderata.

Loar (2003/2017, pp. 303-304) makes a plausible case for directedness being an experientially manifest phenomenal feature of our experiences. He invites us to

⁵³ Note that, although Loar's 'Phenomenal Intentionality as the Basis of Mental Content' was originally published in *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (2003), I will here use the pagination from a more recently published collection of Loar's essays, *Consciousness and Meaning: Selected Essays* (2017).

⁵⁴ Loar (2003/2017) allows that a mental state may be about something by successfully referring to it, and maintains that this would require the state to bear a relationship to some actual thing in the external world. However, when it comes to the 'aboutness' that a mental state possesses just by virtue of its phenomenal character, the unique phenomenal feature of 'directedness' is solely responsible.

imagine several indistinguishable lemons being one-after-the-other brought to our visual attention, and to imagine thinking something about each lemon, such as ‘that’s yellow’. We are then told that some of the apparent lemons were hallucinations, and asked whether our experiences and thoughts all visually presented their objects in the same way. Loar claims that the natural answer is yes. I agree that this is the intuitive answer, for we are supposed to have undergone several experiences *as of* lemons which are indistinguishable from one another, and in each case had a thought about the lemon-as-experienced. The way in which a lemon-as-experienced is visually presented appears to remain the same even when there is in fact no actual lemon, in which case our experience is not veridical and our thought fails to refer to an actual lemon.

Loar then argues that this identified sameness—in the way in which the thoughts and experiences visually present an object—must be sameness in a feature of the thoughts and experiences which is *intentional*, in the sense that it brings *aboutness* or *directedness* to the thought or experience. He reasons that all of the thoughts and experiences “purport to refer” to a lemon (Ibid.), regardless of whether they actually succeed in doing so or not. He also rejects two ways of trying to accommodate the sameness in something other than an intentional feature.

First, one cannot sufficiently accommodate the sameness in some intentionality-stripped presentation-of-an-object in general, or presentation of

'some object'.⁵⁵ Loar's objection seems to be that this does not sufficiently capture the fact that our thoughts and experiences, in the imagined scenario, are all specifically *as of* lemons with the same particular appearance—a specificity which suggests *aboutness* or *directedness* at play here, as opposed to a tokening of some general concept of 'objecthood'. One could also object to this move on grounds of incoherence, given that it is reasonable to suppose that any presentation-of-an-object in general, or concept of general 'objecthood', could never be entirely devoid of intentionality.

Second, it is implausible that the sameness of our thoughts (e.g. 'that's yellow') consists in the sameness of some existentially quantified representational content, such as 'I am seeing some lemon or other and it is yellow'.⁵⁶ This is a tempting option for the opponent of Loar, because it looks promising as a way of explaining the sameness between the veridical-perception-based thoughts and the hallucination-based thoughts purely in terms of representational content, without needing to appeal to some phenomenal-intentional feature irreducible to representational content, as Loar insists we must. However, this option is implausible for two reasons. First, it seems much more reasonable to understand the thoughts as being *demonstrative* rather than *existentially quantified*: pointing to whatever

⁵⁵ Loar (2003/2017, p. 302): "You cannot capture this common feature by generalizing over objects: 'there is some object that the demonstrative concept visually presents'."

⁵⁶ Loar (2003/2017, p. 302, original emphasis): "And surely the content of those thoughts is not itself existentially quantified: 'I am seeing some object and it is yellow.' The thoughts in question are *demonstrative* and they are not self-consciously reflexive."

'that' is, rather than asserting the existence of a lemon and ascribing the colour yellow to it.⁵⁷ This point is drawn out by Loar in a slightly modified version of the lemon scenario, wherein one is having a lemon-experience without knowing at the time whether it is veridical or hallucinatory. Loar claims, quite reasonably, that one would find themselves torn between treating the lemon as existing (saying, for example, 'that object'), and withholding commitment to the existence of the lemon, resulting in the lemon-as-experienced being felt as "a sort of ontologically neutral object that could have the property of existing or not-existing" (2003/2017, p. 318). Such an experience patently could not be captured in terms of existentially quantified representational content asserting the existence of an actual yellow lemon.⁵⁸ Second, it seems more plausible that the thoughts are simple demonstrative ascriptions of the colour concept 'yellow', as opposed to the more complex self-reflexive states which would be required for one to ascribe to oneself an act of seeing.⁵⁹

⁵⁷ See previous footnote.

⁵⁸ To clarify, we are effectively here presented with a dilemma. On the one hand, existentially quantified representational content cannot accommodate the feeling of an ontologically neutral object, precisely because it makes an assertion about the existence of the lemon (which we would expect to find reflected in one's conscious experience, if representational content is to be at all plausible as an account of conscious experience.) On the other hand, if we remove the existential quantifier, and hence remove the assertion of the lemon's existence, then we have no explanation of one's (conflicted) inclination to treat the lemon as existing.

⁵⁹ The opponent of Loar might insist that, even if the thoughts are more reasonably understood as being simple demonstrative ascriptions rather than existentially quantified self-reflexive states, this does not prevent one from explaining their sameness in terms of representational content. For one could just posit representational content which is simple and demonstrative in nature, rather than existentially quantified and self-reflexive. However, this response overlooks the implications of Loar's lemon example. The example is intended to show that experiences and thoughts can 'purport to refer' to a lemon, regardless of whether or not they succeed in referring to an actual lemon. In other words, it identifies a

In sum, Loar makes a plausible case for there being some experientially manifest phenomenal feature of certain experiences, which he calls ‘directedness’, which is alone sufficient for furnishing mental states with a form of intentionality. Loar’s suggestion therefore fulfils Desideratum A. But in order to be an all-round good candidate for the PSD, ‘directedness’ must also fulfil Desideratum B, by reducing the mysteriousness of phenomenal intentionality.

Sui generis phenomenal directedness ultimately turns out *not* to be a good candidate for the PSD. This is because we face a dilemma. On the one hand, we could uphold the view that phenomenal directedness is genuinely *sui generis*. In order to be genuinely *sui generis*, directedness must be a *basic or ground-level* phenomenal feature, in the sense that it cannot be broken-down or analysed into simpler or more basic elements. On this view, to claim that directedness is the PSD is to explain phenomenal intentionality simply by identifying the unique and basic phenomenal

demonstrative element (a ‘pointing to’ the lemon-as-experienced) which obtains regardless of successful reference, and hence cannot be accounted for in terms of reference. As such, the challenge for Loar’s opponent is to explain this demonstrative element without appealing to reference. But the opponent of Loar would be extremely hard-pressed to flesh out a notion of representational content which is demonstrative in nature, but without appealing to reference. If the thought ‘that’s yellow’ is explicable purely in terms of representational content, then how could one explain the contribution of demonstrative ‘that’, other than by appealing to reference? It seems that the opponent’s best strategy is to avoid demonstratives altogether, and maintain that the ‘purporting to refer’ Loar identifies is not in fact demonstrative, but is explicable in terms of representational content which asserts the existence of a lemon and ascribes an act of seeing to the subject—representational content which is existentially quantified and self-reflexive. It is precisely because this is the best option available to his opponent that Loar chooses to point out its implausibility in comparison to his own account, which posits a phenomenal-intentional feature which is demonstrative (comprising a *purporting-to-refer*) without referring.

feature which comprises it. On this horn of the dilemma, directedness is not a good candidate for the PSD because it is deeply uninformative, leaving phenomenal intentionality highly mysterious. Although Loar's lemon-experience thought experiment may enable us to better pinpoint that phenomenal feature which comprises phenomenal intentionality, we are no closer to understanding how phenomenal intentionality itself works, or how it comes to be instantiated. Instead, the *sui generis* nature of directedness resigns us to accepting that this phenomenal feature is the end-point of our explanation: directedness simply is what it is, and does what it does.

On the other hand, we could drop the insistence that directedness is a *sui generis* phenomenal feature, and allow that it is something complex and composite—something which can be broken-down or analysed into simpler or more basic phenomenal features; features which are not intentional in and of themselves, but only when they come together in the appropriate way.⁶⁰ But in adopting this view we abandon what is distinctive about *sui generis* phenomenal directedness as a candidate for the PSD—we indicate that a better candidate for the PSD is to be found in a complex composition of phenomenal features amongst which there is *no sui generis* directedness. So on this horn of the dilemma, we effectively reject *sui generis* phenomenal directedness as a candidate for the PSD.

⁶⁰ For if we allow that any more basic component of the composite is itself intentional, then of course we encounter exactly the same problem of a *sui generis* phenomenal feature being an uninformative way of accounting for what phenomenal intentionality consists in.

The dilemma presented illustrates that *sui generis* phenomenal directedness does not fare well as a candidate for the PSD. The phenomenal feature of directedness is either basic or composite. Taking it to be basic leaves us with an uninformative account which does nothing to reduce the mysteriousness of phenomenal intentionality, and taking it to be composite contradicts the very suggestion that the PSD is a *sui generis* phenomenal feature. Therefore, we either endorse *sui generis* phenomenal directedness as an uninformative account of phenomenal intentionality, or do not endorse it at all.

This is not to say, of course, that the notion of a phenomenal feature of *directedness*, independent of any assertion that it is *sui generis*, is uninformative. As we have seen, directedness fulfils Desideratum A independently of any assertion that it is a *sui generis* feature. And if we opt for the second horn of the dilemma set out above, by considering directedness to be a *composite* phenomenal feature, we avoid the outright uninformative nature of positing a *basic sui generis* feature. Although we thereby eliminate *sui generis* phenomenal directedness from the running as a candidate for the PSD, the notion of phenomenal directedness itself, and Loar's phenomenological means of picking it out, provide us with a promising lead on a better candidate for the PSD. By observing that phenomenal intentionality resides in *the way in which an object is visually presented*—which remains constant across hallucinations and perceptions of indistinguishable objects, and across non-referring and successfully-referring demonstrative thoughts about indistinguishable objects—Loar homes in upon that aspect of experience responsible for phenomenal intentionality, and thereby narrows our search for whatever compositional feature

of experience might constitute the PSD. In sum, although phenomenal directedness is not itself a convincing candidate for the PSD, it does help establish where, in the space of all possible experiences, we can sensibly begin looking for a compositional feature of experience which might constitute phenomenal intentionality.

4.4 'Taking', Cognitive Intent, or Cognitive-Experiential Character

Galen Strawson (2008) argues that the PSD consists in what he variously refers to as 'taking', 'cognitive intent' and 'cognitive-experiential character'. This is supposed to be an experientially manifest phenomenal feature of experience, albeit one which is, according to Strawson, often overlooked.

Strawson claims that those analytical philosophers who have been willing to admit the existence of any phenomenal experience at all—of any *way it feels* for one to undergo an experience—have not paid due attention to the full gamut of phenomenal experience. He maintains that they have focused almost exclusively upon *sensory-affective* experience—spanning sensory experiences, sensations, emotions and moods—but have neglected *cognitive* experience, when the latter in fact accounts for most of our total experience. Counting as 'cognitive experience' "every aspect of experience that goes beyond sensory-affective experience considered just as such," Strawson claims that "the vast bulk of experience would classify as cognitive. It saturates everything—swimming and digging as much as philosophising" (2008, p. 292). According to Strawson, it is this *cognitive* dimension of our phenomenal character which provides the phenomenal conditions required for phenomenal intentionality.

More specifically, Strawson thinks that phenomenal intentionality is brought onto the scene by the cognitive experience of 'taking'. According to Strawson, undergoing an intentional experience—whether by perceiving something, hallucinating something, consciously thinking about something, or consciously desiring something—always involves, as part of its cognitive dimension, *taking* the experience to be about whatever it is about.

Strawson supplies some examples to illustrate what this 'taking' consists in. He focuses primarily on the experience of thinking about, visualizing or seeing a moose, all of which include the cognitive experience of *taking* the experience to be about a moose, regardless of whether one is actually causally connected to a moose, as in the case of a veridical perception and subsequent thoughts, or not causally connected, as in the case of a hallucination and subsequent thoughts. Such taking can also occur, Strawson points out, if one sees a rock obscured by mist, and mistakenly *takes* it to be a moose. He also provides some examples of contrast-cases, in which one's experience changes as soon as *taking* is added. For example, he describes the addition of *taking* to experience as "a bit like looking at one of those pictures where you can't see what it is a picture of, and then suddenly you see" (2008, p. 301), or as seeing a scene as "just an array of colours ... without any grasp of it as experience of anything at all", and then coming to see the colours as "buildings, leaves, whatever it is." Elsewhere, Strawson identifies another example of 'taking' in the "automatic and involuntary *taking* of sounds or marks *as* words and sentences that one understands and that represent something's being the case" (1994, pp. 7-8, original emphasis). He highlights this in action using the example of a monoglot

anglophone and a monoglot francophone listening to the French news—intuitively, the francophone will experience a *taking* which the anglophone lacks.⁶¹

Strawson makes a plausible case for the existence of the sort of experientially manifest phenomenal feature he calls ‘taking’, and as such his suggestion fulfils Desideratum A. However, ‘taking’ is ultimately unpromising as a candidate for the PSD, because it turns out to be a *sui generis* phenomenal feature. Strawson reveals his notion of ‘taking’ to be one of a *sui generis* feature when he provides his answer to the question of *how* ‘taking’ determines what an experience is about. Strawson maintains that “it just can” (2008, p. 301). Strawson acknowledges that this way of thinking about intentionality may “leave a feeling of dissatisfaction, partly because it

⁶¹ One might object that Strawson’s ‘taking’ is not a coherent candidate for the PSD because it is not alone sufficient for determining what an experience is about, but rather must be paired with the subject’s causal context. We can answer this objection as follows. Although the objector is correct that ‘taking’ alone is not supposed to be sufficient for determining what an experience is about, it is nonetheless supposed to be a purely phenomenal feature which occurs in one’s cognitive experience. It is therefore something which can be experienced identically by phenomenal duplicates in entirely different causal contexts. Strawson uses the example of Lucy and Louis (2008, pp. 295-299), both of whom have an identical phenomenal experience which involves *taking* their experience to be about Mandy the Moose, but which differs in its causal context, with Lucy’s experience genuinely caused by the actual Mandy, and Louis’ caused by a freakish brainstorm. Strawson maintains that, whereas Lucy’s causal context pairs with her taking to make her experience determinately about Mandy, Louis’ experience fails to be about Mandy, precisely because it is not paired with the causal context required to make this so. Nonetheless, Lucy and Louis can be understood to share phenomenal intentionality in the sense that they share a ‘taking’ which *partly* determines what their experience is about, in the sense that it can pair with any possible causal context to determine what the experience is about. In other words, the ‘taking’ provides a determination of what the experience is about for any given causal context. In this sense, the ‘taking’ alone is sufficient for a degree of intentionality, even if it is not fully determinate. And given that ‘taking’ is a purely phenomenal feature, the degree of intentionality it is sufficient for qualifies as *phenomenal* intentionality. Therefore, ‘taking’ is a coherent candidate for the PSD.

terminates in something primitive, not further analysable” (p. 302), namely the ‘taking’ under consideration. This feeling of dissatisfaction is, of course, a symptom of the unformativeness of positing a *sui generis* phenomenal feature as the PSD. One’s queries about how phenomenal intentionality works and how it comes to be instantiated are met with bald assertions that it simply does work, and simply is instantiated. As such, the *sui generis* cognitive experience of ‘taking’ is not a viable candidate for the PSD.

However, as with Loar’s ‘directedness’, we can perhaps salvage *something* informative from the ashes of Strawson’s ‘taking’. Just as we did with ‘directedness’, we could drop Strawson’s insistence that ‘taking’ is a *primitive, sui generis* phenomenal feature of experience, and allow that it could be a *composite* phenomenal feature. And as we did with Loar’s ‘directedness’ we could try to utilise Strawson’s notion of ‘taking’, and his means of picking out the experiences which instantiate it, in order to narrow our search for whatever compositional feature of experience might constitute the PSD. However, where Loar observes that phenomenal intentionality resides in *the way in which an object is visually presented*, Strawson merely observes that phenomenal intentionality occurs when we *take* things we experience *to be certain sorts of things*, or when we *recognise* things we experience *as being certain sorts of things*. Whereas Loar narrows our focus onto one particular aspect of experience where we can find phenomenal intentionality instantiated, Strawson merely redescribes the very phenomenon of phenomenal intentionality which we are interested in explaining—the phenomenon of experiencing things to be certain sorts of things, or of having experiences *about*

certain things. As such, Strawson's 'taking' remains uninformative even if we drop the insistence that it is a *sui generis* feature, for it does not help us to home in upon an aspect of experience which may be responsible for phenomenal intentionality, but instead simply redescribes the phenomenon of phenomenal intentionality which we are seeking to explain.

4.5 Transparency

As covered in Chapter 2, one of the central motivations for taking conscious experience to be representational or intentional is the so-called *transparency of experience*: the characteristic experience is alleged to possess, of presenting to our attention or awareness only things the experience represents, and not properties of the experience itself. On initial reflection, the transparency of experience seems to be a promising candidate for the PSD—transparency is supposed to be an experientially manifest phenomenal feature of experience, so it purports to fulfil Desideratum A, and all supposedly transparent experiences are intentional, so perhaps their transparency is what makes them intentional.

As Kriegel (2011, fn. 100) points out, it is counterintuitive to take transparency to be the PSD, because it is natural to think of transparency as a *symptom* of an experience's intentionality. In other words, it is intuitive to think of an experience's intentionality as being logically prior to its transparency. In order for one's attention or awareness to focus exclusively on what an experience represents, it is clearly necessary that an experience represents something. It is thus a precondition of an experience's transparency that it is intentional; if an experience is not intentional,

then it cannot be transparent. It seems obvious that the intentionality of experience is the more basic or primitive fact which explains how it can be transparent. Indeed, this is precisely why representationalists such as Harman (1990, p. 39) and Tye (1995, p. 30) appeal to transparency to establish the intentionality of experience. It is therefore highly counterintuitive to reverse the order of explanation by maintaining that an experience is intentional *because* it is transparent, as one does in claiming that transparency is the PSD.

This counterintuitiveness, however, is the least of our worries concerning transparency's candidacy for the PSD. As we saw in Chapter 2 (§2.4), there is a distinction to be drawn between the claim that experience is *phenomenologically transparent*, and the claim that experience is *metaphysically transparent*. According to the former, experience *seems* to present only properties it represents things as having, and no properties of its own. According to the latter, experience *actually* presents only properties it represents things as having, and no properties of its own. Metaphysical transparency is not a viable candidate for the PSD, because as a state-of-affairs concerning the actual nature of experience, it does not specify any set of phenomenal conditions. In other words, the claim that an experience is metaphysically transparent says nothing about what that experience is like for the subject to undergo, from their own first-person perspective. Therefore, metaphysical transparency does not fulfil Desideratum A. Unlike metaphysical transparency, phenomenological transparency *does purport* to specify a set of phenomenal conditions, for phenomenological transparency consists in it *seeming*, to the subject, that their experience presents only what it represents, and no properties of its own.

However, as we saw in Chapter 3 (§3.3), only *e-seemings* are experientially manifest, and it is a mistake to think that experience *e-seems* to be any particular way with regard to its own metaphysics, including the metaphysical constitution of what it presents to the subject. Therefore, it is not plausible that phenomenological transparency specifies a set of phenomenal conditions which are *genuinely experientially manifest*. As such, phenomenological transparency also fails to fulfil Desideratum A.

Beyond this, there are two other ways of conceiving of transparency which might have more success in fulfilling Desideratum A. However, the choice between these returns us to the same dilemma presented by Loar's *sui generis* phenomenal directedness: a choice between a *sui generis* phenomenal feature which is deeply uninformative, and a composite phenomenal feature yet to be characterised, which effectively constitutes a distinct candidate for the PSD.

The first remaining way of construing of transparency is as a *sui generis* phenomenal feature of experience. This is effectively an identical suggestion to Loar's *sui generis* phenomenal feature of 'directedness', which Loar (2003/2017, pp. 317-18) himself describes in terms of the transparency or diaphanousness of perceptual experience, at one point characterising directedness as "diaphanousness without the actual object", or "diaphanousness towards intentional objects, whether 'mere' or real."⁶² We have already established that *sui generis* phenomenal features such as

⁶² The notion of an 'intentional object' is the notion of the object-as-experienced, or the object-as-thought-about, which is real in the case of veridical perceptions, but 'merely intentional' in the case of hallucinations. As Loar (2003/2017, pp. 305-306) explains, representationalists appeal to this notion in order to extend their

this are unsatisfactory as candidates for the PSD, given that they are deeply uninformative.

The second way of construing transparency is as a composite phenomenal feature, rather than a basic, *sui generis* one. Just like Loar's directedness, transparency is not itself a promising candidate for the PSD, but nonetheless indicates which experiences, if any, must instantiate the composite phenomenal feature which constitutes phenomenal intentionality, and thereby helps to narrow our search for a better candidate for the PSD. We have already seen that Loar's notion of phenomenal directedness points towards *the way in which an object is visually presented*, in both hallucinations and veridical perceptions of an object, as an aspect of experience which can support phenomenal intentionality. So what aspect of experience does the notion of transparency point us towards?

As we saw in Chapter 2 (§2.2), there are certain experiences—such as visually perceiving trees or blue squares—which are cited as paradigm cases of 'transparent' experience, and other experiences—such as blurred vision, phosphene experiences, afterimages and orgasm—which are cited as paradigm cases of less 'transparent', more 'opaque' experiences. In Chapter 3 (§3.3), we saw that representationalists and qualia realists might not actually disagree over the 'phenomenological data' captured in their respective examples of paradigmatically 'transparent' and 'opaque' experiences. Nonetheless, it is undeniable that paradigmatically 'transparent'

metaphysical account of experience to hallucinatory experiences, such that they can maintain that experience is constituted by whatever the experience *represents*, whilst allowing that it might represent an object which is not real.

experiences have an extremely tangible intentionality which paradigmatically 'opaque' experiences lack. It is, for example, extremely clear what a perceptual experience of a tree is about, but it is not clear what phosphene experiences or after-images are about, or even clear that they are about anything at all. This is, no doubt, precisely why representationalists have appealed to paradigmatically 'transparent' experiences to support their metaphysical view, and qualia theorists have appealed to paradigmatically 'opaque' experiences to support theirs. So, what can we glean from the paradigmatic examples of each sort of experience, about what phenomenal conditions make an experience paradigmatically 'transparent' and grant them their extremely tangible intentionality? From the sorts of experience cited as paradigmatic examples of each, we can reasonably conclude that the phenomenal conditions for transparency are those most clearly instantiated when we have a clear perception of something external to us in the world, and that the conditions for opacity are those most clearly instantiated when we have a less clear, more ambiguous sensation, feeling or experience of something internal to us, or some way things are within us. The notion of transparency, then, points us towards clear perceptions of external objects, and away from sensations, feelings and experiences relating to how things are within us.

4.6 Phenomenal Presence

A fifth candidate for the PSD is suggested by Christopher Frey (2013), in an effort to rectify problems he sees with the paradigmatic notion of experiential transparency. Frey argues that the paradigmatic transparency claim makes

unwarranted commitments about the metaphysical characteristics that phenomenal qualities are *appreciated as* having in experience—namely, the characteristics of being public, objective, mind-independent and/or external (pp. 72-75).⁶³ As such, he proposes a different conception of transparency, namely ‘Core Transparency’ (CT), which is free from these commitments (p. 76). On Frey’s conception, transparency consists in phenomenal qualities being appreciated as *other* (CT1), and never being appreciated as aspects of *the self*, or as being about *the self*, considered as an experiential subject (CT2).⁶⁴ Frey argues that the phenomenological observations of the transparency advocates reveal this phenomenally manifest opposition between self and other, and thereby shed light upon a basic and intrinsic structural feature of phenomenal experience, which he calls ‘phenomenal presence’ (pp. 75-79). Frey then claims that phenomenal presence, by ensuring that phenomenal qualities are *appreciated as other*, makes experience “universally and intrinsically *other-directed*”, and thereby “realizes a basic and non-derivative form of intentional directedness” (pp. 86-87, original emphasis). Given that phenomenal presence is put forward as a purely phenomenal feature, any intentionality it supports qualifies as a form of phenomenal intentionality. Phenomenal presence is therefore taken to capture

⁶³ As mentioned in the previous chapter (§3.3, fn. 40), Frey makes his argument in relation to ‘sensuous elements’ rather than ‘phenomenal qualities’, but his examples of sensuous elements suggest that the terms are referentially equivalent.

⁶⁴ Frey stresses that, in his formulation of core transparency, a phenomenal quality being ‘appreciated as’ *the other* or as part of *the self* should not be considered to consist in an attribution of a property or a classification as a kind. This is because these would both qualify as truth-evaluable claims about phenomenal qualities, and hence would make experience vocal about its own metaphysics, when in fact we know it to be silent or neutral in this regard (p. 77).

those phenomenal conditions necessary and sufficient for phenomenal intentionality, and is thereby nominated as a candidate for the PSD.

Phenomenal presence is a coherent candidate for the PSD, in that it purports to fulfil both of our desiderata. As touched upon in Chapter 3 (§3.3), Frey takes great care to separate the metaphysical and the phenomenological, recognising that phenomenological reflection is entirely neutral with respect to vying metaphysical accounts of phenomenal experience. Phenomenal presence is therefore put forward as an experientially manifest phenomenal feature of experience, and hence purports to fulfil Desideratum A. It also purports to fulfil Desideratum B further than the previous candidates. Frey stresses that phenomenal presence is not a relationship instantiated between distinct components, whereby *the other* is presented to *the self*, but is rather the “basic and intrinsic phenomenal structure” of experience (p. 76). Although this means that phenomenal presence is a *sui generis* structural feature of phenomenal experience, it promises not to leave phenomenal intentionality quite so mysterious as Loar’s ‘directedness’ and Strawson’s ‘taking’ do. ‘Directedness’ and ‘taking’ leave phenomenal intentionality mysterious because their characterisations simply presuppose intentionality.⁶⁵ As such, these candidates provide no insight into what makes the difference between the absence and presence of phenomenal intentionality, beyond phenomenal intentionality popping into existence with the

⁶⁵ In the case of ‘directedness’, ‘purporting to refer’ is simply a way of saying ‘being about something’ or ‘having intentionality’ whilst cancelling any implication of successful reference. And in the case of ‘taking’, ‘taking experiences to be about certain things’ or ‘seeing or understanding things as certain things’ seem equivalent to ‘having experiences which are of or about certain things’, or ‘having intentional experiences’.

relevant *sui generis* feature. By contrast, phenomenal presence is characterised in terms of a basic structural opposition in phenomenal experience between *the self* and *the other*, and this characterisation ostensibly does not presuppose intentionality. It thereby holds promise of providing us with an insight into how the basic configuration of phenomenal experience is responsible for bringing phenomenal intentionality into existence. Because it delves into the basic configuration of phenomenal experience, the candidate of phenomenal presence also has an advantage over the candidate of phenomenality itself, because it promises to tell us precisely what it is about phenomenal experience which brings phenomenal intentionality about, without being beholden to a solution to the hard problem of consciousness. In sum, phenomenal presence certainly purports to fulfil our desiderata, and to a greater extent than the previous candidates. Let us now see whether phenomenal presence can deliver on its promises, and provide us with a convincing account of the PSD.

As just established, phenomenal presence ostensibly appears more promising than the previously considered *sui generis* candidates, because it purports to delve into the structural configuration of phenomenal experience, and to potentially explain phenomenal intentionality in non-intentional terms. But on closer inspection, three problems emerge. First, phenomenal presence is only a 'structural' feature in a very tenuous sense, and in fact leaves phenomenal intentionality equally as mysterious as the previous *sui generis* candidates do. Second, Frey's case for the existence of phenomenal presence relies upon phenomenological observations which in fact appear false. Third, phenomenal presence only provides conditions

sufficient for an extremely minimal form of intentionality, and in this regard is less satisfactory than the candidates already considered. I will explain each of these in turn.

Frey's description of phenomenal presence often seems to portray an asymmetric bipartite structure. *The self* and *the other* are said to be in "opposition", with *the other* being "that which is present before the self" and *the self* being "that before which the other is present" (p. 77), and this situation is described as the "bipartite, phenomenal articulation" of phenomenal experience. This kind of structural picture looks like a promising way of explaining phenomenal intentionality, precisely because, as a *structural* picture, it promises a way of *breaking down* phenomenally intentional experience into components which are not themselves intentional, but which produce phenomenal intentionality when configured together in the appropriate way. However, Frey clearly makes a concerted effort to quash any implication that experience can be *broken down* into the elements of *the self* and *the other* which he distinguishes. As mentioned earlier, Frey insists that phenomenal presence does not consist in any kind of relationship between distinct components. He also stresses that the structure is "basic and intrinsic", and characterises experiences as "phenomenally articulate *unities*" (p. 77, emphasis added), thereby emphasising that the structure he describes is not of the sort that can be taken apart or deconstructed. But a structure which cannot, in principle, be broken down, taken apart or deconstructed does not look like much of a structure at all. Rather, in describing a feature which is 'structural' in the sense that it concerns something's make-up or build, but which is not the sort of thing which can be broken down, Frey

seems to be describing a feature more closely analogous to *shape* or *form*. Conceiving of phenomenal presence as the *shape* or *form* of phenomenal experience makes significantly more sense, as such a conception aligns much more closely Frey's characterisation of the feature. Just as being cylindrical is a basic, intrinsic feature of the ten-pence coin in my pocket, phenomenal presence is supposed to be a basic, intrinsic feature of experience. And just as the sides and angles of a triangle are determined in relation to one-another, and stand in unity as a triangle, so *the self* and *the other* are supposed to be determined in relation to one-another, and stand in unity as a phenomenal experience.⁶⁶

Although phenomenal presence, conceived of as the *basic form* of experience, no longer promises to explain phenomenal intentionality by breaking experience down into non-intentional components, it still has the potential to explain phenomenal intentionality in terms of rudimentary or fundamental facts about experience's basic form, which may not themselves assume intentionality. Unfortunately, phenomenal presence does not achieve this, and in fact leaves phenomenal intentionality equally as mysterious as the previous *sui generis* candidates do. Just like 'directedness' and 'taking', phenomenal presence does not *explain* how phenomenal intentionality works or how it comes to be instantiated, but simply *asserts* that a particular kind of phenomenal intentionality is built-into the

⁶⁶ Frey patently considers *the self* and *the other* to be codeterminate in this sense, for he defines them in terms of one-another—"the self—that before which the other is present—and the other—that which is present before the self" (p. 77, original emphasis), and claims that "it is only through the presence of the other that the phenomenal self exists at all" (p. 79).

basic form of phenomenal experience. For phenomenal presence is supposed to consist in the basic, phenomenally manifest opposition between *the self* and *the other*. This opposition is supposed to be asymmetric, in that *the other* is that which is present before the self, and *the self* is that before which the other is present, and Frey clearly envisages this asymmetric opposition as a kind of directedness, calling it “the intrinsic orientation of experiential phenomenality toward the other”, and explaining that “experience is universally and intrinsically *directed* toward the other” (p. 79, original emphasis). And it is precisely this *directedness toward the other* which Frey considers to constitute “an essentially phenomenal form of intentional directedness” (p. 86), or in other words, a kind of phenomenal intentionality. Therefore, the kind of phenomenal intentionality at issue is simply asserted to obtain as part-and-parcel of phenomenal presence, and so phenomenal presence cannot be considered to explain this kind of phenomenal intentionality in non-intentional terms. Rather, Frey’s positing of phenomenal presence is simply the assertion of a *sui generis* feature of all phenomenal experience—its basic and intrinsic form, consisting in an opposition between *the self* and *the other*—which is necessary and sufficient for a kind of phenomenal intentionality consisting in a *directedness toward the other*. So in sum, Frey claims that phenomenal intentionality is a primitive feature of all phenomenal experience. As such, phenomenal presence leaves phenomenal intentionality equally as mysterious as the previous *sui generis* candidates do, making it equally uninformative, and equally unsatisfactory as a candidate for the PSD.

Even more damagingly for phenomenal presence, Frey’s case in support of its existence relies upon questionable phenomenological observations. The crux of

Frey's argument for phenomenal presence is his claim that core transparency *discloses* the basic structural feature of experience that he calls 'phenomenal presence'. Frey appears to take core transparency to 'disclose' phenomenal presence in the sense that phenomenal presence *best explains* the observations of core transparency. The first observation (CT1) is that phenomenal qualities are always appreciated as other, and the second (CT2) is that they are never appreciated as aspects of the self, or as being about the self, considered as an experiential subject. Frey emphasises the absolute scope of these observations (*always* and *never*) by asserting that they hold true even for the *ganzfeld experiences* and *spatially punctiliar experiences* mentioned earlier (§4.2). These are experiences in which phenomenal character obtains in its purest and simplest form—where the phenomenal conditions instantiated are at their most minimal—in the sense that the entire phenomenal character of the experience is in one modality (vision) and is of a single quality (a particular colour). Given that these are possible instances of phenomenal experience in its purest and simplest form, and given that CT1 and CT2 are supposed to have absolute scope over all experiences, including Ganzfeld experiences and spatially punctiliar experiences, it indeed appears that the best explanation of CT1 and CT2, if true, would reside in the basic and intrinsic nature of phenomenal experience. Furthermore, given that CT1 and CT2 observe a universal *appreciation as other*, and a universal absence of *appreciation as self*, their truth would be best explained by a basic and intrinsic feature which directed or oriented all experience toward the other, and away from the self. Therefore, if CT1 and CT2 were true, then phenomenal presence—or something equivalent in all significant respects—would indeed present

itself as the best explanation of their truth. So in this sense, CT1 and CT2, if true, *do* 'disclose' phenomenal presence.

However, we have good reason to think that CT1 and CT2 are false. Their truth is questionable, in part, precisely because of the ganzfeld experiences and spatially punctiliar experiences which are so important for securing phenomenal presence as the best explanation. Frey maintains that, when we have these experiences, and experience a single, pure phenomenal quality, we "appreciate it as being both something other than ourselves and as standing in opposition to ourselves" (p. 77). However, this is far from self-evident. It is highly plausible that there can be some phenomenally manifest opposition between *the self* and *the other*, whereby one feels oneself to be an *experiencer*, confronting *that which is experienced*, or feels oneself to be a *being* confronting *a world*. It is even plausible that such an opposition is manifest in *most* of our experiences. For example, it is plausible that the experience of seeing a lime involves the phenomenal seeming of *otherness*, as well as the phenomenal seeming of *greenness*, and does not involve any kind of seeming *to-be-an-aspect-of-or-to-be-about-oneself-ness*. However, it is not so plausible that this opposition is manifest in *all* experiences, including ganzfeld or spatially punctiliar experiences. In fact, it seems to be in the very nature of these experiences that they are *purely* of a phenomenal quality such as *greenness*, and therefore would not involve any additional phenomenal seeming of *otherness*. Even if Frey objects that this is not in the nature of the particular experiences he is referring to, it seems conceivable that one *could* have a ganzfeld or spatially punctiliar experience of pure *greenness*, without experiencing any *otherness*. Therefore, CT1 appears false in

claiming that phenomenal qualities are *always* 'appreciated as other'. It also appears that CT2 is false in claiming that phenomenal qualities are *never* appreciated as aspects of the self, or as being about the self, considered as an experiential subject. Consider the counterexamples of being drunk or feeling cloudy-headed. These seem to be cases in which the phenomenal qualities we experience *can be* appreciated as aspects of oneself as the experiential subject, for one can appreciate *the way in which one is experiencing, or the state or condition one is in as an experiencer*. One might even be able to do this when sober or clear-headed, by appreciating the contrasting condition they now find themselves in as an experiencer, compared to the night before. We therefore have good reason to think that both of Frey's core transparency claims are false, and this completely undermines his claim that core transparency 'discloses' phenomenal presence.

Finally, and most damagingly of all, phenomenal presence does not adequately capture the conditions necessary and sufficient for phenomenal intentionality in general, and is therefore unsuitable as a candidate for the PSD. This inadequacy results from the fact that phenomenal presence is supposed to realise a very *specific* kind of phenomenal intentionality, in combination with the fact that it is supposed to be a *basic* feature of experience. As established above, Frey takes phenomenal presence to realise a particular kind of phenomenal intentionality consisting in *directedness toward the other*. In consisting in directedness *toward the other*, this kind of phenomenal intentionality consists exclusively in experiencing something specifically *as other*. But the phenomenal intentionality we wish to explain is the general capacity of experiences to be about *something*, just by virtue of their

phenomenal character, where this ‘something’ stands in for any particular thing. We thus wish to explain how one can experience something *as a mug*, *as a dolphin*, or *as an insult*, and not merely how one can experience something *as other*.⁶⁷ Experiencing something *as other* may of course be a component of how we experience things *as mugs* or *as dolphins*, and in this respect it would be useful to bear in mind potential explanations of this other-directedness. It is tempting to think that one could identify the phenomenal conditions necessary and sufficient for phenomenal intentionality in general *within* the phenomenal conditions necessary and sufficient for other-directedness specifically—for how could phenomenal conditions suffice for phenomenal intentionality of a *specific* kind, without sufficing for phenomenal intentionality in *general*? This line of reasoning would be legitimate, were it not for the fact that the directedness toward the other with which Frey is concerned is supposed to be a *basic* feature of phenomenal experience. This means that the phenomenal conditions necessary and sufficient for other-directedness cannot be broken down. We therefore have no hope of isolating, within them, a more general set of phenomenal conditions. In sum, even if phenomenal presence were without its other faults, it would only provide us with an explanation of those phenomenal conditions necessary and sufficient for *other-directedness* specifically, and would be

⁶⁷ Unlike phenomenal presence, the previous candidates accommodate this generality because they do not specify a particular target for the phenomenal intentionality they are supposed to realise; ‘phenomenality itself’ allows one to have experiences about anything which can be phenomenally presented, Loar’s ‘directedness’ allows for experience to be directed towards anything, Strawson’s ‘taking’ allows one to ‘take’ something as anything, and ‘transparency’ allows one to transparently experience anything which is not an aspect of experience itself.

of no help in establishing the conditions necessary and sufficient for phenomenal intentionality in general.

In conclusion, phenomenal presence is unsuitable as a candidate for the PSD, because despite its initial promise, it in fact posits a *sui generis* feature which is deeply uninformative, lacks a convincing case in support of its existence, and cannot adequately capture the conditions necessary and sufficient for phenomenal intentionality in general, which is precisely what is required of the PSD.

4.7 A Further Problem with *Sui Generis* Candidates for the PSD

Before moving on to my final and preferred candidate for the PSD, it is worth recognising a further reason that any *sui generis* candidate for the PSD will be unsatisfactory, besides being deeply uninformative. This problem with *sui generis* candidates is a rather intricate one, and hence deserves the attention of its own section.

To hold that the PSD is *sui generis* is to maintain that the necessary and sufficient conditions for phenomenal intentionality consist in the instantiation of a particular phenomenal feature which is basic or fundamental, in the sense that it cannot be broken-down into simpler composite parts or elements. And we have so far recognised three different suggestions as to what *sui generis* feature the PSD consists in: directedness, taking, and phenomenal presence. In each case, the relevant *sui generis* feature is supposed to be responsible for the phenomenal intentionality of all phenomenally intentional experiences. In other words, all phenomenally intentional experiences are supposed to have the relevant *sui generis*

feature in common, for it is this feature which is supposed to make them phenomenally intentional.

As noted in the previous section, the phenomenal intentionality we wish to explain is the general capacity of experiences to be about *something*, just by virtue of their phenomenal character, where this ‘something’ stands in for any particular thing. We wish to explain how, just by virtue of their phenomenal character, experiences are capable of being about many different specific things, such as mugs, dolphins and insults. It stands to reason that, if an experience is about some specific thing, just by virtue of its phenomenal character, then there must be something in its phenomenal character which makes it about that specific thing. And if there is a difference in what specific things two phenomenally intentional experiences are about, then this must be due to a difference in their phenomenal character. In other words, there must be some set of phenomenal conditions which an experience must fulfil in order to be about a mug, and a different set of conditions which an experience must fulfil in order to be about a dolphin. There must be a different set of phenomenal conditions for every possible thing that a phenomenally intentional experience could be about, for it is in virtue of their different phenomenal characters that phenomenally intentional experiences are about different things.

We have just established that, if two phenomenally intentional experiences are about different things, then there must be something which their respective phenomenal characters do not hold in common—namely, whatever makes one experience about one thing, and the other experience about a different thing. And before this, we established that a *sui generis* PSD would be a basic phenomenal

feature *common to all* phenomenally intentional experiences, which grants each of them their phenomenal intentionality. Therefore, whatever particular *sui generis* phenomenal feature the PSD is supposed to consist in, this feature must be separate from whatever phenomenal features make different experiences about different things—for whilst the former must be held in common across phenomenally intentional experiences, the latter must differ. In other words, if the PSD is *sui generis*, then the phenomenal basis for phenomenal intentionality (the PSD), must be separate from the phenomenal basis for being about a specific thing. But this separation is immensely problematic.

The problem with this separation can be seen most clearly if we consider the aforementioned phenomenal bases as *explanations* given in terms of phenomenal character. The PSD, or the phenomenal basis for phenomenal intentionality, can be considered an *explanation* of why phenomenally intentional experiences are phenomenally intentional. And the phenomenal basis for a phenomenally intentional experience being about a specific thing can be considered an *explanation* of why a given phenomenally intentional experience is about a specific thing. From this perspective, it is clearly inexpedient to separate one's explanation of why an experience is *about* something from one's explanation of why an experience is about *something in particular*. It is inexpedient because the explanations are concerned with describing exactly the same state-of-affairs, simply with a different emphasis. The state-of-affairs in question is the same in both cases because any actual instantiation of phenomenal intentionality—of phenomenal *aboutness*—must necessarily be about *something in particular* in order to qualify as a genuine case of

aboutness.⁶⁸ Therefore, a phenomenally intentional experience will always be phenomenally *about* something *by virtue of* being about *something in particular*. In other words, an experience qualifies as having phenomenal intentionality *because* it is phenomenally about something in particular—a dolphin experience is phenomenally intentional *because* it is phenomenally about a dolphin, and the same goes for phenomenally intentional experiences about mugs, insults, or any other specific things.

To translate into philosophical terms-of-art, it makes sense to regard phenomenal intentionality as a determinable, of which phenomenal mug-aboutness and phenomenal dolphin-aboutness are determinates. Any instance of the determinable *must* be an instance of one determinate or another, just as any instance of the determinable colour red must be an instance of one more determinate shade or another, such as maroon or scarlet. Given that phenomenal intentionality is a determinable which can only ever be instantiated as one or another of its determinates, such as mug-aboutness or dolphin-aboutness, whatever phenomenal feature provides an experience with its determinate aboutness will *constitute* the phenomenal feature which provides an experience with its determinable phenomenal intentionality. For example, a phenomenally intentional experience

⁶⁸ I am not claiming that there cannot be any generality or vagueness in what a phenomenally intentional experience is about, but merely pointing out that there must be *some* demarcation of what the experience is about, even if the demarcation is broad or fuzzy. By analogy, in order for a person to qualify as *pointing at something*, they must at the very least be extending a finger or gesturing in one direction or another, even if it is unclear precisely where their finger is pointing, or even if they gesture generally towards everything in front of them.

about a mug will have its phenomenal intentionality by virtue of whatever phenomenal feature makes it about a mug. The phenomenal basis of mug-aboutness will *constitute* the phenomenal basis of that experience's phenomenal intentionality. Therefore, it is unreasonable to think that one could separate the phenomenal basis of an experience's phenomenal intentionality (the PSD) from the phenomenal basis of an experience's determinate aboutness. But as shown above, this is precisely what one must do if one wishes to maintain that the PSD is a *sui generis* phenomenal feature. Hence, any *sui generis* phenomenal feature is unsatisfactory as a candidate for the PSD.

There is a way in which one might attempt to uphold the view that phenomenal intentionality is *sui generis* in nature, without falling foul of the problem just explained. One might claim that there is no common *sui generis* phenomenal feature comprising the phenomenal intentionality of all phenomenally intentional states, but rather a unique *sui generis* phenomenal feature for each phenomenally intentional state about a different specific thing. However, this would mean denying that there is anything phenomenally in common between all phenomenally intentional experiences, by virtue of which all of them are phenomenally intentional. And denying this is equivalent to denying the existence of a PSD—of a common phenomenal signature shared by all phenomenally intentional experiences, which is constitutive of their phenomenal intentionality. Therefore, no satisfactory candidate for the PSD is to be found here either.

In conclusion, *sui generis* phenomenal features are unsatisfactory as candidates for the PSD, because they either require a problematic separation of the

phenomenal basis of an experience's phenomenal intentionality (the PSD) from the phenomenal basis of an experience's determinate aboutness, or they require denying that a PSD—a common phenomenal signature shared by all phenomenally intentional experiences—even exists. And this is all in addition to the problem identified in the previous sections, that *sui generis* candidates are deeply uninformative, seeing as they do nothing to reduce the mysteriousness of phenomenal intentionality. For these reasons, *sui generis* phenomenal features should not be taken seriously as candidates for the PSD.

4.8 The Structural Organisation of Experience

Our final candidate for the PSD is the *structural organisation* of experience. Katalin Farkas (2013) and Farid Masrour (2013) both argue that this structural feature is constitutive of a certain kind of phenomenal intentionality, which Farkas calls “perceptual intentionality” (pp. 100-101), and Masrour terms “phenomenal objectivity” (pp. 117-18). This is the kind of phenomenal intentionality paradigmatically exhibited by perceptual experiences, which consists in the apparent presentation to us of *an experience-independent external world*; a world made up of objects and properties, external to our bodies and minds, whose existence does not depend on us experiencing them. It takes little effort to recognise that perceptual experiences do indeed seem to present to us such a world. Consider whether it seems to you that this thesis, or the screen upon which it is displayed, exists in front of you, externally to you, and has certain properties. Now close your eyes for a moment, or look in the opposite direction. Perhaps put the thesis down so that you're not

touching it. Then please resume reading. Consider whether, throughout the experiential episode you just had, the essay or the screen seemed to continue to exist and retain their properties. Be sure to focus only upon how things *naturally seemed* to be, rather than reflecting more deeply on what substantive knowledge or evidence you had from moment to moment. I take it that your response to these considerations will be unequivocally affirmative; our perceptual experiences *do* seem to present to us an experience-independent external world.⁶⁹ This without doubt constitutes a kind of phenomenal intentionality, for it is an instance of phenomenal experience *being about* an external world, just by virtue of its phenomenal character. I will adopt Farkas' moniker of 'perceptual intentionality' for this kind of phenomenal

⁶⁹ It is important to recognise that this seeming 'presentation of an experience-independent external world' is not the same as *phenomenological transparency*, which consists in it seeming, to the subject, that their experience presents only what it represents, and not properties of its own. As we saw in Chapter 3 (§3.3), phenomenological transparency is not experientially manifest because it is a mistake to think that experience *e-seems* to be any particular way with regard to its own metaphysics, including the metaphysical constitution of what it presents to the subject. The seeming 'presentation of an experience-independent external world' is different, because it can be easily cashed-out in a way which remains neutral or silent on the metaphysics of experience. For example, upon seeing a tomato, it could plausibly e-seem to one that there is a red, round thing located out in the world at some distance from one, and after looking away from it and then looking back, that thing could plausibly e-seem to have persisted in its external location, and in its being red and round, despite the temporary absence of these e-seemings. There is no obvious reason not to accept that there could be such experientially manifest e-seemings, given that none of them say anything about the metaphysical constitution of experience itself. As such, we can account for the seeming 'presentation of an experience-independent external world' just in terms of e-seemings, and so, unlike phenomenological transparency, we can admit it as an experientially manifest e-seeming.

intentionality,⁷⁰ though it is important to recognise that it may occur beyond instances of perception, in cases of hallucination, imagination or dreaming, for example.

It is also important to recognise that perceptual intentionality is not exhaustive of phenomenal intentionality. Consider the tangible aboutness exhibited by our conscious cognitive states, such as thoughts and beliefs. It seems that we could have a conscious belief that $2 + 2 = 4$, or a conscious thought about ghosts, without either presenting us with an experience-independent external world. Although these states lack perceptual intentionality, they are nonetheless about things by virtue of their phenomenal character. Thus, perceptual intentionality does not exhaust phenomenal intentionality.

Although Farkas and Masrour confine themselves to the domain of perceptual intentionality, arguing only that this particular kind of phenomenal intentionality is constituted by the structural organisation of experience,⁷¹ I will argue that the structural organisation of experience is a promising candidate for the phenomenal feature which constitutes phenomenal intentionality in general—that is, a promising candidate for the PSD. I will first explain what is meant by the structural organisation of experience, and how it is supposed to constitute perceptual intentionality. I will

⁷⁰ Masrour's term "phenomenal objectivity" refers to exactly the same phenomenon. As such, I employ 'perceptual intentionality' as a direct translation of Masrour's "phenomenal objectivity".

⁷¹ Although Masrour (2013, p. 118) suggests that one *could* argue that phenomenal intentionality in general is constituted by that which constitutes perceptual intentionality (or "phenomenal objectivity'), he does not himself argue for this, recognizing that one might "refrain from equating the phenomenal-intentional with the phenomenal-objective."

then show how it fulfils our desiderata, making it a coherent candidate for the PSD, before explaining why it is in fact the most promising candidate out of all of those considered.

What I am calling 'structural organization' consists in the way in which the various phenomenal features of perceptual experience ordinarily hang together. Masrour's paradigmatic example of phenomenal features hanging together with structural organization is that of veridical visual perception of a tree as one walks towards it. As Masrour points out, "your representation of the visual angle through which you see the tree and your representation of your distance from the tree co-vary with each other in a law-like manner" (p. 123), or in other words, the region of one's visual field that the tree occupies and the distance one experiences being from the tree change in accordance with one-another, as if according to a rule. To put it simply: as you walk towards a tree, the tree occupies a larger region of your visual field and you experience being closer to it; as you walk away from the tree, it occupies a smaller region of your visual field, and you experience being further from it. Farkas deploys a similar example relating to fire, describing how, as you move towards or away from a fire, the smell of burning, its brightness, the sound of crackling, and its warmth all correspondingly increase or decrease in intensity "in a way we are accustomed to in the perception of objects" (p. 109).

The point that Masrour and Farkas are trying to make is that our conscious experiences in ordinary perception do not simply consist in a disorganised flood of phenomenal qualities. It is not the case that ever-fluctuating patches of colour move around our visual field, accompanied by a stream of random sounds to which they

bear no discernible relation, along with an array of taste and smell sensations which come and go without any kind of regularity or order. Rather, these various phenomenal qualities do ‘hang together’ in a certain manner; they are “organized into a systematic, cross-modally coherent and predictable order” (Farkas, p. 110).

Farkas and Masrour both argue that this structural organisation of experience is constitutive of perceptual intentionality, and they both do so in part by showing that the absence and presence of perceptual intentionality correlates with the absence and presence of structural organisation. For example, both highlight afterimages as an example of an experience which does not exhibit perceptual intentionality—where “there is no overwhelming impression that some experience-independent object is presented” (Farkas, 109), such that, when we experience an afterimage whilst looking at a blank page, “we do not experience the afterimage as an objective pattern on the page” (Masrour, p. 126). And both point out that afterimages lack the structural organization which perceptual experiences ordinarily have. For example, an afterimage does not respond to our movement in the way that the rest of our experienced world does (Farkas, p. 109), and rather than reappearing immediately after we blink, an afterimage is absent for a short period, before gradually reappearing (Masrour, p. 126).

Masrour goes on to more thoroughly defend the view that there is a “perfect correlation” between perceptual intentionality and the structural organisation of ordinary perceptual experience,⁷² and argues that the best explanation for this

⁷² Masrour (2013, p. 123) calls this structural organisation “schematic dynamical unity”.

correlation is that the latter *constitutes* the former (p. 134). Farkas argues similarly, supplementing the afterimage example with a thought-experiment about pain. She establishes that we do not normally regard our actual pain experiences as perceptions of external objects, but rather as internal sensations. But she also describes two plausible scenarios in which we *would* regard our pain experiences as perceptions of external objects: one in which we experience pain on our skin only when it touches a specific patch of a leaf, and one in which we experience a headache whenever passing through a particular corridor (p. 103). Farkas makes the reasonable claim that such experiences would be regarded as perceptions of a ‘pain-patch’ on the leaf, and of a ‘headachy’ corridor—that is, perceptions of experience-independent external objects possessing the property of being painful, which therefore qualify as having perceptual intentionality. She then identifies the structural organisation of experience as a good explanation as to why the imagined pain experiences possess perceptual intentionality and our actual pain experiences do not: whereas the imagined pain experiences correspond to and co-vary with the other phenomenal qualities we experience in the same way that our experiences of a spatially located property would, our actual pain experiences tend not to exhibit the same kind of co-variance and correspondence, often persisting for a period regardless of the other phenomenal qualities we experience.⁷³

⁷³ To elucidate with more specific examples: whereas our pain in the imagined ‘pain-patch’ scenario co-varies with our experiences of our skin making contact with the specific patch of the leaf (such as the haptic sensation of touching the patch, the visual experience of seeing our skin touch the patch, or the faint auditory experience of hearing our skin brushing against the leaf), our actual pain when we stub our toe persists regardless of what other experiences we have or continue to have (whether

Farkas and Masrour make a good case for perceptual intentionality being constituted by the structural organisation characteristic of perceptual experiences. But why should we consider the structural organisation of experience to be a good candidate for the PSD—a good candidate for the phenomenal feature which constitutes phenomenal intentionality in general? First, it is important to recognise that ‘structural organisation’ is itself a highly generic notion. Farkas and Masrour use a variety of examples to illustrate what specific *sort* of structural organization they are concerned with, and Farkas describes the relevant kind of structural organisation as a “clustering around objects” (p. 110). We have little reason to think that this *specific sort* of structural organisation constitutes phenomenal intentionality in general, because it is not at all clear that non-perceptual phenomenally intentional experiences, such as consciously believing that $2 + 2 = 4$, or consciously thinking about ghosts, exhibit the ‘clustering around objects’ characteristic of perceptual experiences. Nonetheless, it does seem clear that such experiences have a clarity, distinctiveness and coherence suggestive of a structural organisation of *some sort*, even if not of the object-clustered kind. Furthermore, it is plausible that the various ways in which phenomenally intentional experiences are structurally organised share something in common, by virtue of which they all instantiate phenomenal intentionality in their various ways—even if there are specific kinds of phenomenal intentionality, such as perceptual intentionality, which are themselves united by an even more specific common structural organisation. Therefore, it is at least coherent

or not we see our toe hitting something, feel our toe touching something, or hear a forceful crunch).

to suggest that all phenomenally intentional experiences share some aspect of their structural organisation in common. I will now explain why an aspect of structural organisation, shared by all phenomenally intentional experiences, is a viable candidate for the PSD.

Some sort of structural organisation of experience is a viable candidate for the PSD because it fulfils both of our desiderata. It clearly purports to fulfil Desideratum A, because the structural organisation in question is the structural organisation that an experience is supposed to have *from the perspective of the experiencing subject*. In other words, it is precisely the notion of a sort of structural organisation which is *experientially manifest* to the subject. This is clear from the way that Farkas and Masrour's examples work: they describe how things seem from the perspective of the experiencing subject in a particular situation (e.g., moving around a fire), and highlight the fact that what is experientially manifest has its own structural organisation (e.g., the fire's smells, sounds and warmth correspondingly increase and decrease in intensity). One would be hard pressed to deny that there is any such experientially manifest structural organisation to our experiences. For it is undeniable that a wide array of distinct phenomenal qualities can be experientially manifest, and if this is so, then those qualities must either be manifest to us as a chaotic and disorganised flood, or as something less chaotic and disorganised, with at least some degree of structural organisation. Given that our experiences do not present us with a chaotic and disorganised flood of phenomenal qualities, it stands to reason that our experiences have some sort of experientially manifest structural organisation. As such, Farkas and Masrour's observations of experientially manifest

structural organisation are eminently plausible. Furthermore, there is no obvious reason for suspecting that their observations are mistaken. As such, we can consider Desideratum A sufficiently fulfilled at this stage.

Structural organisation also promises to fulfil Desideratum B. It would be informative to explain phenomenal intentionality in terms of the structural organisation of experience for two reasons. First, any particular structural organisation of experience is not a *sui generis* phenomenal feature, but rather an overarching feature comprised of the relations which various elements of the experience bear to one another. Therefore, it is analysable into more basic or fundamental components of experience. Second, the structural organisation of experience looks like something we can make sense of without having to presuppose the existence of some kind of intentionality. Therefore, structural organisation looks promising as a way of explaining what constitutes phenomenal intentionality in non-intentional terms. For these two reasons, the structural organisation of experience looks promising as a way of reducing the mysteriousness of how phenomenal intentionality works, and how it comes to be instantiated. So, given that it fulfils both of our desiderata, some sort of structural organisation of experience is a viable candidate for the PSD. But why should it be considered our *most promising* candidate?

There are three reasons why structural organisation should be considered our most promising candidate for the PSD, which I will explain in turn. First, because structural organisation is not a *sui generis* feature, it does not suffer from the serious problem described in the previous section. Rather, it entirely accommodates

phenomenal intentionality as a determinable *type* of experience, which is always instantiated as one or another determinate *token* experiences, such as an experience with phenomenal mug-aboutness. This is because any given structural organisation of experience is concerned with how different elements of experience relate to one another. As such, an aspect of structural organisation can be fulfilled in many different ways, by experiences which differ from one-another in many ways other than the relevant aspect of structural organisation. Therefore, we can coherently maintain that experiences with phenomenal mug-aboutness and phenomenal dolphin-aboutness both possess phenomenal intentionality by virtue of some shared aspect of structural organisation, whilst holding that they are about different things due to phenomenal differences in other respects. By analogy, a scarf and a jumper can both qualify as knitwear due to the structural organisation of thread which they have in common, but are nonetheless different items of clothing by virtue of differences beyond that structural organisation of thread. Because structural organisation doesn't suffer from the problem that *sui generis* candidates do, it is immediately more promising than most candidates we have considered.

The second reason why structural organisation is our most promising candidate is as follows. In earlier sections, I suggested that the notions of *directedness* and *transparency*, although not good candidates for the PSD in themselves, might nonetheless be helpful to us by indicating what sort of composite phenomenal feature might constitute phenomenal intentionality. Structural organisation is our most promising candidate for the PSD, partly because it is the only candidate that provides precisely that composite phenomenal feature which the

notions of both directedness and transparency both gesture towards. Paradigmatically transparent experiences possess precisely the kind of structural organisation Masrour and Farkas describe, and the structural organisation of experience is a coherent way of cashing out Loar's notion of *the way in which an object is visually presented*, which remains constant across hallucinations and perceptions of indistinguishable objects, and across non-referring and successfully-referring demonstrative thoughts about indistinguishable objects. In other words, structural organisation provides a constructive way of preserving the insights of directedness and transparency—it offers us a fine-grained explanation as to *what it is* about the phenomenal character of certain experiences which makes them 'transparent' or 'directed'. Thus, our reasons for originally considering transparency or directedness to be the PSD now lend support to the candidate of structural organisation. And given that the latter does not commit the error of the former two, of positing a sui generis PSD, it survives them as a highly promising candidate for the PSD.

Third, structural organisation is more promising than the first candidate we considered, of phenomenality itself, because it fully accommodates the possibility of experiences of a single phenomenal quality in isolation—such as ganzfelds and spatially punctiliar experiences—which do not possess phenomenal intentionality. If we hold that phenomenal intentionality is constituted by some sort of structural organisation, then we can explain why experiences of a single phenomenal quality in isolation lack phenomenal intentionality precisely in terms of their lack of structural organisation. Given that such experiences consist in a single phenomenal quality in

isolation, there is no sense in which the experience can exhibit any kind of structural organisation *amongst* various phenomenal elements. As such, taking structural organisation to be the PSD provides us with a coherent explanation of the possibility of experiences of a single phenomenal quality in isolation, and hence is preferable to the candidate of phenomenality itself.

4.9 Conclusion

In conclusion, some sort of structural organisation is our most promising candidate for the PSD, because it is less problematic and has more explanatory power than the other candidates. *Phenomenality itself* was revealed to be an unconvincing candidate, as it could not accommodate the possibility of experiences of a single phenomenal quality in isolation, which we have good reason to believe are possible. All suggestions that the PSD consists in a *sui generis* phenomenal feature, whether construed as *directedness*, *taking*, *transparency* or *phenomenal presence*, were shown to be unsatisfactory by virtue of being deeply uninformative. *Phenomenal presence* in particular also lacked a convincing case in support of its existence, and could not adequately capture the conditions necessary and sufficient for phenomenal intentionality in general, which is precisely what is required of the PSD. We also saw that there is a further serious problem with taking any *sui generis* phenomenal feature to be the PSD. Some sort of *structural organisation* emerges as our most promising candidate for the PSD because it does not have the aforementioned problems, because it explains what it is about the phenomenal character of certain experiences which makes them 'transparent' or 'directed', and because it explains

how it is possible to have experiences of a single phenomenal quality in isolation
which do not possess phenomenal intentionality.

5. Minimising the Risk of Error about how Phenomenal Intentionality is

'Given'

In the previous chapter, we reviewed some suggestions as to precisely what phenomenal conditions are constitutive of phenomenal intentionality. We found that the most promising suggestion was that phenomenal intentionality is constituted by some sort of *structural organisation* of more basic phenomenal elements which are not themselves intentional. In this chapter, I will assess this suggestion in light of the lessons of Part I, regarding the mistakes we are liable to make in what we take to be 'given' in phenomenally conscious experience. I will consider how these mistakes might be made specifically in relation to how we take phenomenal intentionality to be 'given', and assess whether the *structural organisation* account manages to avoid these mistakes. In the previous chapter, we saw that it is eminently plausible that the structural organisation of experience is experientially manifest or 'given' in experience. In this chapter, we will query the plausibility of this more rigorously, and consider whether it would be a mistake to identify structural organisation as the means by which phenomenal intentionality is 'given' in experience.

In order to explore how we might make mistakes in how we take phenomenal intentionality to be 'given' in experience, I will return to the distinction drawn between i-seemings and e-seemings in Chapter 3, and consider those seemings whose nature as either an i-seeming or e-seeming is unclear. We will see that so-called 'high-level properties' fall within this grey-zone of seemings, and that the dispute over whether or not high-level properties are experientially manifest is of central importance to the matter of how we take phenomenal intentionality to be

'given' in experience. I will show that, amidst significant disagreement over precisely what is 'given' in experience, there is widespread implicit agreement that structural organisation is 'given', and that it underpins the phenomenal intentionality of experiences. Given that the structural organisation account of phenomenal intentionality only commits to claims about what is 'given' which are broadly agreed upon, and is neutral regarding more contentious claims, I conclude that it carries minimal risk of error concerning what is 'given' in experience.

5.1 The Grey-Zone of Seemings

We have so far characterised e-seemings using the archetype of a phenomenal quality, which is patently manifest in our experience, and characterised i-seemings using the archetype of an inclination to believe in some metaphysical state-of-affairs, in which case that state-of-affairs is patently not manifest in our experience. But there is plausibly a significant grey-zone between these two archetypes, of seemings whose nature, as either an e-seeming or an i-seeming, is not immediately clear.

We can see more clearly where straightforward i-seemings end, and where the grey-zone begins, using Wilfrid Sellars' (1963) distinction between the "manifest image" and the "scientific image" of man-in-the-world. According to Sellars (p. 5), each image is a different conception of the world and our place within it. The manifest image roughly equates to the conception of the world we employ on an everyday basis, populated by perceptible objects such as "persons, animals, lower forms of life and 'merely material' things, like rivers and stones" (p. 9). By contrast, the scientific

image is the conception generated by our best scientific theories, which is populated by the imperceptible entities scientists have postulated to explain the behaviour of perceptible things, such as the quarks and leptons of particle physics (p. 7). Given that these are “equally non-arbitrary” conceptions (p. 5), “each of which purports to be a complete picture of man in the world” (p. 2), it falls to the philosopher to “attempt to see how they fall together in one stereoscopic view” (p. 5). At present, I am not concerned with fitting the two conceptions together, but rather with discerning which elements of each image are and are not manifest to us, in our phenomenally conscious experience.

As one would expect given Sellars’ dichotomy, the scientific image is patently not manifest to us, precisely because it is populated by imperceptible entities—entities which are not manifest to us in our experiences. As Montague (2016, pp. 38-39) points out, even though electrons partly constitute the objects that we see and feel, and as such are very much involved in our perceptual experiences, they are not part of what is “phenomenologically given in those experiences” (p. 38), for although we see and feel the objects, we do not thereby see or feel the electrons. Rather, we believe objects to be partly constituted of electrons, on the basis of scientific theory. In our terminology, it cannot e-seem that there are electrons, but it can i-seem that there are. And the same can be said for all postulates of the scientific image.

It is of course tempting to assume that, by contrast, the manifest image *is* experientially manifest to us. But in fact, this is not a cut-and-dried issue. It is controversial as to whether much of what we consider part of the manifest image is in fact *experientially manifest* to us. This controversy is no clearer than in the ongoing

dispute over which properties we perceptually experience, to which I will now turn my attention.

5.2 The High-Level / Low-Level Dispute

The dispute is over whether we only perceptually experience ‘low-level’ properties, or whether we also perceptually experience ‘high-level’ properties. As

Heather Logue explains:

“Low-level properties are the ones pretty much everyone agrees we can [perceptually] experience. ... By contrast, high-level properties are those that everyone agrees can figure in [perception]-based belief, but it’s controversial whether we can literally [perceptually] experience them.”^{74,75}

(Logue, 2013, pp. 1-2)

As such, the low-level properties of visual experience include colour, shape, size, location or position, distance or depth, and motion.⁷⁶ In auditory experience, they

⁷⁴ I have substituted ‘perceptually’ where Logue uses ‘visually’, because although Logue’s paper focuses exclusively on what properties we visually experience, her characterisation applies equally well to the wider debate about perceptual experiences in general.

⁷⁵ This might seem like an arbitrary way to distinguish between high-level and low-level properties, but as Grace Helton (2016, p. 852) points out, this arbitrariness is not problematic. The real aim of the debate is “to determine which features we can perceive”, with the high-level/low-level distinction used as “short hand for whether those features are among those that have traditionally been recognized as percepts.”

⁷⁶ Lists of the low-level properties of visual experience, varying in length, can be found in Siegel (2006, p. 482; 2016, §4.3), Bayne (2009, p. 385), Macpherson (2011, p. 8), Logue (2013, p. 2), Helton (2016, p. 852). Bayne and McClelland (2018, p. 2) list the most low-level properties: “Regarding vision, it is relatively uncontroversial that the following properties can figure in the contents of visual experience: colour, shape, illumination, spatial relations, motion, and texture.”

include volume, pitch and timbre, and in gustatory experience, they include sweetness and sourness.⁷⁷ These are all properties which feature in the manifest image of the world—in the world as we conceive of it on an everyday basis—and which we can, uncontroversially, also consider to be experientially manifest. As such, we can happily admit motion, timbre and sourness to the category of e-seemings: a violinist's bow can e-seem to move, their violin's sound can e-seem piercing, and the violin's rosin can e-seem sour.

High-level properties would also standardly feature in our everyday, manifest image of the world, but it is controversial as to whether they are experientially manifest. Logue provides a comprehensive list of example high-level properties:

- natural kind properties (e.g., being a banana),
- artifactual kind properties (e.g., being a table),
- semantic properties (e.g., experiencing a bit of text, say, 'las bananas son amarillas', as meaning that bananas are yellow),
- causal properties [e.g., experiencing a ball's impact as the cause of a window shattering]
- dispositional properties (e.g., being edible),
- others' mental states (e.g., being sad), and
- evaluative properties, such as
 - moral properties (e.g., being morally wrong, being kind), or
 - aesthetic properties (e.g., being graceful).

(Logue, 2013, p. 2)

It is clear that all or most of these properties feature in our everyday manifest image of the world. It is evident that, according to our everyday conception, the world

⁷⁷ I take these examples of low-level auditory and gustatory properties from Bayne (2009, p. 385). Helton (2016, p. 852) also lists volume and pitch as paradigm low-level auditory properties.

contains bananas we can eat, tables we can dance upon, meaningful text (such as this sentence), things which cause other things, people who can be sad or kind, and cats which can be graceful. But it is not clear whether any of these properties are *experientially manifest* to us in the same way that colour, timbre and sweetness are. To clarify with an example, it is clear that yellowness and curviness of a banana can be experientially manifest to one, and it is clear that one can be inclined to believe that certain yellow, curvy objects *are* bananas. However, it is not clear that an object's *being a banana* can be experientially manifest to one, over and above low-level properties such as yellowness and curviness. In other words, it is not clear that there is such a thing as experientially manifest *bananeness*, alongside yellowness and curviness. It is clear that objects can e-seem yellow and curvy, and can i-seem to be bananas, but it is not clear that objects can e-seem to be bananas. And this generalises to every other high-level property: although it is clear that an object can i-seem to be a table, that a ball can i-seem to cause a window to shatter, and that a person can i-seem sad or kind, it is not at all clear that an object can e-seem to be a table, that a ball can e-seem to cause a window to shatter, or that a person can e-seem sad or kind.

Given that their nature as either an i-seeming or an e-seeming is unclear, high-level properties fall well within the grey-zone of seemings. Should we consider all the high-level properties to be like the properties of the scientific image, which we can be inclined to believe in but cannot literally experience, or should we consider any of them to be like the low-level properties, which we do literally experience? The debate over which view to take is characterised by Tim Bayne (2009, pp. 385-89) as one

between “conservatives” who hold that only low-level properties are “phenomenally represented” in perceptual experience, and “liberals” who hold that at least some high-level properties are also phenomenally represented. Similarly, Kriegel (2007, p. 118)—who uses the term “phenomenologically manifest” where I use the term ‘experientially manifest’—characterises the debate as one between “*phenomenological deflationism*[:] the tendency to go exclusive and starve the phenomenology” and “*phenomenological inflationism*[:] the tendency to go inclusive and bloat the phenomenology”.

5.3 Distinguishing What is Experientially Manifest from What Is Perceived

Before we see how the debate just described bears upon the matter of how phenomenal intentionality is ‘given’ in experience, it is important to distinguish the question which we are interested in, of ‘*what is experientially manifest in experience?*’, from the question of ‘*what properties do we perceive?*’. Because the debate just described focuses heavily on perceptual experience, it would be easy to conflate these two questions. And in fact, the issue at the heart of the debate is sometimes characterised as a question of whether or not we *perceive* high-level properties, rather than as a question of whether high-level properties are *experienced*.⁷⁸ At other times, it is simply assumed that what can be perceived must

⁷⁸ Helton (2016), for example, states that “several theorists have proposed that we can *perceive* a range of high-level features”, and that “some theorists deny that we *perceive* any—or many—high-level features” (pp. 851-52, emphasis added).

be experientially manifest.⁷⁹ The question of whether or not we perceive high-level properties is a legitimate and interesting one, but it is not the question we are interested in, so it would be wise to avoid conflating them. Fortunately, my distinction between i-seemings and e-seemings helps us to retain focus on the question of ‘*what is experientially manifest in experience?*’ and keep it separate from the distinct question of ‘*what properties do we perceive?*’. Let us briefly see why the questions differ, and how my distinction helps to separate them.

The difference between the two questions is most obvious when we consider the possibility of extra-perceptual experience. It is broadly accepted, for example, that we can consciously experience feelings associated with emotions or moods, such as feelings of joy, excitement, anger and fear.⁸⁰ Furthermore, many philosophers think that there are phenomenally conscious experiences associated with our propositional attitudes and thoughts—with our states of believing, desiring, judging and thinking. If there are such extra-perceptual experiences, then there will be instances in which what is experientially manifest to one is not some property which

⁷⁹ Siegel (2006, pp. 481-82), for example, claims that “if one can perceive that a surface is round, roundness, the property, can be represented in visual experience.” I take it that anything which is ‘represented in visual experience’ would be considered to be experientially manifest in visual experience. Although I think it is uncontroversial that roundness is indeed an experientially manifest property, I take issue with Siegel’s assumption that it must be experientially manifest, given that it is perceived.

⁸⁰ For example, Kriegel (2007), states that “there is, doubtless, a phenomenology of *emotional* and *somatic* experiences, such as feeling angry or pleasantly ticklish” (p. 119, emphasis in original), and Bayne and Montague (2011, p. 1), enumerating what they take to be phenomenally conscious states, claim that “there are the conscious states associated with emotions and moods, such as feelings of elation, boredom, fear and anxiety.”

is being perceived. In other words, it is plausible that there can be non-perceptual e-seemings. Equipping ourselves with the notion of an e-seeming, distinct from the notion of what is perceived, enables us to respect the fact that what is experientially manifest isn't necessarily something which is perceived.

The difference between the experientially manifest and the perceived is more subtle when it comes to considerations of what is experientially manifest just within perceptual experience, for it is tempting to think that everything which is experientially manifest in perceptual experience is perceived, and vice versa. We can use an example from Kriegel (2007, p. 116) to illustrate that this is not necessarily the case.⁸¹ When we see a coin, there is a sense in which we see it as being round. We see it as round even if, as is often the case, we are looking at it from an angle, such that the light reflected from the coin casts an elliptical shape on our retina. In such a case, we can ask whether the coin is *perceived* as elliptical or round. One might hold that we have an elliptical sensation in our visual field, but perceive the coin as round. But even if we accept this view, it remains a further question as to what is experientially manifest—is it only the elliptical sensation, or only the perceived roundness, or is it both the ellipticality and the roundness?⁸² Does the coin e-seem

⁸¹ Although I will use the example to establish roughly the same point, I disagree with a number of claims Kriegel makes along the way (see following footnote), and so my use of the example here is significantly different.

⁸² Kriegel does not agree that there remains a further question about what is experientially manifest. He maintains that if you have an elliptical sensation, but perceive the coin as round, then it is roundness that is, in his terminology, “phenomenologically manifest”. Similarly, he thinks that if you perceive the coin as elliptical, but believe it to be round, then it is ellipticality that is manifest. In fact, the only case in which he thinks there remains an open question about what is manifest

elliptical, round, or both? One's answer to the question of what is perceived does not settle one's answer to the question of what is experientially manifest, or how things e-seem, and as such they are revealed to be different questions. Having the notion of an e-seeming, distinct from the notion of what is perceived, enables us to respect the fact that the question of what is experientially manifest is distinct from the question of what is perceived, even when it comes to perceptual experience.

The notion of an i-seeming also helps us to respect the difference between the experientially manifest and the perceived. Specifically, it helps us to accommodate the fact that what is perceived is not fully determined by what is experientially manifest. For example, when I see my wife and I am not wearing my glasses, blurriness is experientially manifest to me, but I do not perceive my wife to be blurry. Nor do I strictly *perceive* my wife to have a determinate physical outline, because there is no determinate physical outline experientially manifest to me in the way that the blurry outline is. Nonetheless, there is still a sense in which I take my

is that in which one perceives both ellipticality and roundness, due to the apparent impossibility of a phenomenal 'superposition' of an ellipsis and circle.

I reject Kriegel's position for three reasons. First, he gives no reason for thinking that sensations cannot be experientially manifest. In the absence of any such reason, the view that sensations are not experientially manifest is entirely indefensible, given that sensations such as coldness, blueness or sweetness are paradigmatic instances of experientially manifest properties. Second, his stance on the first two cases described above seems to be guided by the implicit view that what is perceived must be experientially manifest. However, he provides no reason to support this view, and in any case the view is inconsistent with his claim that there remains an open question about what is manifest when both ellipticality and roundness are perceived. Third, Kriegel claims that we do not experience a superposition of an ellipsis and circle, without providing any reason for thinking that this is the case. In the absence of any such reason, I do not find it implausible that we experience such a superposition.

wife to have a determinate physical outline. The notion of an i-seeming allows us to explain the sense in which I non-perceptually take my wife to have a determinate physical outline: I have an i-seeming or an inclination to believe that that she has a determinate physical outline, even though she e-seems blurry. Furthermore, the fact that I do not straightforwardly perceive my wife to be blurry, despite the e-seeming blurriness, can be explained in terms of what I perceive (and do not perceive) being informed by what I am inclined to believe, or what i-seems to be the case. This also explains how one might perceive more than what e-seems to be the case. For example, it seems plausible that I could perceive a wall to be impermeable, despite never having touched or prodded it to check that it is not merely a convincing piece of set-design which I could put my hand right through. This can be explained in terms of the perception of impermeability being informed by what one is inclined to believe, or what i-seems to be the case, perhaps on the basis of visual e-seemings that the wall shares with walls one knows to be impermeable. In providing the explanations to accommodate these situations in which one doesn't simply perceive what is experientially manifest, the distinction between i-seemings and e-seemings again helps us to separate the matter of what is perceived from the matter of what is experientially manifest.

Despite the risk of conflating what is perceived with what is experientially manifest, the debate concerning high-level and low-level properties nonetheless seems to focus primarily on the question we are concerned with, of what is experientially manifest, even if it is sometimes articulated in terms of what is perceived. For as Helton (2016) points out,

“The claim that perception represents high-level features is ... generally construed as the claim that perception represents high-level features in a way that contributes to the ‘phenomenal’ or subjectively accessible aspect of perception. Thus, the claim that we can visually perceive some object as a cat is committed to the claim that the attribution of being a cat to that object makes a constitutive contribution to ‘what it is like’ for us to see that object...”

(Helton, 2016, p. 852)

As such, even though the debate is sometimes presented as a debate over whether or not we perceive high-level properties, in practice it tends to turn upon the issue of whether or not high-level properties are manifest in our phenomenally conscious experience. This is reflected in the fact that Siegel (2006) considers herself to be addressing the question of what properties are “represented in visual *experience*” (p.482, emphasis added), and the fact that Bayne (2009) is concerned with the “admissible contents of perceptual *phenomenality*” (p. 386, emphasis added). Similarly, Logue (2013, p. 2) makes efforts to emphasise that the key question is what is actually experienced, twice characterising the core controversy as whether or not we “literally visually experience” high-level properties.

Therefore, it is perfectly reasonable to consider the debate to be primarily concerned with our question of what is experientially manifest, and perfectly reasonable to take those properties considered ‘high-level’ to occupy the grey-zone of seemings—seemings whose admission to the category of experientially manifest e-seemings is controversial. Having clarified that there is genuine controversy over whether or not high-level properties are experientially manifest, let us now see how

this controversy bears upon the matter of how phenomenal intentionality is ‘given’ in experience.

5.4 Phenomenal Intentionality in The High-Level / Low-Level Dispute

We have seen that there is an ongoing dispute over precisely what is experientially manifest to one, in phenomenally conscious perceptual experience. It is controversial as to whether low-level properties, such as *yellowness* and *curviness*, are all that is experientially manifest, or whether high-level properties, such as the natural kind property of *being a banana*, are also experientially manifest. I will now explain why this dispute is of central importance to the matter of how we take phenomenal intentionality to be ‘given’ in experience.

The dispute is important, because if high-level properties *are* experientially manifest, then their experiential manifestation would seem to deliver phenomenal intentionality. In other words, if high-level properties are experientially manifest, then the phenomenal intentionality of an experience *of* or *about* a banana would seem to be constituted by the experiential manifestation of the high-level property of *being a banana*.

This reasoning extends across the whole range of high-level properties we saw listed earlier by Logue (2013, p. 2). All of the high-level properties are properties of being one *sort of thing* or another, or of *belonging to one category* or another.⁸³ For example, an object may be the particular sort of natural kind we call a ‘banana’,

⁸³ Bayne and Montague (2011, pp. 17-18) use the term “categorical perception” for the act of perceiving something *as* a particular sort of thing.

and might also belong to the category of things we consider to be 'edible'. Alternatively, the object might be the sort of artefact which we call a 'sledgehammer'. It might belong to the category of things which we consider aesthetically 'ugly', and be the sort of thing which we are able to pick up in our hands and hit against things. It might also be the sole member in the category of things which caused the window to shatter. Similarly, a piece of text or speech may be the sort of text or speech we would consider meaningful, or as meaning something in particular. Further still, a person may perform an act of the sort we consider to be 'generous', and may be in a mental state of the sort we call being 'happy'.

For these categorisations to be experientially manifest is for a form of phenomenal intentionality to be experientially manifest. To experience something *as* a particular sort of thing is to have an experience *about* something of that particular sort, and hence to have an experience with *intentionality*. For example, if it is experientially manifest to one that an object is an edible banana, one thereby has an experience *about* an edible banana. And if it is experientially manifest to one that an object is an ugly sledgehammer which shattered the window, one thereby has an experience *about* an ugly sledgehammer which shattered the window. Furthermore, because such categorisations are supposed to be experientially manifest, or 'given' in phenomenally conscious experience, any intentionality they carry would be exhaustively constituted by the experientially manifest phenomenal conditions, and would therefore qualify as *phenomenal* intentionality. Therefore, if it is true that high-level properties are experientially manifest, then it is also true that the experiential manifestation of those high-level properties delivers phenomenal

intentionality. In other words, if it is true that high-level properties are 'given' in experience, then it is true that phenomenal intentionality is 'given' by means of those properties being 'given'. Conversely, if it is true that only low-level properties are experientially manifest, then phenomenal intentionality is either 'given' in some way which requires only low-level properties, or is not in fact 'given' at all.

In putting these options on the table, the dispute maps out some of the possible truths concerning how phenomenal intentionality is 'given' in experience, if it is given at all. Of course, possibly true claims are also possibly false, and so in mapping out these claims, the dispute provides an insight into the various ways we might be mistaken in how we take phenomenal intentionality to be 'given' in experience. As we saw in Chapter 3, we are liable to mistake mere i-seemings, which are not 'given' in experience, for e-seemings which *are* 'given'. We are also liable to make mistakes if we blindly trust that things e-seem as they introspectively i-seem to e-seem, when in fact the i-seeming may be misleading. The dispute over high-level and low-level properties, between phenomenological liberals and conservatives, shows some of the ways in which these mistakes may play out specifically in relation to how we take phenomenal intentionality to be 'given' in experience.

If high-level properties are 'given', then it won't merely be the case that one is inclined to believe that things have such properties, but it will be experientially manifest to one that things have such properties. That is, things won't merely i-seem to have such properties, but will e-seem to have them. However, in claiming that high-level properties are experientially manifest as e-seemings, one might be mistaking mere i-seemings for e-seemings—mistaking mere inclinations to believe

that something is a banana with experientially manifest *banananness*. Alternatively, one may be trusting that things e-seem as they i-seem to e-seem—trusting that there is experientially manifest *banananness* because introspection inclines one to believe that there is—when in fact that i-seeming is misleading. Equally, one could be making the same sort of mistakes in claiming that *only* low-level properties are experientially manifest as e-seemings. One could be mistaking genuine high-level e-seemings for mere i-seemings, or could be blindly trusting misleading introspective observations. Furthermore, one might be committing the same mistakes in claiming that phenomenal intentionality is ‘given’ in experience at all.

Given this potential for error, it would be risky to commit to the substantive claims about what is and is not ‘given’ in experience made by the phenomenological liberal and the phenomenological conservative, respectively. But as we will shortly see, careful analysis of how the dispute plays out reveals a point of implicit agreement between both sides, concerning the experiential manifestation of structural organisation and its relationship to the phenomenal intentionality of experiences. As such, there is a way of moving forward, regarding how we take phenomenal intentionality to be ‘given’ in experience, which minimises the risk of error: we can commit to what is widely agreed upon about what is ‘given’, whilst remaining neutral on the more contentious claims of the phenomenological liberals and conservatives.⁸⁴

⁸⁴ Such a strategy does not, of course, render us *immune* to the sorts of error we are prone to, but it does *minimise* the risk of error. Consider the widespread disagreement with regard to the colour of ‘The Dress’ in a photograph which became widely discussed in February 2015, as reported by the BBC (2015). Some claimed that

5.5 The Point of Implicit Agreement

The point of agreement implicit in the dispute between phenomenological liberals and conservatives can be discovered through careful analysis of how the dispute plays out. As Kriegel (2007) and Logue (2013) point out, the dispute often produces stalemates—situations in which the best arguments each side can give are ultimately inconclusive, leaving their opposing views equally plausible. As Logue (2013) says, “...we find ourselves with initially compelling but ultimately inconclusive arguments for and against [phenomenological liberalism]” (p. 8),⁸⁵ to the extent that we should be worried that we are faced with an “irresolvable impasse” (p. 7). Kriegel concurs that

“Phenomenological disputes have a way of leading to apparent deadlocks with remarkable immediacy. Disputants reach the footstomping stage of the dialectic more or less right after declaring their discordant positions.”

(Kriegel, 2007, p. 122)

However, if we follow through precisely what is required in order for the kind of arguments deployed to result in stalemates, we will see that, beneath the explicit

the dress was white and gold, whilst others claimed it was black and blue. Prior to the dress being revealed to be black and blue, one could have minimised one’s risk of making an erroneous claim by committing only to what was widely agreed upon—that the dress had two different colours.

⁸⁵ Logue (2013) in fact makes her remarks exclusively with respect to the thesis that “we can visually experience natural kind properties” (p. 3), as opposed to phenomenological liberalism more generally. However, the arguments deployed for and against the aforementioned thesis are of precisely the same sort as the arguments deployed for and against phenomenological liberalism in general. As such, Logue’s points can be extended to arguments against phenomenological liberalism in general.

disagreement between phenomenological liberals and conservatives, concerning whether or not high-level properties are experientially manifest, there is a substantive point of implicit agreement concerning the experiential manifestation of structural organisation and its relationship to the phenomenal intentionality of experiences. Logue (2013) draws attention to three sorts of argument which produce stalemates.⁸⁶ I will now show, for each sort of argument, that it only produces stalemates if disputants agree on the aforementioned point. I will first explain how each sort of argument leads to a stalemate, and then explain why this requires the substantive point of agreement.

5.5.1 Three Stalemate-Inducing Arguments

The first sort of argument is that from *epistemological role* (Logue, 2013, p. 3). This sort of argument begins with the premise that our phenomenally conscious experiences or e-seemings play a role in our acquisition of knowledge, by justifying many of our beliefs. For example, my belief that there is a banana in the fruit bowl could be justified by the phenomenally conscious experience I have when I look at the fruit bowl and it contains a banana. The liberal points out that my visual experience of a banana justifies my belief that there is a banana in the fruit bowl, and claims that it could only do this if the high-level natural kind property of *being a banana* is experientially manifest. But in response to this argument, the conservative

⁸⁶ Logue (2013) in fact distinguishes four different sorts of argument, but I take the fourth sort of argument she distinguishes—a sort of argument based on a ‘Twin Earth’ scenario—to simply be a variation on arguments from phenomenal contrast, considered below. As such, I will not attend to it separately here.

claims, equally plausibly, that my belief could be justified by an experience in which only the low-level properties of *yellowness* and *curviness* are experientially manifest, combined with my justified background belief that yellow, curvy things are bananas. For whatever epistemological function the liberal may assign to experientially manifest high-level properties, the conservative can explain how that function could be fulfilled without experientially manifest high-level properties. As such, arguments from *epistemological role* produce stalemates.

The second sort of argument is that from *phenomenal contrast* (Logue, 2013, pp. 4-5). Phenomenal contrast arguments identify a difference between the phenomenal character of two phenomenally conscious experiences, and then, by inference to the best explanation, conclude that the difference is due to the experiential manifestation of a high-level property. For example, the most widely discussed phenomenal contrast argument, put forward by Siegel (2006), identifies a difference between the phenomenal character of one's experiences of looking at pine trees *before* and *after* learning to recognise pine trees from sight. Siegel claims that the best explanation of this phenomenal contrast is that the natural kind property of *being a pine tree* is experientially manifest *after* one has learned to recognise pine trees, but not *before*.^{87,88} Arguments from *phenomenal contrast* lead

⁸⁷ Siegel (2006) does not express her argument in terms of what properties are experientially manifest, but in terms of what properties are 'represented'. Nonetheless, for the reasons given towards the end of §5.3, I consider Siegel to be primarily concerned with the question of what properties are experientially manifest, when she talks about what properties are 'represented'.

⁸⁸ I have enormously simplified Siegel's argument for concision. Nonetheless, my simplification does justice to the overall form of the argument, and the points I will make apply equally to the unabridged argument.

to stalemates, because for whatever contrast a phenomenological liberal may identify, and claim to be best explained by the experiential manifestation of a high-level property, the phenomenological conservative can give an equally plausible explanation exclusively in terms of low-level properties. For example, Price (2009, pp. 513-14 & p. 516) suggests that once one learns to recognise a high-level property, such as the natural kind property of *being a pine tree*, one's experience might phenomenally differ due to a change in one's pattern of attention. Such a change of attention might mean, for example, that whereas one's phenomenally conscious experience was previously dominated by the low-level properties of being brown and green in colour, and large in size, one's experience now gives increased prominence to the low-level properties distinctive of pine trees, such as the particular shape of its needles, and the particular texture of its bark, thereby making for a phenomenal difference. Given that liberals and conservatives can provide equally plausible explanations for any given phenomenal contrast, arguments from *phenomenal contrast* produce stalemates.

The third sort of argument are arguments from *illusion* (Logue, 2013, p. 6).⁸⁹ This argument, from the phenomenological conservative, attempts to mount a *reductio ad absurdum* against liberalism. The supposedly absurd consequence of liberalism is drawn out as follows. Consider being presented with two objects: a real banana and an exact visual replica of that banana made from plastic. One veridically experiences both objects' identical low-level properties, such as *being yellow* and

⁸⁹ For purposes of clarity, my presentation of the argument, below, makes some cosmetic alternations to the argument as presented by Logue (2013, p. 6).

being curvy. If liberalism is true, such that the high-level natural kind property of *being a banana* can be experientially manifest, then in visually perceiving the banana, one should experience it *as being a banana*. Given that the replica is an exact visual copy of the real banana, one should also experience it *as being a banana*, because it is visually indistinguishable from the real banana. However, this means that one's experience of the replica is illusory, because one experiences it to have a high-level property it does not in fact have. The conservative highlights this as an absurd consequence of liberalism. They claim that it is absurd to think that there is anything illusory about an experience which correctly presents an object's low-level properties. It is more reasonable, they maintain, to hold that one's *experience* of the banana and the replica are equally veridical, but that one is mistaken in one's belief, or inclination to believe, that they are both real bananas. In other words, the conservative thinks it is more reasonable to hold that the replica's seeming to be a banana is an i-seeming based on veridical e-seemings, rather than an illusory e-seeming.

As Logue (2013, p. 6) points out, this argument is unconvincing, because it turns entirely upon the conservative's intuitions regarding when it is reasonable or absurd to classify an experience as illusory. The liberal might not share these intuitions, and so could simply insist that it is entirely reasonable to classify one's experience of the replica as an illusory experience, wherein the property of *being a banana* is experientially manifest to one, and yet the object is not really a banana. As such, we have a stalemate. Furthermore, removing intuition from the equation does nothing to improve the situation. For without appeal to intuitions, any claim

regarding what it is reasonable or absurd to classify as an illusion simply begs the very question at issue, concerning what properties are experientially manifest. For one cannot adjudicate on what it is reasonable or absurd to classify as an illusion without assuming a position regarding what properties are and are not experientially manifest. As Siegel (2006, p. 483) suggests, whether an experience is illusory or not should be settled on the basis of what properties are experientially manifest in that experience, rather than the other way around. In sum, consideration of what experiences it is reasonable or absurd to classify as illusory does not settle the matter between conservatives and liberals, but instead results in a stalemate. The liberal's case for classifying one's experience of the replica banana as an illusory experience involving an experientially manifest high-level property is just as plausible as the conservative's case for classifying it as veridical experience involving only low-level properties.

5.5.2 The Agreement Required for the Stalemates

I will now explain why, in order for the sorts of argument just described to produce stalemates between the phenomenological liberal and conservative, the two must implicitly agree upon the experiential manifestation of structural organisation and its relationship to the phenomenal intentionality of experience.

As we just saw, each sort of argument, from *epistemological role*, *phenomenal contrast* and *illusion*, effectively identifies something which it is incumbent upon the liberal and conservative to explain. Each sort of argument leads to a stalemate, because the liberal and conservative can provide equally plausible explanations with

and without committing to the experiential manifestation of high-level properties, respectively. The point of implicit agreement between the liberal and conservative becomes apparent if we consider what they must each be committed to in order to uphold their respective, equally plausible explanations in each case.

All explanations provided by the liberal of course implicate the experiential manifestation of high-level properties. But these explanations only make sense provided high-level properties are experientially manifest *in a systematic manner*. Let me explain what I mean by this. The liberal claims that the experientially manifest high-level property of *being a banana* can play a role in one's acquisition of the knowledge that there is a banana in the fruit bowl, by justifying one's belief that there is a banana in the fruit bowl. But if this is the case, then the experiential manifestation of the high-level property of *being a banana*, in perceptual experience, must reliably correspond to the presence of a genuine banana. Only if the experiential manifestation of a high-level property reliably corresponded to the genuine instantiation of that property could the experiential manifestation of the property justify one's belief in its instantiation. Patently, if the experiential manifestation of a high-level property in perceptual experience did not reliably correspond to the genuine instantiation of that property, the former could not justify belief in the latter. If the experiential manifestation of the high-level property of *being a banana* occurred on only one in every ten occasions when a banana was present, and on all occasions when an orange or peach was present, the experiential manifestation of that property would patently not justify one's belief that a banana is present. Rather, in order to justify beliefs in the way the liberal envisages, high-level properties must

be experientially manifest in a *systematic* manner, such that their manifestation reliably corresponds to the instantiation of the relevant property.

The same goes with regard to the liberal's explanation of *phenomenal contrast* cases. If experiential manifestation of the high-level property of *being a pine tree* is supposed to explain the phenomenal difference between seeing a pine tree before and after learning to recognise pine trees, then it must be experientially manifest in the same *systematic* manner, such that its manifestation reliably corresponds to the presence of pine trees, only *after* one has learned to recognise them.

Even the liberal's explanation in response to arguments from *illusion* requires there to be some *systematicity* to the experiential manifestation of high-level properties. For it is only reasonable to suppose that one could have an illusory experience of a high-level property in circumstances such as those of the banana-replica scenario. It is reasonable to suppose that one might experience the replica *as being a banana* because it is visually indistinguishable from the real banana. As such, the liberal's explanation of such illusory experiences requires the experiential manifestation of high-level properties to systematically occur on occasions when an object is sufficiently indistinguishable from objects which genuinely instantiate the relevant property.

To recap, all three of the aforementioned sorts of argument result in stalemates, where the phenomenological liberal and conservative both have equally plausible explanations to offer. As we have just seen, the plausible explanations which the liberal can give, in terms of experientially manifest high-level properties,

only make sense provided that those high-level properties are experientially manifest in the *systematic* manner just described. To take this a step further, the liberal's explanations will only be entirely coherent provided that the *systematicity* of experientially manifest high-level properties can be explained. In other words, the liberal position will only be entirely coherent so long as it is committed to some explanation of the *systematicity* of experientially manifest high-level properties.

A straightforward explanation of this systematicity is that high-level properties only become experientially manifest on occasions when a certain *pattern* or *structure* obtains amongst the low-level properties that are experientially manifest. For example, the natural kind property of *being a banana* only becomes experientially manifest on those occasions when the low-level properties of *being yellow, being curvy, etc.*, are experientially manifest to one as properties of a single object. This explains how high-level properties such as *being a banana* or *being a pine tree* could manage to be experientially manifest *systematically*, so as to reliably correspond to the presence of genuine bananas and pine trees. Furthermore, if the obtaining of a certain *pattern* or *structure* amongst experientially manifest low-level properties is sufficient for prompting the experiential manifestation of a particular high-level property, this explains why an exact visual replica of a banana can produce an illusory experience of the replica as *being a banana*. For so long as one's visual experience of the replica instantiates the requisite *pattern* or *structure* amongst experientially manifest low-level properties (i.e. *yellowness* and *curviness* in the same object), the relevant high-level property will become experientially manifest.

Given that the liberal position will only be entirely coherent so long as it is committed to some explanation of the *systematicity* of experientially manifest high-level properties, and given that the straightforward explanation articulated above is available to them, I take the liberal to be implicitly committed to this explanation. And given that the liberal is committed to there being *patterns* or *structures* amongst experientially manifest low-level properties, I take them to be committed to the view that experience has an experientially manifest structural organisation. Furthermore, the liberal is implicitly committed to this experientially manifest structural organisation being implicated in the phenomenal intentionality of experience. As established earlier (§5.4), if it is true that high-level properties are experientially manifest, then it is also true that the experiential manifestation of those high-level properties delivers phenomenal intentionality. Therefore, if the liberal implicitly holds that high-level properties only deliver phenomenal intentionality in a systematic way by virtue of patterns or structures amongst low-level properties, then they implicitly hold that the structural organisation of our experiences plays an integral role in determining the phenomenal intentionality of our experiences.

Now, let us turn our attention to what the conservative must be committed to, in order to uphold their explanations in relation to the three arguments surveyed above. All explanations provided by the conservative of course implicate the experiential manifestation of exclusively low-level properties. But these explanations only make sense provided that instantiations of experientially manifest low-level properties can systematically instantiate particular *patterns* or *structures*. For example, according to the conservative, my belief that there is a banana in the fruit

bowl need not be justified by the experientially manifest high-level property of *being a banana*, but could be justified by the experiential manifestation of the low-level properties of *yellowness* and *curviness* together in one object, combined with my justified background belief that yellow, curvy objects are bananas. This explanation clearly requires that experientially manifest low-level properties can instantiate *patterns* or *structures*—such as *yellowness* and *curviness* together in one object— which are systematically instantiated so as to reliably correspond to the presence of a certain sort of thing, in this case a banana.

The instantiation of *patterns* or *structures* amongst experientially manifest low-level properties is also required in order for the conservative's denial of an illusory experience, in the banana replica case, to make sense. The conservative maintains that the experience itself is entirely veridical, but that one has the false belief or the inclination to falsely believe that the replica is a genuine banana. It only makes sense that the visually indistinguishable banana and replica would produce the same belief or inclination to believe in the presence of a genuine banana if, in being visually indistinguishable, both caused the instantiation of the same *pattern* or *structure* of experientially manifest low-level properties—a *pattern* or *structure* which one automatically takes to indicate the presence of a banana.

The same goes for the conservative's explanation of *phenomenal contrast* cases. According to the conservative, the phenomenal difference between looking at a pine tree before and after learning to recognise pine trees is due to a change in one's pattern of attention when one sees pine trees, perhaps to attend to the distinctive shape of their needles and texture of their bark. This patently requires that

experience is capable of systematically instantiating a particular *pattern* or *structure* of experientially manifest low-level properties, whenever a pine tree is seen *after* one learns to recognise them.

In sum, the explanations which the conservative gives, in relation to the three sorts of argument surveyed above, only make sense provided that experientially manifest low-level properties are capable of systematically instantiating *patterns* or *structures*. Just as with the liberal, given that the conservative is committed to there being *patterns* or *structures* amongst experientially manifest low-level properties, I take them to be committed to the view that experience has an experientially manifest structural organisation. Furthermore, just like the liberal, the conservative is implicitly committed to this experiential manifest structural organisation being implicated in the phenomenal intentionality of experience. As established earlier (§5.4), if it is true that only low-level properties are experientially manifest, then phenomenal intentionality is either 'given' in some way which requires only low-level properties, or is not in fact 'given' at all. Even if the conservative wishes to withhold judgment on whether or not phenomenal intentionality can really be said to be 'given' in experience, it is clear that a conservative's best explanation as to how phenomenal intentionality *could* be 'given' in experience, would be that it is 'given' by means of *structures* or *patterns* amongst experientially manifest low-level properties. Just like the liberal, the conservative implicitly holds that, in so far as phenomenal intentionality is 'given' in experience, the structural organisation of our experiences plays an integral role in determining the phenomenal intentionality of our experiences.

5.6 Conclusion

By looking at the dispute between phenomenological liberals and conservatives, over whether high-level properties are experientially manifest, we saw how we might make the sort of mistakes identified in Chapter 3 specifically in relation to how phenomenal intentionality is 'given' in experience. Having established the risk of error in committing to the contentious claims of either the liberal or the conservative, I proposed to show that there is a substantive point of implicit agreement between both sides, regarding how phenomenal intentionality is 'given' in experience, which we could commit to whilst remaining neutral on more contentious claims, thereby minimising the risk of error.

Through careful analysis of the stalemates which the dispute produces, and of the implicit commitments required on the part of both the conservative and the liberal for those stalemates to obtain, I have established that conservatives and liberals do indeed implicitly agree on a substantive point. They implicitly agree that experience has an experientially manifest structural organisation, and they implicitly agree that, in so far as phenomenal intentionality is 'given' in experience, the structural organisation of our experiences plays an integral role in determining the phenomenal intentionality of our experiences.

Therefore, if we wish to minimise the risk of error in how we take phenomenal intentionality to be 'given' in experience, and commit only to what conservatives and liberals agree upon, whilst remaining neutral on their more contentious claims, we cannot go far wrong in endorsing the view that phenomenal intentionality is constituted by the structural organisation of experience. In endorsing the structural

organisation account of the constitution of phenomenal intentionality, we commit to experience having an experientially manifest structural organisation, and we commit the view that the structural organisation of our experiences determines their phenomenal intentionality, but we can remain neutral on the more contentious issue of whether or not high-level properties are experientially manifest. In sum, we can move forward with the structural organisation account as a promising account of the constitution of phenomenal intentionality, given that it minimises the risk of error regarding how phenomenal intentionality is 'given' in experience.

PART III

Introduction to Part III

In Chapter 4, we reviewed some suggestions as to *precisely what phenomenal conditions* are constitutive of phenomenal intentionality. We identified that the most promising suggestion was that phenomenal intentionality is constituted by some sort of *structural organisation* of more basic phenomenal elements which are not themselves intentional. Then, in Chapter 5, we assessed the structural organisation account in light of the lessons of Part I, and concluded that it remains a promising account, given that it carries minimal risk of error about what is 'given' in experience. In this third part of the thesis, I will flesh out the structural organisation account, articulating my own more detailed theory as to how non-intentional elements of phenomenal character come to constitute phenomenal intentionality when structurally organised in the appropriate way.

Besides simply articulating a theory, I will also provide positive reasons for endorsing it. The idea that phenomenal intentionality is constituted by the structural organisation of more basic phenomenal elements emerged from the previous chapters as a promising candidate primarily because it remained standing where others fell. As such, my reasons for pursuing this idea are, as it stands, mainly negative. In this chapter, I aim to show that the idea is more than simply the best of a bad bunch, fleshing it out in a way which provides good positive reasons for pursuing it.

I will articulate my theory primarily in relation to perceptual experience, in part because the notion of a perceptual experience is more accessible and less controversial than notions of other sorts of experience, and in part because everything I need to say about phenomenal intentionality can be said in relation to perceptual experience, and then extended to other sorts of experience. In other words, all aspects of phenomenal intentionality which need explaining are present in perceptual experience, so explaining those aspects will suffice for explaining phenomenal intentionality across the board. These aspects, I maintain, are twofold.

First, there are *impressions*. *Impressions* are patterns in what we experience, which are distinctive of particular things (objects, properties, states-of-affairs, etc.). These are perhaps most easily recognisable in the case of visual perception, as ‘looks’—the distinctive ‘look’ of a domestic cat, for example, which the vast majority of actual domestic cats share in common, and would also share with an accurate wax model of a domestic cat. But we can, with a little more effort, also recognise impressions in other modalities. The ‘sound’ of a nightingale’s call, or the ‘sound’ of a combustion engine are examples of auditory impressions. Again, the ‘sound’ in question is some pattern in what we hear, shared by all nightingale calls or all combustion engines respectively, and which would also be shared by an undistorted audio recording of the call or engine. The ‘feeling’ of wind against one’s skin might also be considered an impression—perhaps a multimodal one incorporating the ‘sound’ of the wind in one’s ears. Impressions are undoubtedly part of the phenomenal character of our experience, for their entire nature consists in the fact that, like all experiences with phenomenal character, there is something it is like to

undergo them. As such, any intentionality or aboutness possessed by these impressions will of course count as *phenomenal* intentionality.

Second, there is *taking*. *Taking* is what occurs when one *recognises* something in experience *as* some particular thing, or *takes it to be* a particular thing. For example, when one sees a creature in the distance and correctly *recognises* it as a moose, or sees an oddly shaped rock in the mist and mistakenly *takes* it to be a moose. Or, for non-visual examples, when one hears a sentence in a language one understands, one *recognises* the language and *takes* the spoken words to mean something. More mundanely, one may hear a sound *as* the sound of a door closing, see the trees on the hill *as* pine trees, smell an aroma *as* the aroma of coffee, or feel a brush against one's leg *as* the brush of a cat's tail. Unlike *impressions*, it is not immediately clear that *taking* is part of the phenomenal character of our experience. But the fact that it is can be drawn out by revisiting the contrast-cases provided by Strawson (2008), which were touched upon in Chapter 4 (§4.4). These examples draw out the fact that *taking* is part of the phenomenal character of experience by identifying the contrast in phenomenal character between experiences which lack taking, and experiences which include it. As one example, Strawson describes the addition of *taking* to experience as "a bit like looking at one of those pictures where you can't see what it is a picture of, and then suddenly you see" (2008, p. 302). There seems to be a significant difference in the phenomenal character of the experience before and after one recognises what the picture is of, which we can identify as the contribution of *taking* to the experience. As another, Strawson considers a monoglot anglophone and a monoglot francophone listening to the French news. The

francophone, equipped with an understanding of French, will have an experience of considerably different phenomenal character from that of the anglophone, who doesn't understand a word. Again, we can identify the difference as the contribution of *taking* to the francophone's experience. As such, *taking* does appear to be part of the phenomenal character of our experience, and so any intentionality or aboutness possessed by *taking* will count as *phenomenal* intentionality.⁹⁰

I take these two aspects, *impressions* and *taking*, to cover everything which might reasonably be considered part of the phenomenal intentionality of perceptual experiences. My perceptual experience of my cat is *of* or *about* a cat in part because I have distinctly 'catty' impressions—I experience a catty look, and a catty sound if she meows, or a catty texture if I stroke her. And it is in part an experience *of* or *about* a cat because I *take* my cat *to be* a cat—I *recognise* my cat *as* a cat. There does not seem to be anything else in perceptual experience which could reasonably be considered to make an additional contribution to its phenomenal intentionality. In other words, experiencing *impressions* distinctive of X, and *taking* X to be X, seem to exhaust the ways in which a perceptual experience could *be about* X, in virtue of its phenomenal character.

⁹⁰ Acknowledging this does not compromise our neutral stance on the question of whether high-level properties are experientially manifest, which we adopted in Chapter 5, because it does not commit us to a view on whether or not *taking* qualifies as the manifestation of high-level properties. This question will remain open even once we have explained what phenomenal conditions constitute *taking*, somewhat similarly to the way in which the question of whether a Jaffa Cake qualifies as a cake or a biscuit remains open, even once we have established exactly what constitutes a Jaffa Cake. As such, I will remain entirely neutral on this issue throughout, and focus instead on the primary issue of what phenomenal conditions *constitute* taking.

I will not argue, up front, that explaining these two aspects will suffice for explaining the constitution of phenomenal intentionality in *all* cases of phenomenally intentional experience, including those which are non-perceptual. Rather, the proof will be in the pudding. I will focus my efforts primarily on accounting for *impressions* and *taking* in terms of the structural arrangement of basic phenomenal elements in cases of perceptual experience. Then, once the account is on the table, I will show that it can feasibly be extended beyond perceptual experience, to explain the constitution of phenomenal intentionality across the board.

In Chapters 6 and 7, I will show how *impressions* can be accounted for in terms of the structural arrangement of basic phenomenal elements. I maintain that basic phenomenal elements are structurally arranged so as to depict things, in much the same way that paints can be structurally arranged on a canvas so as to depict things. Why is this relevant to my account of impressions? Because I maintain that impressions are precisely those arrangements of basic phenomenal elements which depict things. It is by *depicting* things that impressions manage to be *about* things, and hence contribute to the phenomenal intentionality of an experience. And just as paints themselves do not depict, but can be structurally arranged into a painting which *does* depict, so I maintain that basic phenomenal elements are not themselves intentional, but can be structurally arranged into a phenomenally intentional impression. In Chapter 6, I will elucidate the notion of depiction through consideration of depictive artefacts, and thereby lay the groundwork for treating impressions as phenomenal depictions. In doing so, I will also identify a positive reason for pursuing my account of phenomenal intentionality. Then, In Chapter 7, I

will clarify the sense in which impressions are phenomenal depictions, before showing that it is plausible that the phenomenal character of experience has the capacity to instantiate such depictions. I will show that the brain can plausibly support a mental paint palette of the sort required to supply the basic phenomenal elements of all possible experiences, and then explain how these phenomenal building blocks are brought to instantiate the depictive structural arrangements which constitute *impressions*.

In Chapter 8, I will account for *taking* in terms of the higher-order structural organisation of *impressions* themselves. Then, I will show how the complete account of *impressions* and *taking* can feasibly be extended beyond perceptual experience, to explain the constitution of phenomenal intentionality in all types of phenomenally intentional experience.

6. Learning from Depictive Artefacts

In the next chapter, I will flesh out the idea that certain structures in phenomenal character are a sort of depiction in themselves. But first, in this chapter, I wish to elucidate the notion of depiction itself, through consideration of the depictive artefacts we encounter out in the world. Through consideration of these artefacts, I will show that depiction is a form of representation, or a way of being about things, where what the depiction represents or is about depends upon *experienced structure*. I will thereby elucidate the notion of depiction central to my account, whilst also identifying a positive reason for pursuing my account of phenomenal intentionality. For this account is perfectly equipped to explain the dependence of a depiction's *aboutness* upon *experienced structure*, in terms of the deeper dependence of phenomenal intentionality upon the structure of the basic, non-intentional phenomenal elements of experience itself.

6.1 What is Depiction?

Before we can engage in considerations about depictive artefacts, we must establish a rudimentary idea of what depiction is, and what qualifies as a depictive artefact. In other words, we must establish a conception of 'depiction' which identifies the phenomenon of interest. The conception will align with colloquial use of the word 'depiction' insofar as the word is commonly used to refer either to the phenomenon of an artefact *being about or representing* something, or to such an artefact itself. But the conception will not align with exactly what every person means whenever they use the word 'depiction', for the word is sometimes used to refer to

phenomena we are not interested in. Rather, the conception will identify a particular phenomenon which is often called ‘depiction’, and thereby single out the more specific phenomenon for which I will reserve the word ‘depiction’ as a technical term.⁹¹ As I will shortly make clear, I think that there are many artefacts which can be rightly considered depictions, in the technical sense, even though they are not pictures. Nonetheless, pictorial depictions are undoubtedly the most discussed of all artefacts which are depictive in the technical sense, and so it makes sense to start with pictorial depictions as a means of establishing the technical conception of depiction.

6.1.1 Pictorial Depiction

Pictorial depiction, in the technical sense of ‘depiction’, is a specific form of pictorial representation wherein pictures represent by presenting a visual image *of* the very thing that they represent. Photographs and portraits are paradigmatic examples of this. A photograph or portrait of President Trump—that is, a piece of paper covered with ink, or a canvas covered with paint—represents President Trump by presenting a visual image of him. Not all pictorial representation works in this way,

⁹¹ One may object that it is question-begging to put forward my own conception of ‘depiction’, and then draw conclusions about how depiction works. However, such an objection is misguided. For the conclusions I will draw are not assumed in my initial conception. Rather, the purpose of putting forward my own conception is to identify a genuine phenomenon which occurs in the world, by means of an initial rudimentary characterisation of that phenomenon. Then, with the phenomenon identified, labelled, and roughly characterised, I will be able to engage in more focused consideration of the phenomenon, and from there draw conclusions about how it works.

for a picture can represent something without presenting a visual image of it. For example, I might pictorially represent President Trump's inauguration by littering a canvas with chaotic smudges of blacks, blues and greys. I may be considered to have pictorially represented the inauguration in an abstract manner which might count as 'symbolic' representation. But I have certainly not *depicted* the event, for I have not painted a visual image *of* the event, as I would have had I visited the Capitol on the day and painted the visible scene before me.

This specific form of pictorial representation is also singled out by Peacocke (1987), Hopkins (1995) and Sedivy (1996).⁹² Hopkins (1995, p. 425) singles it out as "a form of representation that is distinctively exhibited by pictures and that it is distinctive of pictures to exhibit."⁹³ As such, one should be able to home in upon the

⁹² Although only Peacocke and Hopkins explicitly use the term 'depiction', with Sedivy instead using 'pictorial representation', it is clear that Sedivy is addressing the very same phenomenon, but classifying it in a different way. Sedivy (1996, p. 103) states that she is concerned with "the 'realistic' subset of images such as 'realistic' paintings, drawings and photographs". This set certainly only contains *depicting* artefacts according to my characterisation of depiction, for any painting, drawing or photograph must surely present a visual image of something if it is to qualify as 'realistic'. In other words, it seems that the criterion of 'realism' can only be met, in any of these forms, by presenting a visual image of something, or *as of* something, which looks sufficiently similar to the way that thing does or would look in reality. And given this, it seems that Sedivy's demand for realism can be factored out, because it is precisely by looking similar to the way something does or would look in reality that a visual image qualifies as a visual image of that thing. So, it is clear that Sedivy is only concerned with artefacts which represent things by presenting visual images of them, and hence is concerned with depiction as I characterise it.

⁹³ It is, of course, tautological that *pictorial* depiction will be a form of representation distinctive of pictures, but this does not mean that it is trivial to draw attention to the fact. Analogously, it is tautological that refugees are people just like you and I, but it is often not trivial to highlight the fact. Furthermore, it is worth noting that what we are considering 'pictorial depiction' is considered 'depiction' *simpliciter* to Hopkins, Peacocke and Sedivy (although not in those terms for Sedivy—see previous footnote). Therefore, from their perspective, the claim that depiction is a form of

phenomenon of depiction by distinguishing it from other forms of representation which can occur in pictures, but which are not *distinctive* of pictures, in the sense that they can also occur in representational artefacts which aren't pictures. Peacocke, Hopkins and Sedivy provide examples to help distinguish depiction in this way, all of which support the characterisation of depiction that I have outline above.

Hopkins and Peacocke provide examples which distinguish symbolic representation from depiction. Peacocke (1987, p. 383) points out that "a painting may depict a lamb, and the lamb may represent the man Christ; but it does not depict the man Christ." In other words, one can paint a picture of a lamb wherein the lamb symbolises Christ, without thereby painting a picture which depicts Christ. This is well explained by my characterisation of depiction: a painting depicts a lamb by presenting a visual image of a lamb, and that visual image can symbolise Christ without a visual image of Christ being presented, and hence without Christ being depicted.

Similarly to Peacocke, Hopkins (1995, p. 425) considers a painting which "shows a seated woman [who] symbolizes Despair." He shows that symbolic representation is not distinctive of pictures, because "a written description ... might describe a seated woman who symbolized Despair." Hopkins then homes in upon depiction by contrasting the similarities and differences of the forms of

representation distinctive of pictures does not have the appearance of a tautology. (Of course, in claiming that 'depiction' *simpliciter* is distinctive of pictures, Hopkins contradicts my earlier claim that there are many artefacts which can be rightly considered depictions even though they are not pictures. I will address this in the following subsection (§6.1.2) on 'Depiction in General'.)

representation at work in the written description and the painting: although they “represent Despair in the same way—by representing something that itself symbolizes Despair”, they “represent a seated woman in different ways”—by linguistic description in one case, and by depiction in the other. Again, my characterisation makes sense of this: the painting depicts a seated woman by presenting a visual image of a seated woman, but it does not depict Despair, and the written description does not depict Despair *or* a seated woman, because in none of these cases is a visual image of the thing represented presented.⁹⁴

Like Hopkins, Sedivy also highlights the distinction between linguistic description and depiction, pointing out that we have an intuitive appreciation of this distinction:

“I hold it as theoretically important that when we see a [depicting] picture, it seems to us that we are seeing things pretty much the way they are, as when we see an actual scene—and certainly not as when we encounter a description of a scene. My objective is to follow up on our intuitive appreciation of [depicting] pictures as being like the visible world and unlike linguistic descriptions.”⁹⁵

(Sedivy, 1996, p. 105)

⁹⁴ I take it that ‘Despair’ is not the sort of thing one could have a visual image of. One could have a visual image of a woman *experiencing* despair, but even in so far as this is true, the visual image is just of a woman with a certain expression. The expression might *communicate* despair, but the despair is not captured in the visual image in the same way that the woman and her expression are.

⁹⁵ I have modified this excerpt for clarity. Whereas I follow Peacocke and Hopkins in treating depiction as a specific form of pictorial representation, which potentially allows for the possibility of pictures which do not depict, Sedivy reserves the name “pictures” exclusively for pictorial artefacts which depict, and hence does not make efforts to clarify that she is always talking about depicting artefacts, whenever she is talking about ‘pictures’.

Sedivy also draws out our intuitive appreciation of the distinction between depiction and the form of representation by virtue of which a noun represents its referent, asking the question:

“Why does it seem obvious that the word ‘elephant’ stands in a purely arbitrary conventional relation to its object and why does it not seem obvious that that is also the relation between a picture of an elephant and its object?”

(Sedivy, 1996, p. 109)

Both of the intuitive distinctions Sedivy draws out are well-accommodated by my characterisation of depiction. Seeing a depiction of a scene is like seeing an actual scene precisely because the depiction presents a visual image of that scene, and it is unlike encountering a description of the scene because the description does not present us with any such visual image.⁹⁶ Similarly, it does not seem obvious that a picture of an elephant bears an arbitrary conventional relation to its object, an elephant, precisely because the relation is *not* an arbitrary conventional one. A picture *of* an elephant is a picture which *depicts* an elephant, and as such presents a visual image of an elephant.⁹⁷ Therefore, unlike the word ‘elephant’, the picture

⁹⁶ Even if we allow that a linguistic description can conjure a visual image of the scene in the subject’s imagination, it is still true that the description itself does not present a visual image of the scene, and as such is different from a depiction.

⁹⁷ I take it that when we talk about a picture ‘of’ something, we imply that the picture depicts what it is ‘of’, in the sense of depiction as I am characterising it. For example, ‘a picture of the inauguration’ intuitively describes a picture depicting the visible scene on the day of the inauguration, as opposed to an abstract painting of chaotic smudges intended to symbolise the event. Even if, at a push, the latter could be considered a picture ‘of’ the inauguration, the former certainly seems to be the default reading of the locution. This coheres well with Hopkins’ idea that depiction is a form of representation distinctive of pictures, for this would explain why the

doesn't represent an elephant simply because we uphold conventions which arbitrarily assign it this role. Rather, the picture represents an elephant because, by virtue of its very own intrinsic nature, and regardless of convention, it presents a visual image of an elephant.

6.1.2 Depiction in General

Having established a rudimentary conception of *pictorial* depiction, let us now see how we can expand it into a broader conception of depiction in general. We can begin to do so by reassessing the sense in which, as Hopkins claims, pictorial depiction is *distinctive of pictures*.

Because pictorial depiction is a form of representation which represents by presenting a visual image of the represented, it is thereby a form of representation supported only by those representational artefacts which we engage with through the sensory modality of vision. And undoubtedly, the paradigm of a *visual* representational artefact is a picture.⁹⁸ However, it seems that pictorial depiction is just one instance of a more general form of visual depiction, as it happens to be instantiated in the two-dimensional artefacts of pictures. For a three-dimensional

locution 'picture *of*' implies depiction by default. In any case, we are warranted in taking Sedivy's use of 'picture of' in this excerpt to imply depiction, because as I noted in fn. 95, above, Sedivy uses 'picture' to refer exclusively to depictions.

⁹⁸ One might claim that there is another type of visual representational artefact at least as paradigmatic as pictures, namely written language. However, closer consideration reveals that written language is constituted entirely out of many individual pictures, namely the glyphs (e.g. ⟨a⟩, ⟨a⟩, ⟨a⟩, ⟨a⟩) which represent the graphemes (e.g. ⟨a⟩, ⟨b⟩, ⟨c⟩, ⟨d⟩) which make up our words. As such, the example does not counter my claim.

sculpture can equally well represent something by presenting viewers with the ‘look’ of that thing, just as a picture, in presenting a visual image of something, thereby presents the ‘look’ of that thing. This is perhaps exemplified best by the wax models displayed in Madame Tussauds museums across the globe, which achieve a remarkable level of realism in ‘looking like’ the celebrities and notable figures they depict.⁹⁹ Crucially, a sculpture can only do this insofar as it is engaged with via the visual modality, for only then can it present a ‘look’. As such, pictorial depiction, and visual sculptural depiction, should both be considered instances of a more general form of representation, namely *visual depiction*, wherein artefacts engaged with via the visual modality depict things by presenting their ‘looks’.

From here, we can get a foothold on depiction as a more general form of representation. For if *visual* depiction consists in the presentation of *looks* of the depicted, then it stands to reason that there could be equivalent forms of depiction relating to the other sensory modalities—*auditory* depiction through the presentation of sounds, *tactile* depiction through the presentation of shapes and textures, *olfactory* depiction through the presentation of smells, and *gustatory* depiction through the presentation of tastes. We do not have to stretch our imaginations far to see that this reasoning holds up.

⁹⁹ It is testament to the models’ high level of realism that several celebrities have been able to surprise unsuspecting tourists by posing in place of their own wax model. See §3.1 of Wikipedia Contributors (2018). Photographs of the models can be found at (<https://www.madametussauds.com/london/en/whats-inside/world-leaders/>).

The recorded speech and music we commonly hear from speakers and headphones provides a panoply of auditory depictions, presenting us with sounds *of* things depicted, from the sound of a person's voice or the sound of a drum being hit, to the sound of a nightingale's call or the sound of waves crashing. And the sense in which these qualify as 'depictions' is not at all tenuous. Just as the paradigmatic visual depictions of photographs and portraits capture snapshots of the visual world, so do audio recordings capture snippets of the auditory world. And just as we *see* a person's face in a photo or portrait of them, so we *hear* a person's voice in a recording of them. Audio recordings are so closely analogous to visual depictions that it seems entirely warranted to consider them to also be depictions, but of the auditory kind.

We can apply the same reasoning to the other modalities. An accurately shaped and textured scale model of a building can be legitimately considered to be a tactile depiction of it, when engaged with through touch.¹⁰⁰ And a fragrance which recreates the scents of a particular substance or environment may legitimately be considered to depict it.¹⁰¹ In fact, a quick online search will reveal a number of

¹⁰⁰ Outside of the Scottish National Gallery in Edinburgh, at its north-eastern corner, one will find such a tactile depiction of Edinburgh itself. The three-dimensional bronze model of the city, according to its inscription, "was crafted to enable people with impaired vision to enjoy the grandeur of the city" (Canmore, n.d.). Photographs of the model can be found on the website of the sculptor, David Westby (<http://davidwestby.co.uk/commissions>).

¹⁰¹ A particularly good example of a real fragrance intended to depict in this way is that of 'Eau de Toast', which was launched by the UK's Federation of Bakers in 2013 "to encourage young women to think differently about bread" (Culliney, 2013), with samples distributed to models at catwalk shows during London Fashion Week. The fragrance, produced by The Aroma Company, was "developed using yeast top notes with a hint of caramel and malty base notes to give the overall impression of freshly toasted bread" (Ibid.).

perfume retailers and reviewers using the term “olfactory depiction” to describe just such products.¹⁰² Gustatory depictions are more difficult to identify, perhaps because we do not generally acquaint ourselves with the tastes of most things. (I know how bread and apples taste, but I have yet to sample the ink in my pen or the surface of my garden wall.) But environments with denizens whose tastes we *are* acquainted with can plausibly be depicted in a meal, as in the case of Heston Blumenthal’s dish ‘Damping Through the Boroughgroves’. Comprised of mushrooms, beetroot, blackberry, fig leaf, meadowsweet, melilot, oakmoss and black truffle, this dish contains a selection of strong earthy flavours, and may legitimately be considered a gustatory depiction of a forest or wooded environment.¹⁰³

Having enumerated various forms of representation which legitimately qualify as ‘depiction’, we now require a general conception of ‘depiction’ which unites them. So, what do the above forms of representation all have in common? In every case, the artefact represents something by presenting an *impression* of that thing, specific to the modality through which the artefact is engaged. Visual

¹⁰² Take, for example, this excerpt from JTD (2013) of Scenthurdle, on the fragrance ‘Traversée du Bosphore’ by L’Artisan Parfumeur, created by Perfumer Bertrand Duchaufour: “Turkish delight is flavored with classically aromatic ingredients such as rose water, orange blossom, citrus and mastic. It takes elements that the Western nose identifies as scents, not flavors, and makes food with them. Duchaufour steals them back and makes an olfactory depiction of the aromatic confection.” Similarly, Lo (2014) of Asia Spa describes the ‘Marrakech Intense’ fragrance by Aesop as “[a]n olfactory depiction of the city of Marrakech ... a woody oriental blend that captures the city’s colourful traditions and the evocative scents of spices used in Moroccan dishes.”

¹⁰³ The dish and its ingredients are listed by Sims (2016). Gen.u.ine.ness (2016) provides photographs and a description of the dish (<https://genuiness.wordpress.com/2016/03/19/the-fat-duck-bray-magical-heston/>).

depictions present looks, auditory depictions present sounds, tactile depictions present shapes and textures, olfactory depictions present smells, and gustatory depictions present tastes. So, in conclusion, our rudimentary conception of depiction is as follows. Depiction is a form of representation wherein an artefact represents something by presenting an impression of that thing, specific to the modality through which the artefact is engaged.

6.2 Depiction Depends Upon Experienced Structure

This characterisation of depiction we have established helps us to pinpoint the phenomenon in question, but it does not provide us with a clear understanding of how depiction *works*, which is what we will require in order to see how depiction depends upon *experienced structure*. Our characterisation leaves the functioning of depiction largely opaque, because it does not explain how a depiction manages to present an *impression* of something. So, let us now look more closely at what it means for a depiction to present an impression of something, to get a clearer understanding of how depiction is orchestrated. As before, I will focus first on pictorial depictions, given that they are by far the most discussed in the existing literature. I will first establish how pictorial depictions manage to present the looks of what they depict, before generalising out to all forms of depiction.

6.2.1 Dependence Upon Experienced Structure in Pictorial Depiction

It is tempting to say that a depiction presents the look of something by itself *resembling* that thing. And it is tempting to say that things resemble one-another by

virtue of being similar to one-another in some respect. So, giving into temptation in order to get the ball rolling, we might say that a depiction depicts something by bearing similarities to it. However, this account quickly reveals its slapdash assembly by casting its net for depictions far too wide. For not just *any* similarities between objects are relevant to successful depiction. You and I share about 14,000 genes with the acorn worm.¹⁰⁴ This is certainly a similarity, even one which might be considered a genetic *resemblance*. But fortunately, the similarity doesn't render you or I a walking depiction of an acorn worm.

We can home in upon the similarities that *are* relevant to successful depiction by applying similar considerations to artefacts which are paradigm instances of depiction. Consider two printed photographs—one of the Brooklyn Bridge, and one of my cat. The photographs are both printed at a size of four by six inches, using the same inkjet printer, onto the same photographic paper with the same glossy finish. As such, the printed photographs resemble one another very closely as physical objects. And yet, these photographs do not depict one-another. Rather, one depicts the Brooklyn Bridge, and one depicts my cat. So the sort of physical similarities which the photographs bear to one-another cannot be the sort of similarities in virtue of which each photograph successfully depicts what it does. The irrelevance of such physical similarities is even clearer in reverse: one photograph depicts the Brooklyn Bridge without being made of limestone, granite and steel, and the other depicts my cat without being made of flesh, bone and fur.

¹⁰⁴ Simakov et al. (2015, p. 460).

Sedivy frames the situation as follows:

“The notion of natural similarity picks out a class so superabundant that it in fact constitutes a ubiquitous background of objective relations between objects. Consequently, similarities that we find salient must owe that salience to some additional factor other than the objective fact that they obtain.”

(Sedivy, 1996, p. 106)

In other words, there are lots of similarities which objects naturally bear to one-another, but only some of those similarities are responsible for one object counting as a depiction of another. Therefore, the similarities responsible aren't responsible because of the mere fact that they are similarities. Rather, there must be something about those particular similarities by virtue of which they are responsible for depiction. There must be some fact of the matter which 'singles out' the relevant similarities.¹⁰⁵

In the case of the photographs, the features by virtue of which each print is relevantly similar to what it depicts must be features which the two prints do not share in common. One print must have a certain set of features which makes it relevantly similar to the Brooklyn Bridge, and the other must have a different set of features which makes it relevantly similar to my cat. Of course, I have stipulated that the printed photographs are physically the same in almost every respect. The only substantive respect in which they still differ is in the *arrangement* or *distribution* of

¹⁰⁵ Hopkins (1995, p. 440) expresses the same concern: “It is not true, for instance, that we see a depiction of something as resembling it in the complex overall way in which identical twins are alike. If picture and object look alike they must do so in some more particular respect.”

ink on the surface of the paper—that is, in the placement of different colours of ink alongside one another on the two-dimensional plane of the paper’s surface. As such, the features which make each print relevantly similar to what it depicts must be bestowed upon it by the particular arrangement or distribution of ink on its surface. The same thought experiment can of course be applied to paintings and drawings. We can imagine two paintings identical in almost every way, but which depict different things by virtue of different arrangements of paint on the canvas. And we can imagine two drawings identical in almost every way, but which depict different things by virtue of different arrangements of lead or ink on paper.¹⁰⁶ So it would seem that the relevant similarities between a depiction and what it depicts—at least in the case of typical photographs, paintings and drawings—must be similarities between features which can be instantiated by the arrangement of materials on a two-dimensional plane.

Sedivy brings us even closer to ‘singling out’ the relevant similarities by characterising them as those we find *salient*. Features which aren’t salient to us—features which we cannot notice, recognise or experience—will not appear to us in our perceptual experience of the world, any more than entirely absent features would. Imagine if all features of the printed photographs were non-salient to us. The

¹⁰⁶ This is not to say, of course, that *any* difference in arrangement of ink, paint or lead in these examples will result in something different being depicted. For two different photos of the Brooklyn Bridge, or a drawing and a painting, can of course have very different arrangements of ink, paint or lead, and still depict the Brooklyn Bridge. In other words, I maintain that what is depicted *supervenes* upon the arrangement. This means that one cannot depict something different without changing the existing arrangement of ink, paint or lead, but it does not preclude one from depicting the same thing with a different arrangement.

photographs would not even appear in our experience as objects in their own right, let alone as objects which depict something else. Rather, they would be entirely invisible to us, just like radio waves or magnetic fields. Even if the photographs themselves were salient physical objects, they would fail to depict either the Brooklyn Bridge or my cat if we covered their printed surfaces with a layer of opaque black paint. For the photographs are relevantly similar to what they depict by virtue of the two-dimensional arrangement of ink on their surface, and all features of that arrangement are rendered non-salient by the black paint. As such, we can see that the similarities relevant to successful depiction must be similarities between *salient* features of the depiction and the depicted.

In the case of *visual* depiction specifically, the relevant similarities are of course between *visually salient* features of the depiction and the depicted—between features which we can visually experience. In other words, depictions depict by exhibiting features similar to those which we would visually experience if we were to actually look at the thing depicted. As Hopkins (1995, p. 429) puts it, “Whatever can be depicted could be seen.” The aforementioned photographs do not represent the weight of the Brooklyn Bridge, or the brain activity of my cat, precisely because these are not features we would visually experience when looking at the bridge or my cat. But they do capture the way in which the Brooklyn Bridge and my cat would appear to us, were we to look at them. To again quote Hopkins (p. 434): “this sort of representation [i.e. visual depiction] represents things via representing their appearances.”

So, visual depictions depict by exhibiting features similar to those which we would visually experience if we were to actually look at the thing depicted. In other words, visual depictions imitate the appearance of the thing depicted. Seeing the depiction is, at least in some respects, like seeing the depicted. This is precisely the sense in which a visual depiction presents the 'look' of what it depicts. But crucially, we can now elucidate what it means for a visual depiction to present a 'look' of the depicted in terms of similarities in visual experience. We can say that a depiction presents the 'look' of what it depicts by causing one to have a visual experience similar to the visual experience the depicted would have caused. Thinking in these terms allows us to ask: in what respect must the visual experiences be similar, in order for them to involve the same 'look'?

We know that the relevant similarities between a pictorial depiction and the depicted must be features which can be instantiated by the arrangement of materials on a two-dimensional plane. And given that the features instantiated must be *visually salient* ones, we know that the materials arranged to instantiate those features must be visible ones. Therefore, we know that a pictorial depiction's imitation of the appearance of the depicted, such that it presents the 'look' of the depicted, must be achieved through the arrangement of visible materials on a two-dimensional plane.

Peacocke and Hopkins concur with this line of thought, both highlighting two-dimensional *outline shape*, from some particular perspective, as a significant respect in which pictorial depictions are relevantly similar to what they depict.¹⁰⁷ Outline

¹⁰⁷ Peacocke (1987, pp. 384-88), and Hopkins, (1995, pp. 440-43). Concluding his remarks on p. 443, Hopkins writes: "the thought is that to see something O in some

shape is, of course, a feature of objects which can be rendered on a canvas or some paper using lead or ink, and which is automatically captured in an undistorted photograph—it is a feature which can be instantiated by the arrangement of visible materials on a plane. It is true that the photographs of the Brooklyn Bridge and my cat both capture the outline shapes of those objects, which I would see if I were to look at them from the perspectives at which the photographs were taken. And it is certainly plausible that paintings or drawings could depict the bridge or my cat by virtue of capturing, to a sufficient degree of accuracy, the same outline shape.¹⁰⁸

Of course, outline shape is not the only feature which can be rendered in ink, paint or lead on paper or canvas, and it is not the only feature which can contribute to successful pictorial depiction. Even if a painting heavily distorted the outline shape of the Brooklyn Bridge, the arrangement of composite shapes on the canvas, and their respective colours and textures, might still recreate enough visible characteristics of the Brooklyn Bridge to sufficiently reproduce its ‘look’, and hence qualify as a depiction of it. Such a case would not be so different from a photograph of my cat which is distorted enough to drastically change its outline shape, whilst leaving its body parts in the correct spatial arrangement and its fur visible, such that it is still has something of the ‘look’ which makes it recognizable as my cat. Although

part P of a surface S is to see P as resembling O in outline shape.” Although Peacocke only explicitly mentions “outline” once (p. 386), his focus upon regions of the visual field taken up by a shape, and his use of the example of a silhouette (pp. 385-86), indicate that outline shape is a fair characterisation of the feature he had in mind.

¹⁰⁸ Peacocke (1987, p. 386) concurs that the outline shape of a depiction need not be *precisely* the same as the outline shape of the depicted, stating that “no more than approximate congruence is necessary.”

these depictions do not preserve the outline shape of the depicted, they seem to depict by preserving some visually salient *structural feature* or other. The painting of the Brooklyn Bridge, for example might feature an arrangement of two vertical rectangles, in parallel, with a space between them traversed by a thinner horizontal rectangle, thus preserving the basic spatial arrangement of the bridge's towers and deck, when viewed side-on. Similarly, the distorted photograph of my cat preserves the spatial arrangement of its body parts. Depictions might also preserve structural features in terms of the relationships of similarity and difference between the colours and textures of different regions of the depiction. And it would seem that a pictorial artefact entirely absent of any such structural similarities would fail to reproduce the 'look' of the depicted, and hence fail to depict. For example, a canvas covered entirely in one textured colour, and hence lacking in any articulation of structure, would fail to depict the Brooklyn Bridge, even if the visible colour and texture accurately matched a visible part of the bridge. For even if this visible colour and texture accurately reproduces the 'look' of a part of the bridge, it patently fails to reproduce the 'look' of the bridge as a whole. So, it seems that *visually experienced structural features* are the sorts of features with respect to which a pictorial depiction must bear similarity to something, in order to present its 'look' and thereby depict it.

This makes a great deal of sense, for it explains how arrangements of visible materials on a two-dimensional plane are capable of depicting scenes which they are dissimilar to in potentially every respect other than their *visually experienced structural features*. For example, Sedivy (1996, p. 119) points out that a painting can depict a scene even if it "reflects light rays whose absolute values are very different

from those reflected within the three-dimensional layout [i.e. the depicted scene].” Despite this optical dissimilarity, we can explain how the painting nonetheless manages to present the ‘look’ of the scene in terms of it “preserv[ing] the same colour and brightness ratios” amongst the shafts of light available to the eye (p. 121), and hence preserving the visually experienced differences in the colour and brightness of the light reflected from different parts of the depicted scene—a visually experienced structural feature. To give another example, a depiction of railway tracks continuing away towards the horizon might feature two converging lines. As Hochberg (1984, p. 844) points out, the actual railway tracks do not in fact converge, and hence there is a significant dissimilarity between the actual tracks and the depiction of them. Nonetheless, the converging lines present the ‘look’ of the tracks because the lines have the same structural arrangement in our visual field as the actual tracks do, when we visually experience them.

Having previously established that depictions depict by presenting impressions of what they depict, we have now established more precisely how this works in the specific case of pictorial depiction. Pictorial depictions present the ‘look’ of what they depict by virtue of possessing *experienced structural features* similar to the depicted, of which outline shape is a paradigmatic example. What remains to be seen is that this point generalises, such that *all* forms of depiction manage to present impressions of what they depict by possessing similar *experienced structural features* to the depicted.

6.2.2 Dependence Upon Experienced Structure in Depiction in General

I aim to make the case that all forms of depiction manage to present impressions of what they depict by possessing similar *experienced structural features* to the depicted. Though it is unfeasible to prove such a generalisation beyond doubt, examples of each kind of depiction strongly support the generalisation, and thereby give us sufficient reason to endorse it.

Let us first expand our considerations minimally, from pictorial depiction to *visual* depiction. Earlier (§6.1.2), sculpture was used as an example of a kind of depicting artefact which is not a picture, but which can nonetheless support visual depiction. That is, a sculpture can, when engaged with through the sensory modality of vision, present the ‘look’ of some particular thing, and thereby depict that thing. The aforementioned wax models of Madame Tussauds are a clear example of this. As with the photographs of Brooklyn Bridge and my cat, the importance of preserving similarity in the *structural features* experienced can be drawn out if we consider what impact different changes to the depicting artefact will have on what is depicted.

Let us take a wax model of Queen Elizabeth II, for example. If we altered the spatial arrangement of the model’s limbs, switching the head and a leg, the model will have a ‘look’ significantly less like that of the Queen, and could not be considered to depict the Queen with the same degree of success. One could reasonably say that the model still depicts some variation on the individual that is the Queen, given that it still possesses a face and individual body parts which ‘look’ like those of the actual Queen, albeit a fictional Queen who has perhaps undergone some bizarre surgery. If we proceeded to further alter structural features, such as the spatial arrangement of

facial features, and the shape of the model's limbs, it is clear that the model would no longer reproduce anything like the 'look' of the actual Queen at all, and hence could not be considered to depict her. By contrast, if we were to take all colours used on the original model, and keep their relative brightness levels constant, whilst changing the hue of every colour to blue, the results would be different. It would seem that the model would reproduce enough of the 'look' of the Queen to qualify as a depiction of her, albeit a *blue* depiction. Because the similarity of so many structural features (the entire shape and spatial arrangement of the model, the relative proportions of its parts, the relative brightness levels of its surfaces) is preserved, the dissimilarity in colour between the actual Queen and the blue model does not prevent the model from sufficiently reproducing the 'look' of the Queen, and thereby depicting her. Altering the experienced structural features a degree further, though, by jettisoning the relative brightness levels of the surfaces, and making the entire model one shade of blue, would potentially change the 'look' of the model a degree too far. For we might no longer be able to easily distinguish the model's eyes from the surrounding skin, or distinguish the skin from the clothing. It would be much less easy to recognise this model *as* a model *of the Queen*, to the extent that it should be considered a poor depiction, if not an entirely unsuccessful one.

What these considerations reveal is that visual sculptural depiction depends upon *experienced structural features* in the same way that pictorial depiction does. And this gives us reason to believe that the dependence holds of *visual* depiction in

general. Having established this, let us now expand our considerations to other forms of depiction.

As an example of auditory depiction, let us consider an audio recording of David Cameron's resignation speech. So long as we can still make out the words spoken, speeding up or slowing down the audio, or shifting its pitch up or down, would not seem to alter the 'sound' significantly enough for it to no longer depict the speech. Even littering such changes in speed and pitch sporadically throughout would not seem to alter the recording's status as a depiction. However, if we were to alter the primary structural feature of the recording, namely the ordering of sounds in time, then this would without doubt have a damaging impact. Rearranging phonemes within words would entirely alter the 'sound' of those words, and rearranging words would alter the 'sound' of sentences and phrases. A rearrangement of the words into new grammatical sentences might depict a fictional speech, but certainly would not depict the original speech.¹⁰⁹ Given these considerations, it seems that auditory depictions also manage to reproduce the 'sound' of what they depict by preserving similarities to what is depicted in terms of the *structural features* experienced. As such, the dependence of depiction on *experienced structural features* clearly extends to auditory depiction.

To briefly extend the same reasoning to tactile depiction, let us reconsider the accurately textured scale model of a building. Altering the scale of the model will

¹⁰⁹ See, for example, Casetteboy (2014)—a video in which words from a number of Cameron's speeches have been rearranged to produce a rap (<https://youtu.be/OYBumQHPAeU>).

not affect what it depicts, and even a complete change in the texture of the model is not out of the question, so long as the model retains the shape of the building it depicts. For example, a palm-sized plastic scale model of a church with a steeple could be considered a tactile depiction of the church, even though the walls of the model feel smooth, and the walls of the church feel rough. However, altering the spatial structure or shape of the church would without doubt change what it depicts. A model lacking a steeple, or with a steeple rearranged to sit horizontally at its side, would no longer constitute a tactile depiction of the church. So again, in the case of tactile depiction, depiction seems to depend upon *experienced structural features*.

In the cases of olfactory and gustatory depiction, the experienced structural features are perhaps harder to identify. Nonetheless, one clear example of an experienced structural feature is that of temporal concurrence or sequencing of scenting or tasting notes. For example, the tasting notes of beetroot, mushroom or blackberry alone do not depict a forest, but when experienced *alongside* the tasting notes of the other ingredients in Heston Blumenthal's aforementioned dish 'Damping Through the Boroughgroves', the experience comprises a collection of strong earthy flavours distinctive of a forest or wooded environment, and hence depicts such an environment. We alter this structural feature of concurrence or sequencing of certain scents or tastes if we separate them out individually, as just described, or if we alter some of the scents or tastes involved. If we add sherbet, chilli and fish sauce to Blumenthal's dish, then it would no longer seem to constitute a gustatory depiction of a forest. The same reasoning can be applied to fragrances which constitute olfactory depictions of environments or substances.

These considerations clearly show that the importance of similarity in *structural features* extends to depiction in general, but it remains to be emphasised that the relevant structural features are, just as in the case of pictorial depictions, *experienced* ones. To reiterate, any structural features *not* experienced are as good as anything invisible or undetectable to us—they cannot constitute a part of any impression presented to us, precisely because their absence from experience means that they present nothing whatsoever to us. The importance of exclusively *experienced* structural features is evident from the fact that artefacts are capable of depicting things which they are dissimilar to in potentially every structural respect other than certain *experienced* ones. Visual or auditory depictions of crashing waves need not share the molecular structure of water, and as such I can watch or listen to them without getting wet. A tactile depiction of a person, in the form of an accurately shaped and textured model, need not have the physiological or cellular structure of a real person. Equally, an olfactory depiction of toast need not have a physical structure capable of supporting butter, and a gustatory depiction of a forest need not replicate the aboveground biomass, canopy height or plant density of a forest. All that needs to be replicated by the depicting artefact is the *experience* of seeing or hearing waves crashing, of touching a person, of smelling toast, and of tasting things in the forest. And as we have seen from the considerations above, the features of the experience which it is essential to replicate are *structural* ones.

In sum, consideration of examples of various kinds of depiction strongly supports the generalisation that *all* forms of depiction manage to present impressions of what they depict by possessing similar *experienced structural features*

to the depicted. As such, we have good reason to believe that it is true, across the board, that what a depiction *represents* or *is about* depends upon experienced structural features. Or, to put it more simply, we should accept that *depiction depends upon experienced structure*.

6.3 Structured Experience as an Explanation of Depiction's Dependence on Experienced Structure

In §6.1, we established a rudimentary technical conception of depiction, on which depiction is a form of representation wherein an artefact represents something by presenting an impression of that thing—a pattern in what we experience which is distinctive of that thing. Then, in §6.2, we established that an artefact presents an impression of something, and thereby depicts it, by possessing experienced structural features similar to those possessed by the thing in question. In other words, a depictive artefact causes one to experience structural features similar to those one would experience upon perceiving the actual thing. Let us now see how these points can provide us with a positive reason for pursuing my account of phenomenal intentionality, according to which phenomenal intentionality is constituted by a structural arrangement of the basic phenomenal elements of experience.

Let us concentrate on the role that *aboutness* plays in the phenomenon of depiction. Depiction is a particular way of *being about* things. As such, depictive artefacts are, by definition, *about* things. In other words, depictive artefacts possess

a kind of *aboutness*. And given the points just established, we may say that the *aboutness* of a depictive artefact depends upon *experienced structure*.

It is worth here emphasizing a point which has been touched upon, but not explicitly stated. Because the aboutness of depiction depends upon *experienced structure*, it depends upon *experiences*. To see this, it is helpful to think of depiction as a process, involving several parts, which can either fail or succeed. As we have established, the process succeeds when an artefact causes someone to have an experience of structural features similar to those they would experience if they were to perceive the thing depicted. However, the process does not succeed where an artefact simply exists in isolation. Without the contribution of experience, it is no longer clear that the artefact *depicts* or *is about* anything—it is no longer clear that it carries the same aboutness. In other words, it looks like depictions cannot possess their distinctive form of aboutness without a contribution from experience. One might object that a depictive artefact in isolation *does* retain its aboutness, by virtue of the fact that it *would* cause the relevant sort of experience, were a person to engage with it in the relevant way. But far from being a substantive objection, this is more of a concession of the very point I am making: that the aboutness of a depiction is entirely beholden to a contribution from experience, whether that is the experience an artefact *does* cause, or the experience an artefact *would* cause.

So, the *aboutness* of depiction depends upon *experienced structure*, not just in the sense that the artefact must be such that it will cause experiences of certain structural features, but also in the sense that experiences (or merely *possible* experiences) of those structural features are required for the artefact to genuinely

succeed in *being about* what it depicts. The fact that this is so demands an explanation—an explanation which my account of phenomenal intentionality is perfectly equipped to provide.

According to my account, the sorts of *experienced structure* which depiction is dependent upon consist in *structured experience*. That is, to experience structures of the sort which depiction is dependent upon is to undergo experiences whose basic phenomenal elements are structurally arranged in a particular way. And this sort of structural arrangement of basic phenomenal elements, I maintain, is precisely what constitutes phenomenal intentionality.

This account provides us with a very clear explanation as to why the aboutness of depictions depends upon *experienced structure*, with regard to both the *experienced* component and the *structure* component. First, the aboutness of depictions is so beholden to a contribution from experience because it is *derived from* the phenomenal intentionality which originates in the intrinsic structure of experience. In other words, a depiction manages to be *derivatively* about something by causing an experience which, by virtue of its own intrinsic phenomenal-intentionality constituting structure, is *non-derivatively* about that thing. Second, the aboutness of depictions is dependent upon experienced *structure* precisely because the basic phenomenal elements of experience must be *structurally arranged* in a particular way so as to instantiate the phenomenal intentionality from which the aboutness of a depiction derives.

One might agree that the explanatory payoff here is significant, but disagree that we need to explain *experienced structure* in terms of *structured experience* in

order to secure the payoff.¹¹⁰ One might argue that we can take phenomenal intentionality to be constituted by those phenomenal conditions under which certain sorts of structure are experienced, without making the further claim that experiencing those sorts of structure consists in experience itself being structured. We can thereby secure the same explanatory payoffs, one might argue, as follows: we can say that the aboutness of depictions is beholden to experience because it is *derived from* experiences-of certain sorts of structure, and we can say that the aboutness of depictions is dependent upon experienced *structure* precisely because experiences must be experiences-of an appropriate sort of structure so as to instantiate the phenomenal intentionality from which the aboutness of depictions derives.¹¹¹ One might thereby contend that we can secure the same explanatory payoffs without having to commit ourselves to any claims about the intrinsic structure of experience itself.

However, there are two reasons for preferring my explanation, on which phenomenal intentionality is constituted by *structured experience*, over the objector's explanation, on which phenomenal intentionality is constituted by *experienced structure*. First, the objector's explanation is less complete. The phenomenal conditions constitutive of phenomenal intentionality are supposed to

¹¹⁰ To be clear, *experienced structure* refers to structure which is experienced, and *structured experience* refers to experience which is structured.

¹¹¹ I have hyphenated 'experiences-of' so as to more clearly communicate the sense of 'of' I hear have in mind. On the sense I have in mind, to have an 'experience of' a certain structure is to experience a certain structure. On the sense I do not have in mind, to have an 'experience of' a certain structure is to have an experience which itself has a certain structure.

be fulfilled when an appropriate sort of structure is experienced, but it is unclear exactly what these phenomenal conditions are. In other words, it is not clear exactly how the experiencing of an appropriate sort of structure translates into some set of conditions which phenomenal character can fulfil. So, the explanation attempts to secure explanatory payoffs by stating that phenomenal intentionality originates in experienced structure, but fails to follow through with an explanation of precisely how phenomenal intentionality so originates, leaving the explanation strikingly incomplete. There is, of course, an easy way for the objector to complete their explanation. They can take experiences-of the appropriate sort of structure to consist in experience itself being structured, and take phenomenal intentionality to originate in the latter. But in doing so, they concede entirely to my explanation.

Second, the objector's explanation is circular. For it attempts to explain the origin of intentionality in terms of experienced structure. That is, it attempts to explain the origins of intentionality in terms of experience *of* or *about* an appropriate sort of structure. It thereby presupposes a form of intentionality in its explanation of the constitution of phenomenal intentionality, and as such cannot be considered to have fully explained the origin of phenomenal intentionality, by locating a point at which this phenomenally-constituted form of intentionality arises from the non-intentional. By contrast, my explanation does not presuppose any form of intentionality in its account of the origins of phenomenal intentionality, for it maintains that phenomenal intentionality is constituted by the structural arrangement of the basic phenomenal elements of experience, *which are themselves non-intentional*. It thereby locates the precise point at which phenomenal

intentionality arises from the non-intentional. For these two reasons, the explanation of depiction offered by my account of phenomenal intentionality is preferable to the explanation offered by the objector.

In conclusion, a fact we have established concerning the phenomenon of depiction can be seen to provide us with a positive reason for pursuing my account of phenomenal intentionality, in the sense that the fact demands an explanation which my account of phenomenal intentionality is perfectly equipped to provide. The fact in question, which demands an explanation, is that the *aboutness* of a depiction depends upon *experienced structure*. According to my account of phenomenal intentionality, phenomenal intentionality is constituted by a structural arrangement of the basic phenomenal elements of experience, which are themselves non-intentional. And as we have just seen, this account is perfectly equipped to explain the dependence of a depiction's *aboutness* upon *experienced structure* in terms of the deeper dependence of phenomenal intentionality upon the structure of the basic, non-intentional phenomenal elements of experience itself. The expedient explanation offered by this account of intentionality thereby provides us with a strong, positive reason to pursue the account further.

6.4 Conclusion

In this chapter, I have sought to elucidate the notion of depiction central to my account of the constitution of phenomenal intentionality, through consideration of the depictive artefacts we encounter out in the world. First, in §6.1, we established a rudimentary technical conception of depiction, on which depiction is a form of

representation wherein an artefact represents something by presenting an impression of that thing—a pattern in what we experience which is distinctive of that thing. Then, in §6.2, we established that an artefact presents an impression of something, and thereby depicts it, by possessing experienced structural features similar to those possessed by the thing in question. As such, we established that depiction is a form of representation, or a way of being about things, where what the depiction represents or is about depends upon *experienced structure*. Finally, in §6.3, we identified a positive reason for pursuing my account of phenomenal intentionality, according to which phenomenal intentionality is constituted by a structural arrangement of the basic phenomenal elements of experience. We saw that my account is well equipped to provide an explanation as to why what a depiction is about depends upon *experienced structure*, in terms of a deeper dependence of phenomenal intentionality upon the structure of the basic phenomenal elements of experience itself. With these considerations under our belt, we can now proceed to see more precisely how depiction features in my account of the constitution of phenomenal intentionality.

7. Depicting with Mental Paint

7.1 Impressions as Phenomenal Depictions

As stated earlier, I maintain that the basic phenomenal elements of experience are structurally arranged so as to depict things, in much the same way that paints can be structurally arranged on a canvas so as to depict things. This is relevant to my account of *impressions*, because I maintain that impressions are precisely those arrangements of the basic phenomenal elements of experience which depict things. In other words, I maintain that impressions are *phenomenal depictions*—depictions built from the basic phenomenal elements of experience, arranged in the appropriate way. But given that we have so far considered a depiction to be an artefact which represents something by causing one to experience an impression of that thing, I am liable to invite confusion in taking impressions themselves to be depictions. So, to prevent any such confusion before we proceed, it is worth clarifying precisely what I mean in claiming that impressions are phenomenal depictions.

According to the explanation of depictive artefacts provided by my account of intentionality, an artefact depicts *X* by causing an experience whose basic phenomenal elements are structurally arranged so as to instantiate a phenomenally intentional experience which is non-derivatively *about X*. The experience is *non-derivatively* about *X* because the phenomenal intentionality it possesses, which makes it about *X*, originates in the intrinsic structure of the experience itself. The artefact does not possess its own original intentionality in the same way, but *derives* its aboutness from the original intentionality of the experience it causes. In other

words, the artefact is *derivatively* about *X* by virtue of the fact that it causes a phenomenally intentional experience which is non-derivatively about *X*.

Earlier, I characterised *impressions* as patterns in what we experience, which are distinctive of particular things (objects, properties, states-of-affairs, etc.) To have such an impression, according to my account, is to undergo an experience whose basic phenomenal elements are structurally arranged so as to instantiate a phenomenally intentional experience which is non-derivatively *about* the thing the impression is distinctive of. As such, an impression doesn't depict *X* in the same sense that a depictive artefact does, by *causing* an experience whose structure makes it non-derivatively about *X*. Rather, an impression depicts in the sense that it *constitutes* the very experience which is structured such that it is non-derivatively about *X*. What makes an impression a *depiction* nonetheless is the fact that it owes its non-derivative aboutness to the *structural arrangement* of basic phenomenal elements, analogously to paints on a canvas being structurally arranged so as to depict things. But of course, whereas paints on a canvas are only *derivatively* about what they depict by virtue of causing an appropriately structured experience, impressions are *non-derivatively* about what they depict, because they *are* an appropriately structured experience. Impressions are, therefore, the source of all depiction, in the sense that depictive artefacts cannot depict without them.

One would be forgiven for harbouring suspicions of circularity here. For I am claiming that impressions should be considered depictions because their phenomenal intentionality depends upon the structure of basic phenomenal elements, analogously to the way in which the aboutness of paradigmatic depictions,

such as paintings or photographs, depends upon the structure of paints on a canvas, or ink on paper. And yet I am also claiming that depictive artefacts derive their aboutness, and hence their status as depictions, from the phenomenal intentionality of impressions, and inherit the dependence of their aboutness upon their structure from the deeper dependence of an impression's phenomenal intentionality upon the structure of basic phenomenal elements. In other words, I am claiming that impressions should be considered depictions because they work like depictive artefacts, but also claiming that depictive artefacts work the way they do entirely because of the way in which impressions work. This undoubtedly looks circular at first glance, but a closer review of the dialectic reveals this not to be so.

To borrow Wittgenstein's famous metaphor,¹¹² we have used examples of depicting artefacts, and observations of the dependence of their aboutness upon their own experienced structure, as a ladder up to a more informed position. From this position, we can see how the ladder somewhat obscured the truth, and justifiably kick it away. Although climbing the ladder involved treating depiction as a form of representation which is primarily instantiated in depictive artefacts, we are now pursuing an account according to which depiction is in fact primarily instantiated in impressions. We have effectively developed a hypothesis about the covert representational mechanics of our experiences, from consideration of the tangible representational objects we encounter in our day-to-day lives. We have used

¹¹² The metaphor can be found in proposition 6.54 of Wittgenstein's *Tractatus Logico-Philosophicus* (1961, p. 151). Although, unlike Wittgenstein, I do not regard the propositions of my ladder to be "nonsensical", I am nonetheless claiming that we must "transcend these propositions" or "climb beyond them".

depictive artefacts as an entry point to considerations regarding the phenomenon of depiction, and those considerations have given us reason to think that the source of this phenomenon might be the structural arrangement of basic phenomenal elements which my account takes to be constitutive of phenomenally intentional impressions. If this is indeed the case, then depictive artefacts manage to depict things precisely by exploiting the way in which phenomenally intentional impressions work: artefacts depict *X* by causing experiences structured in such a way that they are phenomenally intentionally about *X*. In other words, if my account is correct, then the way in which depictive artefacts work is parasitic upon the way in which phenomenally intentional impressions work.¹¹³ This is the sense in which impressions may be non-circularly regarded as *phenomenal depictions*. Impressions are taken to be the root of all depiction, and depictive artefacts are taken to dance to their tune. There is no danger of circularity, because the source of depiction is located firmly in impressions. Nonetheless, the way in which depictive artefacts work is taken to provide an insight into the way in which impressions work, precisely because depictive artefacts are parasitic upon, or dance to the tune of, impressions.

¹¹³ However, this should not preclude us from using the way in which depictive artefacts depict by virtue of their own structure as a helpful analogy in explaining how impressions depict by virtue of their structure. Although I maintain that depictive artefacts are parasitic upon impressions, such that dependence-upon-structure in impressions is explanatorily prior to dependence-upon-structure in depictive artefacts, it is nonetheless the case that dependence-upon-structure in depictive artefacts is significantly easier to recognise and understand. This is unsurprising, given that such artefacts are encountered as tangible objects in our day-to-day lives, and given that the sort of structure on which their aboutness depends is often *physical spatial* structure, which we are naturally adept at thinking about.

In sum, the idea of impressions as *phenomenal depictions* makes sense and is not circular. But it remains to be seen whether it is plausible that such phenomenal depictions are genuinely instantiated. In the remainder of this chapter, I will show that it is plausible that the phenomenal character of experience has the capacity to instantiate structures of the sort I take to constitute the *impressions* we experience, namely: structures of intrinsically non-intentional basic phenomenal elements which manage to depict things. In the following section, I will show that the brain can plausibly support a ‘mental paint palette’ of the sort required to supply the basic phenomenal elements of all possible experiences. Then, in the section after, I will explain how these phenomenal building blocks are brought to instantiate the depictive structural arrangements which constitute *impressions*.

7.2 The Mental Paint Palette

7.2.1 Phenomenal Qualities as Exemplar Mental Paints

One of the two aspects of phenomenal intentionality which my theory needs to account for is that of *impressions*—the patterns in what we experience which are distinctive of particular things (objects, properties, states-of-affairs, etc.). As explained in the previous section, I maintain that impressions are *phenomenal depictions*—depictions built from the basic phenomenal elements of experience, arranged in the appropriate way. Similarly to the way in which a painting depicts by virtue of the structural arrangement of the paints on the canvas which make up the painting, I maintain that an impression depicts by virtue of the structural

arrangement of 'mental paints' which make it up.¹¹⁴ And just as paints are not themselves depictive, but can be structurally arranged into a depictive painting, mental paints are not themselves depictive or intentional, but may be structured into an experience which depicts, and thereby qualifies as a phenomenally intentional impression.¹¹⁵ This notion of mental paint needs elucidating. And so, I turn my attention to the question: what are these non-depictive mental paints? Or, more specifically, what are the non-intentional basic phenomenal elements which make-up experience?

The full breadth of what counts as a mental paint can only be made clear once more of my account is on the table. So I will begin to elucidate the notion with the most clear and accessible example of the sorts of things which constitute mental paints: phenomenal qualities.

As established in Chapter 2 (§2.1.1), phenomenal qualities are those elements of an experience which make up its overall phenomenal character, which can be distinguished from one another by means of the particular quality, character or feel that each contributes to the experience. For example, if you experience a fire engine driving by with its siren sounding, the red of the bodywork can be distinguished from the black of the tyres, precisely by virtue of the qualitative difference between the

¹¹⁴ I of course use the analogy of a painting with the qualifications of the previous section in mind. Despite the qualifications, I think the analogy is warranted, for the reasons given in the previous footnote.

¹¹⁵ The inspiration for considering the phenomenal character of experience to operate like paint, as 'mental paint', comes from Block (1996). However, the account I am articulating is ultimately quite different from that which Block endorses. Whereas Block's mental paint is representational or intentional in and of itself, my mental paint is not.

experienced redness and the experienced blackness. Furthermore, both of these can be distinguished from the sound of the siren, by virtue of the qualitative difference between an aural experience and a visual experience. We have identified at least three elements of the phenomenal character of the experience which contribute something qualitatively distinct from one-another, and hence we have identified three different phenomenal qualities.

Taking these three phenomenal qualities as exemplar mental paints prompts two thoughts. First, three mental paints make for a rather measly mental paint palette. Consider how impoverished (and nightmarish!) your experiences would be, if their phenomenal character was comprised exclusively of redness, blackness and siren-sound. If we are serious about the idea that mental paints are the basic phenomenal elements of experience, then the complete mental paint palette must be diverse enough to produce all of the experiences we could possibly have. Furthermore, it must be diverse enough to distinguish mental paints in a more fine-grained way than our examples do. For the bodywork of the passing fire engine will doubtless catch the light such that one experiences more than one shade of 'redness', and there are plausibly many more such distinguishable shades of redness which one could experience, each constituting a distinct mental paint in the palette.

Second, our three examples provide a taste of the sort of relationships which can obtain between mental paints. There are two different sorts of qualitative similarity and difference recognisable across our three examples. Redness and blackness are distinct qualities, but are nonetheless of a similar *visual* type, which is itself distinct from the *aural* type of quality to which the siren-sound belongs.

Furthermore, it seems that the sort of qualitative difference which exists between redness and blackness, as two different qualities of the same type, might come in varying degrees. One can imagine, for example, a phenomenal quality of light pinkness, which would be more qualitatively similar to redness than blackness is, and more different from blackness than redness is.

In sum, our mental paint palette must be diverse enough to support all possible experiences, and must accommodate the various types of relationship which can obtain between different paints. Continuing with phenomenal qualities as our exemplar mental paints, we can elucidate the idea of a mental paint palette, complete with the diversity and relational organisation required of it, by appeal to the notion of a quality space.

7.2.2 Quality Space Theory

Quality space theory, as expounded by Rosenthal (2015), provides a taxonomical framework for the qualitative character of our mental states. Taxonomising mental paints according to this framework furnishes a mental paint palette which is suitably diverse, and sufficiently accommodating of the various relationships which obtain between paints. To do so is to use Rosenthal's framework for a purpose it was not intended, given that mental paints are categorically *not* the sorts of entity Rosenthal intends to taxonomise. But my use of the framework remains entirely coherent, even though it clashes with Rosenthal's broader theoretical commitments. Setting these clashes to one side, Rosenthal can in fact be

seen to engage in a taxonomical project which is not only consistent with my account of the mental paint palette, but also, to a significant degree, complementary to it.

Recall that mental paints are supposed to be the basic elements of phenomenally conscious experience, or the basic elements of phenomenal character. As such, mental paints are essentially conscious entities, in the sense that they consist in feelings or sensations which are, by their intrinsic nature, consciously experienced, or constitutive of a conscious experience. The notion of a mental paint is, therefore, the notion of an entity which is inextricably both qualitative and consciousness-constituting. This is categorically not the sort of entity Rosenthal would endorse the existence of, for he is committed to a view on which consciousness is independent from the qualitative character of mental states.¹¹⁶ On this view, mental states have qualitative character regardless of whether they are conscious or not, and qualitative character only features in conscious experience when the state possessing that character becomes conscious. Because he takes qualitative character to be independent from consciousness, Rosenthal thinks that we should be able to account for qualitative character without relying upon reference to conscious experience. As such, Rosenthal's quality space theory aims to characterise the qualitative character of our mental states without relying upon any sort of introspection or first-person access which might implicate consciousness. Instead, it aims to characterise the various 'mental qualities' which make up the qualitative character of our mental states exclusively in terms of the role they play in perceptual discrimination, and

¹¹⁶ See Rosenthal (1991, 2010).

hence entirely on the basis of publicly observable and empirically testable behaviour. As such, mental qualities are considered to be “the mental properties in virtue of which an individual can perform perceptual discriminations” (2015, p.34).

What does this mean in practice? Consider presenting a subject with an array of tiles which are identical to one-another, apart from the fact that half reflect a certain set of wavelengths of the incoming light, and half reflect a different set of wavelengths. If the subject is able to successfully sort them into these respective groups by sight, the quality space theorist will conclude that two mental colour qualities have figured in that discrimination, and that the subject was able to perceptually discriminate the two sets of tiles by virtue of these two distinct mental qualities. For according to the quality space theorist, “an individual [able] to discriminate two perceptible stimuli ... must be able to be in two types of perceptual state, each corresponding to one of the two stimuli” (p. 38), which “differ relevantly in respect of their qualitative character” (p. 39), by virtue of instantiating distinct mental qualities.

If we wish to work out the minimum number of different mental colour qualities which must figure in the subject’s perceptual repertoire, we can test for what are known as ‘just noticeable differences’ or JNDs. We could, for example, gradually make the reflectance properties of the aforementioned two types of tiles increasingly similar, until the subject can no longer distinguish between them. Or alternatively, we could begin with two sets of identical tiles, and gradually make one set’s reflectance properties increasingly dissimilar, until the subject can distinguish between them. The point at which there is a ‘just noticeable difference’ for the

subject signals the point at which the difference between the reflectance properties of the tiles is large enough to bring two distinct mental qualities from the subject's repertoire into play—one for each set of tiles. Any smaller difference between the reflectance properties, and both sets of tiles bring into play the same mental quality, meaning that the subject is unable to distinguish between them. If we had the time to repeat this procedure for all possible reflectance properties in all possible lighting conditions, then we could use this method to work out the minimum number of different mental colour qualities which must figure in a subject's perceptual repertoire.

Crucially, besides individuating mental colour qualities, this method enables us to see how those qualities relate to one-another in terms of similarity and difference. Consider carrying out the procedure beginning with two sets of tiles which the subject identifies as all being the same shade of blue, and gradually altering the reflectance properties of one set until half of all the tiles are identified by the subject as red. The JNDs we record along the way will allow us to map out a spectrum of the subject's mental colour qualities, varying in hue, from shades of blue, through shades of purple, to shades of red.¹¹⁷ By mapping out a linear sequence of qualities, ordered

¹¹⁷ There are, in fact, many different ways one could gradually alter reflectance properties to change the perceived colour gradually from blue to red, without passing through purple. One could go in the opposite direction with respect to hue, passing through all of the other hues except purple. Or, one could make alterations with respect to brightness, going gradually from blue to white, and then make further alterations with respect to brightness in the direction of another hue, from white to red. But I take it that the most intuitive way of gradually altering a perceived colour from red to blue would be by changing its hue, in the direction that would pass through the colour purple, and so I opt for this most intuitive method in my example.

with respect to their hue, we thereby construct a one-dimensional quality space.¹¹⁸ Within this space, the relative distance between qualities will correspond to their degree of qualitative similarity and difference, such that a shade of blue will be more similar to a shade of purple than a shade of red. Of course, hue is only one of the ways in which we are capable of discriminating colours from one-another. For we can discriminate many different shades of the same hue by their different levels of brightness, and we can distinguish many different shades of the same hue and brightness by their different levels of saturation. As such, hue, brightness and saturation provide us with three different dimensions along which we can map out a three-dimensional space of colour qualities. As with the one-dimensional quality space of hue, distances between qualities within the space will correspond to their degrees of qualitative similarity and difference, but this time across three dimensions simultaneously. It is this use of space to model qualitative similarity and difference between mental qualities which gives quality space theory its name.

Of course, we have the capacity to perceptually distinguish plenty more than just colour properties, and quality space theory can in principle build quality spaces to accommodate every mental quality required for our full array of discriminatory capacities. Multiple quality spaces will be required, because not all perceptually discriminable stimuli differ from one-another along the same set of dimensions. For example, no sound can differ from any colour with respect to hue, brightness or

¹¹⁸ By 'linear', I simply mean that the qualities mapped out differ from one another continuously, in the same respect and in the same direction. I do not mean that the sequence must trace a *straight* line through quality space, as more dimensions are added.

saturation, and no colour can differ from any sound with respect to pitch, loudness and timbre. Rather, sounds and colours are discriminable from one another precisely because they are distinct types of qualities with distinct qualitative characteristics. As such, mental sound qualities require a distinct quality space defined by the dimensions of pitch, loudness and timbre. The mental qualities involved in olfaction also require a distinct quality space—one with between ten and seventeen dimensions, according to Young et al. (2014). And the same goes for any other distinct types of mental quality. As Rosenthal (2015, pp. 50-51) points out, we can use the methodology of testing for JNDs to establish whether one’s ability to perceptually discriminate between two stimuli requires mental qualities of different types, from distinct quality spaces. For if we cannot, by gradually adjusting one stimulus, create a chain of JNDs which leads to the other stimulus, then the mental qualities responsible for the discrimination of those stimuli must not have any common qualitative characteristics which would place them at different points along any single dimension of variation, as would be required for them to occupy one and the same quality space.¹¹⁹

¹¹⁹ In keeping with his overall aim of characterising the qualitative character of mental states without appeal to introspection or first-person access, Rosenthal (2015, pp. 50-52) advocates using this methodology for individuating quality spaces as a means of individuating the sensory modalities, such that distinct quality spaces correspond to distinct sensory modalities, thereby “providing a uniform, seamless treatment of mental qualities and the modalities” (p. 52). Rosenthal also considers some difficulties the methodology might face, including cases of partially overlapping quality spaces, and defends the methodology by providing ways of overcoming these difficulties (pp. 52-58).

In sum, Rosenthal's quality space theory provides a way of characterising the qualitative character of mental states without relying upon any sort of introspection or first-person access to that qualitative character, and hence without implicating consciousness. The theory begins with the reasonable hypothesis that perceptually discriminating between two stimuli involves occupying two different perceptual states, whose difference from one another is what enables us to make the perceptual discrimination. Extending this to our full capacity for perceptual discrimination, the hypothesis is that we can occupy an extensive array of perceptual states, all of which differ from one another in correspondence with the perceptual discriminations we are able to make. What Rosenthal calls 'mental qualities' are the properties in virtue of which these perceptual states are supposed to differ from one another, and quality space theory provides a methodology for taxonomising those qualities. By testing for just noticeable differences between stimuli, we can establish what mental qualities the subject must have in their perceptual repertoire. If we find a chain of just-noticeably-different stimuli which differ by degree in some particular respect, then we can map out a linear sequence or one-dimensional space of mental qualities corresponding to the discriminable stimuli. And if we find another respect in which we can alter the stimuli in this space, to create chains of just-noticeable-differences orthogonal to the existing chain, then we can map out a space of qualities which differ from one another in two independent respects, giving us a two-dimensional quality space. And for every independent respect in which a set of stimuli can just-noticeably-differ from one another, we can add another dimension to the quality space. As such, a quality space can be mapped out for every set of mental qualities

which differ from one-another with respect to a common set of qualitative characteristics, which are modelled as the dimensions of that space. Mental qualities differ from others in the same space by instantiating variations of the qualitative characteristics common to the space, and differ from mental qualities in other spaces by instantiating different qualitative characteristics altogether.

7.2.3 Drawing Upon Quality Space Theory

To see how we can draw upon quality space theory, it will be useful to distinguish two different things it has to offer. First, it offers the general taxonomical framework of quality spaces, which provides a powerful way of categorising qualitative entities. Not only is this framework capable of grouping qualitative entities into their distinct kinds, by sorting them into groups which share common qualitative characteristics, but it also organises the entities within those groups in a way which captures their various relationships of similarity and difference to one another, for every respect in which they can be similar or different. Second, quality space theory offers the hypothesis that there are ‘mental qualities’—properties in virtue of which perceptual states differ, and in virtue of which we can make perceptual discriminations—and provides a methodology for mapping those mental qualities into the quality space framework—the methodology of testing for JNDs between stimuli.

It is the first element of quality space theory—its general taxonomical framework—which we can apply to mental paints to furnish a mental paint palette of the sort we require: one with a suitably diverse variety of mental paints, and the

ability to accommodate the various relationships of similarity and difference which obtain between paints. The quality space framework certainly provides the diversity we require, for different kinds of mental paint (such as visual colours and olfactory scents) are accommodated in different quality spaces, and each quality space accommodates every variant of that kind of paint in every dimension of variation, meaning that every possible paint of each kind is accommodated. This way of organising mental paints also captures the various relationships of similarity and difference which obtain amongst them, for mental paints may differ from others in the same space, to the smallest possible degree or the largest possible degree, or anywhere in between, along every possible dimension of variation for that kind of paint. As such, two paints of the same kind can be more different or less similar in some respects, and more similar or less different in other respects. The framework therefore accommodates the fact that a shade of red and a shade of green, though different in hue, may be very similar, if not the same, with respect to their brightness or saturation. Furthermore, the framework's grouping of paints with a common set of qualitative characteristics into a single quality space captures the more general similarity-of-kind which such paints bear to one another, as exemplified by the general similarity amongst the colours we experience. Equally, the framework's segregation of groups with different common characteristics captures the more general difference-of-kind between paints with different qualitative characteristics, as exemplified by the general difference between visual colour experience and olfactory scent experience.

The second element of quality space theory—its hypothesis about the role of mental qualities in perceptual discrimination, and its JND-testing methodology for mapping those qualities into a quality space—provides the empirical grounding which Rosenthal requires for his particular purposes. Rosenthal intends to taxonomise the qualitative character of mental states without recourse to introspection or first-person access, and hence without implicating consciousness. As such, he turns to a publicly observable and empirically testable aspect of human behaviour which can plausibly provide insight into the qualitative character of our mental states: our capacity for perceptual discrimination. It is at this juncture that my project diverges from Rosenthal's. For Rosenthal is concerned with a quality space which maps out the differences amongst perceptual states that underpin our capacity for perceptual discrimination, whereas I am concerned with a quality space which maps out the basic phenomenal elements that make up our conscious experience.¹²⁰ Thus, the qualities of my quality space must be experientially manifest in the sense described in Chapter 3 (§3.2), whereas the 'mental qualities' of Rosenthal's quality space needn't be. In other words, Rosenthal is concerned with a Perceptual Discrimination Quality Space (PDQS), whereas I am concerned with an Experientially Manifest Quality Space (EMQS).

However, this is not to say that the two quality spaces are incompatible. In fact, setting aside Rosenthal's broader aversion to experientially manifest qualities or

¹²⁰ For ease of exposition, I will from here on use the singular term 'quality space' to refer to any particular taxonomy using the quality space framework, though any such taxonomy may of course group qualities into more than one quality space.

mental paints, the two are entirely compatible. If a PDQS is a repertoire of non-intrinsically-conscious mental qualities, and an EMQS is a mental paint palette containing the basic phenomenal elements which make up conscious experience, then there is absolutely nothing incoherent about a subject possessing both at once. The two quality spaces simply map out two different sorts of entity. In fact, the coexistence of two distinct quality spaces provides a coherent explanation as to how perception occurs unconsciously,¹²¹ and as to why our capacity for perceptual discrimination appears to exceed our capacity for conscious perceptual discrimination.¹²² We can say that unconscious perception makes use of the PDQS, but not the EMQS, and that the PDQS is more fine-grained than the EMQS, such that there are more discrimination-enabling mental qualities in the PDQS than there are mental paints available in the palette of the EMQS. So Rosenthal's notion of the PDQS is compatible with my notion of the EMQS to the extent that their coexistence can provide a coherent explanation of some of the empirical data.

Furthermore, the fact that Rosenthal's PDQS is empirically grounded makes my claim that there is an EMQS more plausible. For if there are good empirical grounds for taking the human brain to be capable of occupying an array of perceptual states which differ with respect to the qualities mapped out by the PDQS, then this

¹²¹ Rosenthal (2015, p. 36) states that "It is beyond dispute that perceiving occurs not only consciously, but without being conscious as well, in masked priming and other forms of subliminal perceiving."

¹²² Rosenthal (2015, p. 40) claims that "we can expect subjective awareness to discriminate in a less fine-grained way than discrimination that occurs without such awareness, since JND judgements are stable even when individuals take themselves to be guessing", and "somebody's guessing about something is a good indicator that the individual is not consciously aware of it."

lends plausibility to the idea that the brain is capable of occupying experiential states which differ with respect to qualities of a similar level of diversity and which bear similar relationships of similarity and difference to one-another, as mapped out by the EMQS. In other words, if we have good empirical grounds for thinking that the brain has an infrastructure capable of supporting the PDQS, then this makes it more plausible that the brain has an infrastructure capable of supporting the EMQS. Of course, to accept that the brain supports the PDQS, and infer that it supports the EMQS, is still to make a leap of epic proportions, for the EMQS requires phenomenal consciousness, and the fact that the brain supports the PDQS has no impact upon the plausibility of the brain instantiating phenomenal consciousness. But if we grant that the brain does, somehow, instantiate phenomenal consciousness, then our empirical grounds for thinking that the brain has a PDQS-supporting infrastructure undoubtedly make it more plausible that the brain has the infrastructure to support a similarly diverse and relationally organised selection of basic phenomenal elements.

To summarise, Rosenthal makes a reasonable, empirically grounded hypothesis about how our capacity for perceptual discrimination works. If his hypothesis is correct, then our perceptual states differ from one another by possessing properties, or 'mental qualities', as mapped out by the PDQS. And if this is the case, then the human brain must have the infrastructure to support the repertoire of mental qualities that is the PDQS. As a necessary implication of a reasonable, empirically grounded hypothesis, I take the claim that the human brain has a PDQS-supporting infrastructure to be a reasonable, empirically grounded

hypothesis in itself. And if we grant that the brain is capable of instantiating phenomenal consciousness, then the reasonableness and empirical grounding of the hypothesis that the brain has a PDQS-supporting infrastructure lends some degree of plausibility to the idea that the brain has an EMQS-supporting infrastructure. In this respect, Rosenthal's taxonomical project is not only consistent with my conception of the mental paint palette, but complementary to it.

To be perfectly clear, I am not claiming that my application of the quality space framework to mental paint is empirically grounded in the same way that Rosenthal's taxonomy of mental qualities is. In fact, I am putting forward no methodology whatsoever for 'discovering' the basic phenomenal elements of experience. Rather, I am merely advocating the quality space framework as a way of taxonomising the basic phenomenal elements of experience, or in other words fleshing out our notion of the mental paint palette. I have taken it for granted that our experiences involve a diverse variety of phenomenal qualities, and that those phenomenal qualities bear various relationships of similarity and difference to one another. And I have illustrated how the taxonomical framework of quality spaces can furnish a mental paint palette with a sufficiently diverse selection of paints, bearing the appropriate relationships of similarity and difference to one another, such that it can plausibly provide the basic phenomenal building blocks of all possible experiences. I have appealed to the empirical grounding of Rosenthal's taxonomy of mental qualities—the PDQS—purely to show that it lends plausibility to the idea that the brain could support a mental paint palette taxonomised according to the same general quality space framework.

7.3 Painting with Mental Paint

7.3.1 From Mental Paints to Impressions

To recap, one of the two aspects of phenomenal intentionality which my theory needs to account for is that of *impressions*—the patterns in what we experience which are distinctive of particular things (objects, properties, states-of-affairs, etc.). These impressions, I maintain, are *phenomenal depictions*—depictions built from the basic phenomenal elements of experience, structurally arranged in the appropriate way. I have suggested that we conceive of these basic phenomenal elements as ‘mental paints’, akin to paints which depict when appropriately arranged on a canvas, and I have sought to flesh out this notion by drawing upon quality space theory. We have seen that a mental paint palette taxonomised according to the quality space framework can plausibly provide the basic phenomenal building blocks of all possible experiences, and can plausibly be supported by the brain.

However, it remains to be seen how these phenomenal building blocks are built-up into phenomenally intentional experiences. Just as a palette of paints alone will not produce a painting, the availability of a diverse palette of mental paints is not itself sufficient for producing *impressions*. In the same way that a selection of paints from the palette must be structurally arranged upon a canvas for people to view, a selection of mental paints from the diverse mental paint palette must be brought to instantiate some sort of structural arrangement which is experientially manifest to the subject.

Of course, the quality spaces of the mental paint palette are, in a sense, ‘structural arrangements’ of mental paints. But the structural arrangement of a

quality space is not the sort of structural arrangement that is experientially manifest to the subject. Rather, the structural arrangement of a quality space maps out a selection of mental paints, and their relationships of similarity and difference, which each have the *potential* to be experientially manifest. For example, the mental paints of *redness* and *pinkness*, and the similarities and differences between them, will only be experientially manifest if both red and pink occur together in an experience. In sum, the ‘structural arrangement’ of a quality space is a structural element of experience which operates in the background—determining some of the possibilities for what can be experientially manifest, without itself being experientially manifest.

By contrast, the sort of structural arrangement of mental paints which is capable of depiction—of comprising *impressions*—must *itself* be experientially manifest. For *impressions* are patterns in what we experience, which are distinctive of particular things. We have a visual impression of a cat when we visually experience the distinctive ‘look’ of a cat, and we have an auditory impression of cat meowing when we aurally experience the distinctive ‘sound’ of a cat’s meowing. Such patterns occur wholly within phenomenally conscious experience—the entire pattern is ‘given’ or experientially manifest. And if, as I am claiming, these patterns consist in structural arrangements of mental paints, then those structural arrangements must be wholly ‘given’ or experientially manifest. The question we need to answer is: *what are these structural arrangements?* For mental paints to be in a ‘structural arrangement’ is surely for them to bear relationships to one-another as composite parts of some larger complex or whole, where the composition of the whole depends upon the relationships between the parts—between the paints. But if so, *what are*

the relationships between the mental paints, which determine the composition of the whole? That is, in the case of an experientially manifest structural arrangement of mental paints, *what are the experientially manifest relationships between the paints?*

We have already seen *some* of the experientially manifest relationships which mental paints can bear to one-another: they can bear the relationships of similarity and difference mapped out in the quality spaces of the mental paint palette. So, the question which remains is: what *other* experientially manifest relationships can mental paints bear to one-another, so as to instantiate the sort of experientially manifest structural arrangements which are capable of depiction—of comprising *impressions*?

In the case of an actual painting, the analogue of an experientially manifest structural arrangement is the arrangement of paint on the canvas. Just as *some* of the experientially manifest relationships which mental paints can bear to one another are those of the similarities and differences in the quality spaces of the mental paint palette, *some* of the relationships evident between the paints on the canvas, upon viewing the painting, are relationships of qualitative similarity and difference which obtain between the selection of paints of the palette. And with regard to the *other* experientially manifest relationships between mental paints, which we need to account for, the analogues in the case of the real painting are the *spatial* relationships between the paints, which obtain across the two-dimensional surface of the canvas. These spatial relationships are crucial for depiction, for they allow one to depict, for example, a smiling face, by arraying paints of appropriate colours in a spatial layout which has a similar 'look' to that of a smiling face. What,

then, are the analogous, experientially manifest relationships between mental paints, by virtue of which they can instantiate the depictive structural arrangements I am calling *impressions*?

In what follows, I will argue that the experientially manifest relationships in question are those that comprise the phenomenologically observable *spatiotemporal form* of perceptual experience. I will first try to establish a metaphysically neutral characterisation of this phenomenologically observable *spatiotemporal form*, and identify the sorts of experientially manifest relationships between mental paints which comprise this spatiotemporal form. Then, I will show how these relationships enable mental paints to instantiate the depictive structural arrangements I am calling *impressions*.

7.3.2 The Spatiotemporal Form of Perceptual Experience

7.3.2.1 Temporal Form

When I say that perceptual experience has a phenomenologically observable *spatiotemporal form*, I mean that phenomenological reflection reveals both *spatial* and *temporal* aspects to the manner in which perceptual experience is organised. The temporal aspect is captured well in the following passage from William James:

“In short, the practically cognized present is no knife-edge, but a saddle-back, with a certain breadth of its own on which we sit perched, and from which we look in two directions into time. The unit of composition of our perception of time is a *duration*, with a bow and a stern, as it were — a rearward- and a forward-looking end. It is only as parts of this *duration-block* that the relation of succession of one end to the other is perceived. We do not first feel one end and then feel the other after it, and from the perception of the succession infer

an interval of time between, but we seem to feel the interval of time as a whole, with its two ends embedded in it.”

(James, 1890, Vol. I, pp. 609–10, original emphasis)

Although James’ description is presented as a characterisation of how we perceive time, James is here more specifically focused on the manner in which *the present* is experientially manifest to us. And given that *all* perceptual experience is experience of *the present*, or of what is happening *now*,¹²³ we are warranted in taking James’ description as a characterisation of the manner in which things are experientially manifest generally, in perceptual experience. Trimming away the more poetic elements of James’ description, we can identify two key phenomenological observations concerning the temporal aspect of perceptual experience’s organisation.

First, what is experientially manifest to us always has its own experientially manifest duration. That is, anything which is experientially manifest subjectively feels to us to have duration or take up time, regardless of the duration of objective, public

¹²³ Hoerl (2018) characterises this as the idea that “it appears somehow essential to the nature of perceptual experience that it is the seeming disclosure of things as they are in the present” (p. 128). Hoerl takes issue with this idea that perceptual experience is “temporally viewpointed” (p. 147), arguing that we do not experience events *as present* (as opposed to past or future), but instead “simply experience their bare occurrence” (p. 148). Hoerl’s argument does not preclude us from taking James’ description as an accurate phenomenological observation of the temporal aspect to the manner-of-organisation of perceptual experience, because the experientially manifest ‘duration-block’ James describes could be cashed out entirely in terms of the ‘bare occurrence’ of events being experientially manifest to one, without requiring any present-tense temporal viewpoint to also be experientially manifest.

time throughout which it is experientially manifest to us.¹²⁴ Perceptual experience is itself dynamic and always unfolding, as opposed to a rigid succession of static moments. We do not experience successive frozen instances or snapshots, but rather a continuous stream of occurrences. As such, we hear a melody unfold. Rather than hearing a series of discrete notes, a temporally extended portion of melody is experientially manifest to us. Similarly, a falling ball is not visually experienced as suspended at a certain height, then at a lower height, and then a yet lower height. Rather, we experience the dynamic occurrence of the ball *falling*. It is in this sense that perceptual experience has the form of a temporal ‘block’ rather than a ‘slice’—a “saddle-back, with a certain breadth of its own”, rather than a “knife-edge”.

Second, experiencing things occurring or unfolding in this way involves experientially manifest relationships of precedence, succession and simultaneity between experientially manifest elements of the block.¹²⁵ The experiential

¹²⁴ Dainton (2017, §1) calls the sort of duration I am highlighting “phenomenal duration”, and distinguishes it, as I do, from “objective duration”. As Dainton puts it, the former is “the *felt* duration of [an experience]”, and the latter is “time as it is measured by clocks and stop-watches.” According to Dainton, “while it is not wrong to say that phenomenal duration is how long an experience *seems*, this should not be taken to mean that phenomenal duration is determined by how long a subject *judges* an experience to be.” We can express this point in terms of the distinction between two kinds of ‘seeming’ which was drawn in Chapter 3 (§3.3): phenomenal duration is how long an experience *e-seems*, rather than how long an experience *i-seems*.

¹²⁵ I here adopt the view argued for by Hoerl (2009), that “perceptual experience presents us with tenseless relations that obtain amongst events, such as before/after relations”, and that “events are simply presented, say, as unfolding in a certain order, rather than unfolding in that order in the present (as supposed to some other time) or unfolding in that order at a certain tenseless date” (p. 15). In his conclusion, Hoerl expresses this as the view that “perceptual experience, on the face of it, presents us with ... a world containing only temporal relations of precedence and simultaneity” (p. 17). Presumably, ‘precedence’ is here supposed to account for relations of both ‘precedence’ or ‘before’ *and* relations of ‘succession’ and ‘after’, given that they

manifestation of a temporally extended portion of melody, for example, involves experientially manifest relationships of precedence and succession between each note as experienced. Similarly, the perceptual experience of a child screaming *whilst* a ball falls involves an experientially manifest relationship of simultaneity between the scream as aurally experienced and the falling ball as visually experienced. The experiential manifestation of these temporal relationships is precisely what unifies different experientially manifest elements into the sort of experiential episode James calls a 'duration-block'.¹²⁶ Not only does the duration-block unify different experientially manifest elements from one modality, as in the case of the different

effectively express one and the same temporal relationship, just inverted. For clarity, I have opted to distinguish relationships of precedence, succession and simultaneity. Although, in the passage quoted, James (1890, Vol. I, pp. 609–10) mentions only the relationship of "succession", I think the addition of precedence and simultaneity is entirely in keeping with his phenomenological observations.

Independently of many of the persuasive points Hoerl makes, we have reason to adopt Hoerl's view just on the basis that it seems to accurately capture what we can phenomenologically observe to be 'given' or experientially manifest in perceptual experience. Consider visually experiencing looking down at the grass, then up at the sky. There appears to be e-seeming greenness, and e-seeming blueness, and an e-seeming ordering of the two occurring one after the other, but at no point does there appear to be any e-seeming 'presentness' or 'nowness' in addition to the experientially manifest phenomenal qualities and their experientially manifest ordering.

¹²⁶ Of course, this 'duration-block' is limited, in the sense that some events occur too quickly to be experientially manifest within the block, and some events occur too slowly. The movement involved in the growth of a rose bush, for example, is far too slow to feature as an experientially manifest occurrence within the duration-block. We only recognise that the bush has grown when we experience it as being larger than we remember. Conversely, the movement of a bullet fired through the air is too fast, such that all that is experientially manifest is, for example, the appearance of a bullet-hole in a target. See Dainton (2017, §2) for a discussion of these limits, and Hoerl (2009, §4, pp. 8-13) for consideration of how the empirically demonstrable limitations of our perceptual systems, at an information-processing level, relate to the phenomenology of temporal experience.

notes unified into a melody, but it also unifies experientially manifest elements from multiple modalities, as in the case of the scream and the falling ball unified as simultaneous events. That is to say, perceptual experience in *all* sensory modalities appears to be subsumed under one and the same temporal manner of organisation.

7.3.2.2 Spatial Form

Let us now turn to the *spatial* aspect of the manner in which perceptual experience is organised. Unlike the *temporal* aspect of perceptual experience's organisation, which unifies experiences from all sensory modalities into one experientially manifest duration-block, only perceptual experiences in certain sensory modalities have a *spatial* manner of organisation, and it is different for each modality.

The prime example of a spatial manner of organisation is the visual field. Michael Martin (1992, p. 198) attempts to give a metaphysically neutral characterisation of the visual field, starting from a point where "the visual field and visual space are taken to be features of the phenomenology of visual experience, aspects which can be identified independent of a commitment to any specific theory of perception." Martin goes on to list those organisational characteristics of the visual field which he thinks we can phenomenologically observe:

"Normal vision can afford us experience of more than one object simultaneously. Distinct objects are experienced as at distinct locations, and as spatially related to each other. There is also a sense in which the space within which the objects are experienced as located is itself a part of, or the form of, the experience. One is aware of the location of visual objects not only relative to other visually

experienced objects, but also to other regions of the spatial array - regions where nothing is experienced, but where something potentially could be.”

(Martin, 1992, p. 198)

Although Martin’s observations are mostly acceptable, there are two points which require amending. The first claim I take issue with is the claim that “the space within which the objects are experienced as located is itself a part of, or the form of, the experience.” The problem with this claim is that ‘the space within which the objects are experienced as located’ is ambiguous. The claim is only true on a specific conception of the space in question, which differs from the conception Martin employs. Martin claims that “sight ... is experience of objects external to one as arranged in physical space” (p. 210), and states that “we can think of ... visual space as simply a region of public space containing the objects currently seen” (p. 199). As such, it is clear that Martin conceives of ‘the space within which the objects are experienced as located’ as the three-dimensional, public physical space which we take ourselves and the objects around us to occupy, on our everyday conception of the world.¹²⁷ However, whilst it is true that we visually *perceive* objects as being located in three-dimensional space, it is not strictly true that this three-dimensional space is *experientially manifest* to us. For although we can visually perceive objects as having three dimensions, and as occupying a certain location in a three-dimensional environment, it is patently not the case that visual experience presents

¹²⁷ That is, the everyday conception of the world which Sellars (1963) calls “the manifest image”, which was touched upon in Chapter 5 (§5.1).

us with the entire three-dimensional structure of objects and their environments. The entire three-dimensional structure of a house is not simply 'given' to me in the experience I have when I look at it. Rather, I only experience the depth and orientation of its facing surfaces, from my perspective. As Jesse Prinz (2012, pp. 51-52) attests, "objects are always presented to us from a specific point of view ... We consciously experience a world of surfaces and shapes oriented in specific ways at various distances from us." Therefore, the space experientially manifest to us in visual experience can be thought of as two-and-a-half-dimensional, or 2.5D, in that it presents us with something akin to a two-dimensional array of coloured and textured shapes, but with the added sophistication of presenting those shapes as surfaces with depth and orientation.¹²⁸ As such, it is more accurate to characterise visual space as having two-and-a-half dimensions, rather than three. We can therefore uphold Martin's claim that "the space within which the objects are experienced as located is

¹²⁸ The idea that phenomenally conscious visual experience presents a 2.5D space was put forward by Ray Jackendoff (1987), who was drawing upon work by David Marr (1982). Marr held that visual object recognition involved three stages of mental representations: an initial 'primal sketch', followed by a more refined '2.5D sketch', and then a much more refined '3D model'. Jackendoff considers the question: at what stage, in the flow of information processing in the brain, does consciousness arise? Drawing upon Marr's work, Jackendoff argues that it is the intermediate level representation of the 2.5D sketch which corresponds to our conscious visual experience, as part of a broader argument that conscious experience more generally arises at the intermediate level of information processing. Prinz (2012, pp. 52-57) considers whether Jackendoff's proposal is problematic given that Marr's theory of visual object recognition has now been surpassed by more advanced theories, but ultimately concludes that "the prevailing conception of how vision works is broadly consistent with Marr on the points that matter for Jackendoff's conjecture" (Ibid., p.54).

itself a part of, or the form of, the experience” (1992, p. 198), so long as we conceive of that experientially manifest space as being 2.5D, rather than 3D.

The only other claim of Martin’s that I take issue with is his claim that, in visual experience, there are “regions of the spatial array ... where nothing is experienced, but where something potentially could be” (Ibid.). To elucidate, Martin uses the example of looking at a Polo mint—a ring shaped sweet with a hole in the middle—and claims that “nothing need be perceived to be within the hole. One is aware of the hole as a place where something potentially could be seen, not as where something is actually seen to be” (Ibid., p. 199). Martin’s example is itself entirely unproblematic. It describes one of many possible cases in which there is a location in visual space—at a certain location and depth in 2.5D space—which is not *currently* occupied by any object or surface, but which *could* be. However, a location unoccupied by an object or surface is not equivalent to a region wherein *nothing is experienced*. In fact, in the case of the Polo, experiencing some surface at a greater depth than the mint, surrounded by the ring-shaped surface of the mint at a shallower depth, is necessary for experiencing the absence of a continuing surface across the middle of the mint. There is not, as Martin initially seems to imply, a *hole in one’s experience*. Rather, one *experiences a hole*. As such, it seems misguided to consider any region of the visual spatial array to ever be lacking in something experienced. Rather, it would be more accurate to consider every part of the experientially manifest spatial array to be constituted by an experience of something or other.

In sum, visual experience has a spatial manner of organisation in that it presents a 2.5D spatial array of coloured and textured shapes, at various orientations and depths. In the same way that experientially manifest temporal relationships of precedence, succession and simultaneity unify different experientially manifest elements of experience into a duration-block, experientially manifest spatial relationships—to-the-right, to-the-left, below, above, further away, closer, etc.—unify different experientially manifest elements of visual experience—colour patches, edges, etc.—into a visual spatial array or field.

Although the visual field is the most tangible example of a spatial manner of organisation in perceptual experience, phenomenological observation also reveals spatial fields in auditory and tactile experience. We can see that experiences in each of these modalities replicate many of the spatial characteristics Martin (1992, p. 198) observes in the organisation of visual experience, in the passage quoted above.

If, in the vein of Mathew Nudds (2009), we consider the objects of auditory perception to be *sound sources*—the things which produce the sounds we hear,¹²⁹ then it is clear that we can have an auditory experience of more than one object simultaneously. When I try to listen to the radio over the sound of a jackhammer outside, I clearly have an auditory experience of two objects—the radio and the jackhammer—simultaneously. And although it is not as clear how many dimensions are experientially manifest to us in audition, distinct objects are nonetheless

¹²⁹ Some of what Nudds (2009, pp. 69-70) says suggests that he would consider *sounds* themselves to be ‘objects’ of auditory experience, in addition to *sound sources*, but this is of no serious consequence to my argument here.

experienced at different locations, and as spatially related to each other. For my auditory experience can present the sound of the radio as coming from close in front of me and slightly to my right, whilst simultaneously presenting the sound of the jackhammer as coming from far off behind me in the street. As such, auditory experience seems to present, at the very least, the rough direction and distance of objects in relation to us. In this respect, there is also a sense in which the space within which objects are experienced as located is itself a part of the experience, for the rough direction and distance of sound sources is experientially manifest alongside the pitch, loudness and timbre of the sounds they produce. Furthermore, auditory experience can also support experiences of the absence of objects, as in the visual case of the polo mint's hole, in the sense that we can experience in roughly what direction, and for roughly what distance, there is no object currently making a sound. If I stand in a quiet forest with my radio on the floor nearby, not only can I experience the radio close in front of me and to my right, but I can experience the absence of sound sources in every other direction. If a bird tweets in the distance behind me, I experience the absence of a sound source in the space between me and the bird. Given that it exhibits organisational traits extremely similar to those of the visual field, auditory experience can also be considered to have a phenomenologically observable spatial field, in which experientially manifest elements are unified by experientially manifest spatial relationships.

Let us now turn to the modality of touch. We can patently have a tactile experience of more than one object simultaneously, for one could, whilst blindfolded, experience touching a mug in one's hands whilst feeling a cat brush up

against one's leg. Also, the cat and the mug are clearly experienced as occupying different locations. Furthermore, there is a clear sense in which the space within which the cat and the mug are experienced as located is itself part of the experience, for when the cat brushes up against one's leg, and the mug is held in one's hand, one experiences them as located at different locations on the surface of one's body. Finally, tactile experience also supports experiences of the absence of objects, for one experiences the location of the cat and the mug on one's body in relation to parts of the body that are untouched, and hence effectively devoid of any external objects. As such, it is clear that tactile experience exhibits organisational characteristics sufficiently similar to vision to also count as having a phenomenologically observable spatial field,¹³⁰ in which experientially manifest elements are unified by experientially manifest spatial relationships.

7.3.3 The Formation of Impressions

My present aim is to explain what experientially manifest relationships obtain between mental paints, by virtue of which they can instantiate the depictive structural arrangements I am calling *impressions*. I am proposing that the

¹³⁰ My treatment of tactile experience here, in relation to visual experience, differs significantly from that of Martin (1992), who argues the converse: that tactual experience *does not* involve anything analogous to the visual field. However, it does not appear that Martin would disagree with the points I make concerning the organisational characteristics of tactual experience. Rather, it seems that we reach opposing conclusions because we employ different standards for what constitutes an analogue of the visual field. As such, Martin's argument is potentially entirely compatible with my argument here. It may be true that tactual experience does not involve anything analogous to the visual field according to his stricter definition, but does involve something analogous to the visual field according to my definition.

experientially manifest relationships in question are those that comprise the phenomenologically observable *spatiotemporal form* of perceptual experience. Having described this phenomenologically observable *spatiotemporal form*, and identified the sorts of experientially manifest spatial and temporal relationships which comprise it, I will now explain how *impressions* are formed from mental paints bearing those relationships to one another.

As a visual example, let us take the impression of the French Tricolour flag. What constitutes the impression is in part the relationships between the three colours of the flag. In practice, seeing the flag in different lighting conditions, and in conditions of imperfect and unequal illumination, will mean that one experiences multiple shades of blue, white and red, rather than three solid shades. Nonetheless, the impression distinctive of the Tricolour is in part constituted by the general relationships of similarity and difference between the three regions of the flag, as mapped out by the visual colour quality space. Besides this, the distinctive impression is constituted by the spatial relations between the mental colour qualities which are manifest. For example, the three distinctly coloured regions are vertical rectangles of equal size, aligned contiguously in a sequence of blue, white and red, from left to right. This of course describes the spatial arrangement of mental paints in something of an ideal situation, when a stationary flag is seen laid out flat, in the middle distance, directly before one. Nonetheless, I think the idea is clear: the distinctive impression of the Tricolour is comprised of mental paints instantiating a certain pattern of experientially manifest spatial relationships, and relationships of similarity and difference in colour. In practice, such a pattern of relationships between mental

paints, or structural organisation of mental paints, would need to be specified in such a way so as to flexibly accommodate the broad variety of particular instantiations of mental paints which we would want to consider distinctive of the Tricolour—regardless of what angle we were looking at the flag from, and regardless of whether it was laid out entirely flat, or blowing in the wind, etc. But this is a difficulty in practice, rather than in theory. In theory, there is some structural organisation of visual mental paints which is distinctive of the Tricolour, which is in part constitutive of an experience's *being about* the Tricolour.

Having just focused on the way in which experientially manifest spatial relationships are implicated in the constitution of impressions, let us now consider how experientially manifest temporal relationships might be implicated. A clear and simple example can be found in the distinctive sound of a doorbell. The impression distinctive of a particular doorbell's 'ding-dong' involves at least two distinctive auditory mental paints from the auditory quality space, and their experientially manifest relationships of similarity and difference. And crucially, it also involves the experientially manifest temporal relationships of precedence and succession between the auditory paints, as part of one's experientially manifest duration-block. The same basic account, albeit augmented with greater levels of complexity, can be given for the impressions distinctive of mobile ringtones, sections of music, and spoken words. An example of an auditory impression involving the experiential manifestation of *both* temporal *and* spatial relationships between mental paints, would be that of a motor vehicle passing close-by at high speed. In the duration block of one's experience, there would not only be experientially manifest temporal

relationships between auditory mental paints increasing and then decreasing in volume as the vehicle passes, but also experientially manifest auditory spatial relationships between locations where the vehicle sound respectively is and is not heard, which will also gradually alter over time in line with the change in volume as the vehicle passes. To reiterate, it is this pattern of relationships between mental paints, or structural organisation of mental paints, which constitutes the distinctive impression of a motor vehicle passing close-by at high speed, and is in part constitutive of an experience's *being about* a motor vehicle passing close-by at high speed.

I believe the examples given are clear enough to excuse me from rehearsing the same sort of account for impressions in every sensory modality. However, it is worth highlighting the fact that impressions needn't be confined to a single sensory modality. Because impressions consist in patterns of relationships between mental paints, and because perceptual experience in *all* sensory modalities is subsumed under one and the same temporal manner of organisation, impressions can include patterns in experientially manifest temporal relationships between mental paints *in distinct sensory modalities*. As such, there can be *multimodal* impressions. For example, if we add visual experience to the aforementioned example of the auditory impression of a passing motor vehicle, such that it includes patterns in the relationships between visual mental paints as the motor vehicle passes, and patterns in the relationships between those mental paints and the auditory mental paints (e.g., as the mental paints comprising the shape of the car become larger in the visual field, the mental paints comprising the sound of the engine rise in volume), then we

have a multimodal impression distinctive of a passing vehicle. Similarly, we may think of the overall experience of eating an onion bhaji as a multimodal impression comprised of the pattern of relationships between olfactory mental paints as one smells, gustatory mental paints as one tastes, and tactile mental paints as one feels the food in one's mouth.

7.4 Conclusion

One of the two aspects of phenomenal intentionality which my theory needs to account for is that of *impressions*—the patterns in what we experience which are distinctive of particular things (objects, properties, states-of-affairs, etc.). On the account I am proposing, those impressions are constituted of mental paints which are structurally arranged so as to depict things. Having seen that a mental paint palette taxonomised according to the quality space framework can plausibly provide the basic phenomenal building blocks of all possible experiences, we proceeded to consider how these phenomenal building blocks are built-up into phenomenally intentional experiences. In the previous section, I aimed to identify what experientially manifest relationships obtain between mental paints, in virtue of which they can instantiate the depictive structural arrangements I am calling *impressions*. I argued that the relationships are those which comprise the spatiotemporal form of perceptual experience. First, I gave a metaphysically neutral description of the phenomenologically observable spatial and temporal aspects of perceptual experience's manner of organisation, and identified the sorts of experientially manifest temporal and spatial relationships which comprise this spatiotemporal

manner of organisation. Then, I showed how these relationships enable mental paints to instantiate the depictive structural arrangements I am calling *impressions*.

8. Beyond Perceptual Impressions

In order to flesh out the structural organisation account of the constitution of phenomenal intentionality, I have broken phenomenal intentionality down into two aspects—*impressions* and *taking*—each of which needs to be accounted for in terms of the structural organisation of more basic phenomenal elements. Having provided such an account for *impressions*, I will, in the following section, turn my attention to *taking*. I will account for *taking* in terms of the higher-order structural organisation of *impressions* themselves. Then, I will show how the complete account of *impressions* and *taking* can feasibly be extended beyond perceptual experience, to explain the constitution of phenomenal intentionality in all types of phenomenally intentional experience.

8.1 Accounting for Taking

Taking, recall, is something which occurs as part of the phenomenal character of one's experience, when one *recognises* something in experience *as* some particular thing, or *takes it to be* a particular thing. For example, when one sees a creature in the distance and correctly *recognises* it as a moose, or sees an oddly shaped rock in the mist and mistakenly *takes* it to be a moose. The presence of *taking* in the phenomenal character of our experience can be drawn out using the sort of contrast-cases provided by Strawson (2008), which were touched upon in Chapter 4 (§4.4), and revisited in the introduction to this third part of the thesis.

With my account of *impressions* on the table, we can now recognise why an account of *taking* is required for a comprehensive account of the phenomenal

intentionality of perceptual experience. Let us see how far my account of *impressions* gets us in explaining what it is, about the phenomenal character of an experience, which makes it *about* a pine tree. According to my account, one aspect of the experience's *being about a pine tree* consists in the instantiation of a structural arrangement of mental paints, forming a pattern which is distinctive of a pine tree. In the case of visual experience, to put it crudely, the structural arrangement of mental paints can be considered analogous to the structural arrangement of ink on a photograph of a pine tree, but comprised of mental paint arranged in 2.5D visual space. But even if this pattern is distinctive of pine trees, in the sense that it is reliably instantiated, in the subject's visual experience, whenever they set eyes upon a pine tree or a realistic visual depiction of one, the instantiation of that pattern only seems to make the experience *about* a pine tree in a rather weak sense, leaving a much stronger element of *aboutness* unaccounted for.

Consider the two experiences compared in the phenomenal contrast argument put forward by Siegel (2006), which we touched upon in the previous chapter (§5.5.1). Both are experiences had by the subject when they look at pine trees, but one occurs *before* they learn to recognise pine trees by sight, and one *after*. There is a clear sense in which latter experience has an additional element of *being about pine trees* which the former experience lacks. However, it is plausible that both experiences involve the instantiation of the distinctive pine-tree-pattern in the visual field. As such, my account of *impressions* can only explain the relatively weak sense in which both experiences are about pine trees, leaving the latter experience's additional element of aboutness entirely unaccounted for. We require an account of

what it is, in the phenomenal character of experience, which changes when one *takes* the distinctive pine-tree-pattern *to be* a pine tree, or *recognises* it *as* a pine tree.

As such, we need to supplement our account of *impressions* with an account of *taking*, which is also in terms of the structural organisation of the basic phenomenal elements of experience, namely mental paints. Somewhat counterintuitively, the account I propose requires relatively little additional work, leaving most of the heavy-lifting, in terms building phenomenally intentional experiences out of mental paints, behind us in our account of impressions. On the account I propose, *taking* consists in the higher-order structural organisation of *impressions* themselves.

Recall that, in our discussion of the *spatiotemporal form* of perceptual experience, we saw that all perceptual experience seems to share the same temporal form, or be collectively unified in its temporal form. The same, I think, can be said for non-perceptual experience, such as day dreaming or imagining. When we engage with a thought experiment in a philosophy paper, for example, we experience a temporal sequence of first reading the paper's description of a scenario, then imagining the scenario, and then continuing to read. If the imagined scenario is so engaging that we end up distracted from the paper entirely, daydreaming about adjustments or modifications one could make to the original scenario, and then a noise snaps us out of the daydream, we experience first the day-dreaming, and then hearing the noise. This sort of temporal sequencing, which all of our phenomenally conscious experiences seem to be beholden to, is a respect in which *impressions* can be organised with respect to one-another. For example, if we read a sentence which

prompts us to imagine a pine tree, we might *first* experience impressions of letters and words in our visual field in the reading of the sentence, and *second* experience the visual impression of a pine-tree in our visual imagination. The temporal sequencing the impressions in one's experience may, of course, not be clear-cut—the instantiation of two impressions may well overlap or be concurrent—but there will nonetheless be some experientially manifest relationship of concurrence, preceding, or succeeding. It is with regard to their temporal sequencing that impressions can be structurally organised so as to constitute *taking*.

What sort of structural organisation, then, in terms of the temporal sequencing of *impressions*, constitutes the aspect of phenomenal intentionality we are calling *taking*? The sort of structural organisation I have in mind is that of the *coherence* of a temporal sequence of impressions. However, it is not the sort of coherence found in the continuity and constancies of perceptual experience, which we saw Farkas (2013) and Masrour (2013) draw attention to in Chapter 4 (§4.8).¹³¹ This sort of coherence is without doubt essential for the overall impression we have of an external, mind-independent world, which we might call a *macro-impression*. However, it is not the sort of coherence which could account for *taking*, for it is a coherence which could plausibly be possessed by the visual experiences of seeing pine trees both before and after one learns to recognise them by sight.

¹³¹ For example, impressions of objects get gradually larger in our visual field as we move towards them, and gradually smaller as we move away. As we do so, many of the structural relations amongst the mental paints which comprise an impression remain constant, whilst their size in our visual field changes in a smooth, continuous manner.

The sort of structural organisation I have in mind is extremely difficult to describe independently of appeals to the way in which our phenomenally conscious experience is *actually* structured, most of the time. Although this means that there remains a significant amount of work to be done in elucidating and fleshing out exactly what this structure consists in, I do not think this precludes us from appealing to it in our account of *taking*. It is a tangible enough feature of our day-to-day experience that it should suffice, for now, to define it by ostension.

The sort of structural organisation I have in mind consists in a cluster or series of systematically related impressions, perceptual or otherwise, which comprise what we might call 'a train of thought'. For example, a visual impression of a pine tree in perceptual experience, followed by a series of impressions in one's imagination, such as an auditory impression of the spoken words 'pine tree', a visual impression of pine needles, and a tactual impression of pine bark. One might even imagine what it would be like to climb the pine tree, which would consist in an entire episode of impressions, or indeed a multi-modal *macro-impression*, depicting the act of climbing a pine tree. This, I maintain, is the sort of structural organisation of impressions, in temporal sequence, which constitutes *taking*.

It is certainly plausible that this is what makes the difference between one's experiences of looking at pine trees before and after learning to recognise pine trees by sight. Although both experiences involve the visual impression distinctive of a pine tree, only the latter is succeeded by a cluster of impressions relating to pine trees. Not only would this clearly make for a substantive contrast in phenomenal character, it also constitutes a credible phenomenological account of conscious recognition, as

a process in which one's prior knowledge and experience (e.g., what pine trees are called in English, what their parts look and feel like, what it is like to climb them, etc.) is brought to bear on what one presently perceives.

The plausibility of the account also shows through in its provision of credible explanations of the contrast cases used to identify the contribution of *taking* to the phenomenal character of one's experience. Let us look in more detail at one of the cases from Strawson (2008), which was touched upon in Chapter 4 (§4.4):

“Suppose Louis is confronting a real scene in the world, and is seeing it, on account of some temporary mental fugue, just as an array of colours, without even any automatic taking of it as of the real world, without even any grasp of it as experience of anything at all. No intentionality here, I say, none at all. Then he comes to, he sees buildings, leaves, whatever it is.”

(Strawson, 2008, p. 301)

Although, as we saw in Chapter 4 (§4.4), Strawson himself considers *taking* to be a *sui generis* feature in phenomenally conscious experience, as opposed to a feature which could be broken down into a structural organisation of mental paints, Strawson here almost anticipates the sort of account I have provided. In attributing Louis' lack of *taking* to “some temporary mental fugue”, Strawson suggests that the absence of *taking* concurs with the absence of clarity and coherence in one's thought. It certainly makes intuitive sense that the ability to consciously recognise things would require clarity and coherence in thought, and my account provides a substantive explanation as to why this would be the case: what we consider to be clear and coherent thought consists in the sort of temporal sequencing of impressions I have described above,

and this temporal sequencing is what *constitutes* taking. Assuming that the same sort of 'array of colours' is experientially manifest in Louis' visual field both before and after the mental fugue, his recognition of buildings and leaves after coming-to can be explained in terms of those impressions already present in the visual field launching sequences or clusters of other systematically related impressions, in a way which they did not during the fugue.

My account of *taking* also provides *partial* credible explanations of the other phenomenal contrasts highlighted by Strawson, which must be supplemented by explanations drawn from our account of *impressions*. In effect, Strawson's other contrast cases identify phenomenal differences which aren't solely attributable to the temporal sequencings of impressions which constitute *taking*, but rather are partly attributable to the structural arrangements of mental paints which I am calling *impressions*. Strawson's example of a monoglot anglophone and a monoglot francophone listening to the French news, for instance, highlights that the francophone's experience, but not the anglophone's experience involves the "automatic and involuntary *taking* of sounds or marks *as* words and sentences that one understands and that represent something's being the case" (1994, pp. 7-8, Strawson's emphasis). Whilst the contrast between the experiences of the francophone and Anglophone can largely be explained in terms of the francophone alone experiencing a cluster of impressions systematically related to the words and sentences they hear, it is also plausible that the auditory impressions manifest to the francophone differ from those manifest to the anglophone in the first instance, due to differences in linguistic processing. In other words, it might be the case that the

structural arrangement of auditory mental paints differs between the francophone and anglophone, such that the francophone experiences word-impressions whilst the anglophone experiences mere sound or noise impressions. If we allow for this possibility, then a comprehensive account of the phenomenal contrast might need to appeal to differences in structural organisation *both* at the level of mental paints arranged so as to constitute impressions, *and* at the level of impressions temporally sequenced so as to constitute *taking*. The same applies to Strawson's example of "looking at one of those pictures where you can't see what it is a picture of, and then suddenly you see" (2008, p. 301). More concretely, consider seeing an image first merely as a collection of black spots in white space, and then coming to recognise part of the image *as* a Dalmatian. Whilst part of the phenomenal change can be attributed to having a cluster of subsequent impressions systematically related to the initial visual Dalmatian impression, it is also plausible that a change in one's attention to the image alters the structural arrangement of one's mental paints so as to instantiate the Dalmatian impression in the first instance. In sum, some of the instances of 'taking' which Strawson identifies are not exhaustively accounted for by the sort of higher-order structural organisation amongst impressions which I am calling *taking*, but must also be accounted for in terms of the lower-order structural organisation amongst mental paints which is capable of instantiating *impressions*. So, although the two examples just considered do not fit cleanly into that aspect of phenomenal intentionality which I am calling *taking*, they can nonetheless be accounted for entirely in terms of the structural organisation of mental paints.

8.2 Extending the Account Beyond Perception

Having accounted for both aspects of the phenomenal intentionality of perceptual experience—*impressions* and *taking*—in terms of the structural organisations of mental paints, I will now argue that this account can feasibly be extended to explain the constitution of phenomenal intentionality, not just in *perceptual* experience, but in all types of phenomenally intentional experience. Though it is beyond my present remit to explain how the account could be extended to *every* sort of conscious experience distinguished in the literature, I will explain how it can be extended to the sorts of extra-perceptual conscious experience most often discussed, and hopefully this will provide enough of an indication as to how the account could be extended to the sorts of experiences I do not explicitly discuss.¹³² I will first explain how the account can accommodate bodily, emotional and mood experiences, and then explain how it can accommodate cognitive experiences.

8.2.1 Bodily, Emotional and Mood Experiences

So far, I have characterised mental paints exclusively in terms of the phenomenal qualities associated with the five paradigmatic sensory modalities of sight, hearing, touch, smell and taste, but we undoubtedly experience phenomenal qualities which are not clearly associated with any of these senses. For example, bodily experiences such as nausea, orgasm, and pain would seem to involve

¹³² I will not, for example, explicitly address *agentive* conscious experiences—experiences of exercising one’s own will or acting voluntarily. Nor will I explicitly address conscious experiences of believing and desiring, or any of the other so-called ‘propositional attitudes’, independently of my discussion of cognitive experience.

phenomenal qualities beyond those associated with the paradigmatic senses. Some may make the same claim for emotional or mood experiences, such as experiences of elation, sadness, nervousness and anxiety. My account can easily extend to accommodate such phenomenal qualities, simply by allowing that the mental paint palette is not restricted to the sensory quality spaces which I have focused upon in articulating the account, but may in fact include a far broader array of mental paints.

With these additional mental paints included in the palette, the phenomenal intentionality of any experiences involving those paints is accounted for exactly as I have laid out above, in terms of the structural organisation of mental paints into *impressions*, and the higher-order organisation of impressions into *taking*. We have already seen, in the examples from Farkas (2013, p. 103) noted in Chapter 4 (§4.8), that feelings of pain could plausibly contribute to experiences *about* a ‘pain-patch’ on a leaf or a ‘headachy’ corridor, when appropriately structured alongside other basic phenomenal elements of experience. Furthermore, it is broadly agreed that our *actual* pain experiences are intentional *at least in so far as* the bodily location of the pain is experientially manifest to one, in the sense that this makes the experience *about* some occurrence, such as tissue damage, at that bodily location.¹³³ The structural organisation account can explain such location-oriented intentionality in terms of pain-constituting mental paints bearing experientially manifest spatial

¹³³ As Kind (2014, p. 122) points out, although there has been a long tradition of classifying pains—along with other bodily sensations, emotions and moods—as non-intentional, many philosophers have recently argued against this view, establishing that “insofar as pains carry information about the states of our body they can ... be understood as intentional.” See Tye (1995, pp. 111-116) and Crane (1998, pp. 232-238) for examples of such arguments.

relations to other bodily sensation mental paints or absences of them, within a sensory field mapping out one's body, such as that involved in tactile experience, as discussed in the previous chapter (§7.3.2.2). All bodily sensations can plausibly be accommodated in this sort of way.

How about emotions and moods? Is it plausible that emotional and mood experiences are simply experiences involving mental paints distinctive of emotions and moods, which become phenomenally intentional when structurally organised, alongside other mental paints, in the sort of way laid out above? It is first worth acknowledging that, as Kind (2014, p. 116) points out, "there is considerable disagreement about whether emotions and moods should be classified together or treated separately as different kinds of mental states." Though I wish to remain neutral on this issue, I am nonetheless going to treat emotions and moods in exactly the same way, rather than address them individually. This approach is justified for two reasons. First, emotions and moods appear to exhibit fundamental similarities in terms of their phenomenal character. Even if one wishes to insist that emotions and moods are distinct sorts of mental state, it is nonetheless the case that experiencing the emotion of anxiety is phenomenally similar to being in an anxious mood, and experiencing the emotion of elation is phenomenally similar to being in an elated mood.¹³⁴ And given that we are here concerned with emotions and moods insofar as

¹³⁴ I adapt this point from Kind (2014, p. 116), who points out that, "[f]or every mood state, there seems to be a corresponding emotion state, often identified by the very same term. Compare being in an anxious mood with one's anxiety about a particularly steep segment of a hike, or being in a fearful mood with one's fear when a rattlesnake suddenly appears on the trail ahead, or being in an elated mood with one's elation upon reaching the mountain summit."

they are consciously experienced, or have phenomenality, their phenomenal similarity warrants us in treating them similarly. Second, insofar as emotions and moods *are* distinguished, it tends to be either in terms of moods being of lengthier duration, or in terms of emotions *being about* particular things in a way that moods aren't.¹³⁵ Given that my account already accommodates the temporal form of experience, and given that its *raison d'être* is to explain what phenomenal conditions make the difference between an experience which is *about* something and an experience which is not, my account looks perfectly equipped to accommodate a distinction between emotions and moods. As such, we can presently set aside the matter of what, if anything, distinguishes emotions from moods, and focus solely on how my account accommodates emotional and mood experiences, whether distinct or not.

So, is the structural organisation account plausible for phenomenally intentional emotional and mood experiences? It is generally taken to be the case that emotions and moods can be intentional either by being about our own bodily states, about things in the external world, or about a mixture of both.¹³⁶ Therefore, so long as the structural organisation account can explain how emotional and mood experiences can be about both bodily states and the external world, just in virtue of their phenomenal character, it should have no problem accommodating

¹³⁵ Kind (2014, p. 117) identifies these as the two suggestions prevalent in both the psychological and philosophical literature, for a criterion which distinguishes emotions from moods.

¹³⁶ This is how Kind (2014, pp. 118-24) taxonomises the debate.

phenomenally intentional emotional and mood experiences. I maintain that the account can indeed provide the requisite explanations.

With regard to explaining how emotional and mood experiences can be about our own bodily states, we can provide essentially the same account we just did for pain and bodily sensations. The structural organisation account can explain any location-oriented intentionality of such bodily states in terms of emotion-and-mood mental paints bearing experientially manifest spatial relations to other bodily sensation mental paints or absences of them, within a sensory field mapping out one's body, such as that involved in tactile experience. Beyond this, any further phenomenal intentionality—i.e., the experience being about a particular *sort* of bodily state—can be explained in terms of the particular emotion-and-mood mental paints instantiated, and their spatiotemporal structural organisation.

With regard to explaining how emotional and mood experiences can be about the external world, it makes sense that this is accounted for in precisely the way we accounted for the phenomenal intentionality of perceptual experience, but with emotion-and-mood mental paints incorporated into the relevant structural organisations of mental paints. Cases in which an emotional or mood experience is about something in particular—such as feeling elated about winning a ping-pong tournament—can be explained in terms of those mental paints distinctive of elation being implicated specifically in those structural organisations of mental paints which constitute the *impressions* of the event of winning the ping-pong tournament, and the *taking* of this as having occurred. And cases in which an emotional or mood experience is diffuse and general—such as an all-round anxiety—can be explained in

terms of those mental paints distinctive of anxiety being experientially manifest at all times in the background, as various *impressions* and episodes of *taking* come and go.

One might deny that there are distinctive mental paints involved in emotional or mood experiences, or deny that the instantiation of such distinctive mental paints is *all there is* to an emotional or mood experience. Such denials are unproblematic, for we can, if need be, simply adjust the account to accommodate them. We can allow that emotional or mood experiences involve some sort of overall change in the structure of one's experiences—not one which prevents the structural organisations of mental paint constitutive of *impressions* and *taking* from occurring, but one which effectively 'warps' the way in which one's whole experience is structurally organised.

In sum, the structural organisation account I have fleshed out for the phenomenal intentionality of perceptual experience is perfectly capable of extending to accommodate phenomenally intentional bodily, emotional and mood experiences.

8.2.2 Cognitive Experiences

In fleshing out the structural organisation account in relation to perceptual experiences, I characterised mental paints exclusively in terms of the phenomenal qualities associated with the five paradigmatic sensory modalities of sight, hearing, touch, taste and smell. And in extending the account to bodily, emotional and mood experiences, I have allowed that the mental paint palette may in fact include a far broader array of mental paints, including paints distinctive of bodily sensations such as pain, and paints distinctive of emotions or moods such as nervousness, anxiety, elation and sadness. This expanded palette of mental paints exhausts what is often

called “sensory phenomenology”, and it is ostensibly plausible that perceptual experiences and perception-like experiences—such as hallucinations, illusory experiences, dreams, daydreams and imaginings—are exhaustively constituted by sensory phenomenology.¹³⁷

However, some take the view that there are *cognitive experiences*—phenomenally conscious experiences of *having thoughts* or *thinking*¹³⁸—which involve their own kind of phenomenology distinct from sensory phenomenology, namely ‘cognitive phenomenology’.¹³⁹ If this is the case, and my account cannot accommodate cognitive phenomenology alongside sensory phenomenology, then my account of the constitution of phenomenal intentionality will not extend across the board, to include cognitive experience as well as perceptual and perception-like experience. Of course, this would not be a problem if thoughts were non-intentional.

¹³⁷ The ostensible plausibility of this view has resulted in many adopting what Bayne and Montague (2011, p. 3) call “the ‘conservative’ conception of conscious thought”—the view that sensory phenomenology—construed broadly, to also include bodily sensations, moods and emotions—is the only kind of phenomenology there is.

¹³⁸ Eric Lormand (1996) argues that thoughts can be *conscious* without being *phenomenally conscious*. Unlike Lormand, as stated in Chapter 1, I follow Montague (2016, p. 13) and Strawson (2008, p. 281), in that the only notion of ‘consciousness’ which I am concerned with is that of *phenomenally conscious experience*. As such, Lormand’s non-phenomenal conscious thoughts simply do not qualify as aspects of conscious experience, as I conceive of it. See Pitt (2004, pp. 3-4) for a cogent rebuttal of Lormand’s view.

¹³⁹ Horgan and Tienson (2002) and David Pitt (2004) uphold this view, which Bayne and Montague (2011, p. 3) call “the ‘liberal’ conception of conscious thought”. Also, as we saw in Chapter 4 (§4.4), Strawson (2008) claims that there is ‘cognitive experience’ distinct from ‘sensory-affective experience’, and maintains that it is this *cognitive* dimension of phenomenal character which provides the phenomenal conditions required for phenomenal intentionality, in the form of what he variously refers to as ‘taking’, ‘cognitive intent’ and ‘cognitive-experiential character’.

But far from this being the case, thoughts are a paradigm example of a sort of mental state that is *essentially* intentional, in the sense that a mental state can't be a thought without being a thought *about* something, or a thought *that* something is the case. As such, it is incumbent upon me to show that my account can accommodate phenomenally intentional cognitive experience, or the conscious experience of intentional thought. It is beyond my present remit to engage fully with the intricacies of debates concerning cognitive experience, so I will instead consider the three main diverging positions, amongst those who uphold the existence of cognitive experience, and provide a response to each of them.

The three main positions diverge from one another across two orthogonal lines of disagreement. The first disagreement is over whether thought has a *proprietary* phenomenal character. As Bourget and Mendelovici (2017, §5) explain, “[t]hought has a proprietary phenomenal character just in case the phenomenal characters of thoughts are special or unique to thought,” or, in other words, just in case thoughts have a special cognitive phenomenology distinct from sensory phenomenology. The second disagreement is over whether thoughts have *individuating* phenomenal characters—that is, over whether it is the case that “...thoughts with different intentional contents have different phenomenal characters, and thoughts with different phenomenal characters have different intentional contents” (Ibid.).

Given that these two lines of disagreement are orthogonal, they carve up four different positions. However, one of these is the position that the phenomenal character of thoughts is non-proprietary and non-individuating, which is tantamount

to denying that 'cognitive experience' demarcates any distinctive kind of phenomenally intentional conscious experience. As such, this position simply absolves me of any obligation to extend my account to cognitive experience. The remaining three positions do not make the same denial, and so do not absolve me of this obligation. I will now address each of the remaining three positions in turn.

First, let us take the view that there is a non-individuative, proprietary phenomenal character to thought. This view is entirely compatible with my account, as a non-individuative, proprietary cognitive phenomenal character can be accommodated in one of three ways. Because the phenomenal character is supposed to be proprietary, but not individuative, we only need to account for the instantiation of some phenomenal character special to thought in general, and need not worry about that character being different for different thoughts. As such, one way of accommodating such phenomenal character is simply to expand our mental paint palette further, to include mental paints distinctive of undergoing thought. If one does not find it plausible that there are such cognitive mental paints, existing alongside sensory mental paints such as *blueness* and *sweetness*, then one could instead accommodate proprietary, non-individuative phenomenal character in terms of some overall change in the structure of one's experiences, of the sort suggested as a way of accommodating emotions and moods—a change which effectively 'warps' or 'shapes' the overall way in which one's whole experience is structurally organised, but without preventing the structural organisations of mental paint constitutive of *impressions* and *taking* from occurring. Although this sort of overall structural change would not differ between thoughts about different things, and

hence would be non-individuative, it would be distinctive of undergoing thought in general, and hence proprietary. The third way of accommodating the relevant phenomenal character would be to combine the two methods just described.

Second, let us consider the view that thought has a phenomenal character which is individuative, but non-proprietary. This is perhaps the easiest sort of 'cognitive experience' for my account to accommodate, for it does not require my account to posit any additional phenomenal elements special to cognition. Rather, we can maintain that conscious thoughts are individuated in precisely the same way that conscious perceptions are individuated—namely, by virtue of the structural organisation of mental paints which constitute the experience, comprising *impressions* and *taking*. Different structural organisations of mental paints, comprising different *impressions* and different *takings*, make experiences *about* different things, and thereby individuate them.

Third and finally, there is the view that the phenomenal character of thought is individuative *and* proprietary. Whether my account can accommodate such phenomenal character depends upon precisely what one requires of phenomenal character which is both individuative and proprietary. On the one hand, one may merely require that, when one undergoes conscious thought, there is some aspect of the phenomenal character which individuates the particular thought, and some potentially independent aspect of the same phenomenal character which is special to thought in general. If this is the case, then we can employ a variation on the sort of account proposed in response to the first view, wherein there is some change in the overall structure of one's experience which effectively 'warps' or 'shapes' the

overall experience, but without preventing the structural organisations of mental paint constitutive of *impressions* and *taking* from occurring. We can say that the aspect of the phenomenal character special to thought is the overall 'shaping' or 'warping' of the experience which is distinctive to undergoing thought, and that the aspect which individuates thoughts is precisely the same as that which individuates perceptual experiences, namely the structural organisations of mental paints comprising *impressions* and *taking*, which make experiences *about* different things.

On the other hand, one may require that, when one undergoes conscious thought, there is a *single* aspect of the phenomenal character which is both special to thought *and* which individuates the particular thought. In this case, we could still maintain that conscious thoughts are individuated by different structural organisations of mental paints comprising different impressions and takings, but we would need to stipulate that, in the case of thought, these impressions and takings are comprised of an entirely distinct palette of special cognitive mental paints.

If one wishes to uphold one of these three views in relation to cognitive experience, but is not satisfied with my suggestions as to how they can be upheld alongside my account of the constitution of phenomenal intentionality, then at present I have no more to say. Though I hope, at the least, to have illustrated that the structural organisation account has a number of feasible options at its disposal, such that a belief in the existence of cognitive experience does not immediately render the account untenable.

8.3 Conclusion

In order to flesh out the structural organisation account of the constitution of phenomenal intentionality, I have broken phenomenal intentionality down into two aspects. In the previous chapter, I accounted for one aspect—*impressions*—in terms of the structural arrangement of mental paints. I began this chapter by accounting for the other aspect—*taking*—in terms of the higher-order structural organisation of *impressions* themselves, namely their being temporally sequenced into a systematically related cluster or series. Having articulated these accounts primarily in relation to perceptual experience, I then showed how the complete account of *impressions* and *taking* could feasibly be extended beyond perceptual experience, first focusing on bodily, emotional and mood experiences, and then on cognitive experiences.

9. Conclusion

9.1 Summary

This thesis has been concerned with a particular question concerning the relationship between the phenomenality and intentionality of conscious experience, namely: if the intentionality of conscious experience is constituted by its phenomenality, then how does its phenomenality manage to constitute its intentionality? I have argued that a conscious experience's phenomenality manages to constitute its intentionality by being structurally organised in the appropriate way.

In Chapter 1, we saw why one might be motivated to adopt the phenomenal intentionality view, or PIV—the view that intentionality reduces to phenomenality—as a means of upholding inseparatism alongside non-relationism about intentionality. We recognised that such a position requires some account of what phenomenal conditions are constitutive of phenomenal intentionality. The remainder of this thesis has worked towards such an account.

In Part I, we prepared for our inquiry by considering the mistakes we are prone to, and should avoid, when it comes to theorising about conscious experience. In Chapter 2, I highlighted the importance of distinguishing phenomenology from metaphysics. We saw that, in the debate over the 'transparency' of experience, disputants have overlooked this crucial distinction, and as such have given phenomenological claims undue metaphysical weight. Then, in Chapter 3, we saw that disputants have also conflated two different sorts of 'seeming', and as such have made mistakes about what is 'given' to us in experience. I explained in detail why we

are susceptible to mistakes about what is 'given', so as to help us avoid these mistakes in our own inquiry.

In Part II, we began to consider what phenomenal conditions are constitutive of phenomenal intentionality. In Chapter 4, I assessed six suggestions as to what constitutes the 'phenomenal signature of directedness'. We identified, as a promising candidate, the suggestion that phenomenal intentionality is constituted by the structural organisation of more basic phenomenal elements. Then, In Chapter 5, I assessed this suggestion in light of the lessons of Part I, regarding the mistakes we are liable to make in what we take to be 'given' in conscious experience. We considered how these mistakes might be made specifically in relation to how we take phenomenal intentionality to be 'given', and assessed whether the structural organisation account manages to avoid these mistakes. I argued that the structural organisation account carries minimal risk of error concerning what is 'given' in experience, because it only commits to claims about what is 'given' which are broadly agreed upon, and is neutral regarding more contentious claims.

In Part III, I built upon the core structural organisation account, fleshing it out in greater detail. I articulated an account of how non-intentional elements of phenomenal character—phenomenal qualities—come to constitute phenomenal intentionality when structurally organised in the appropriate way. I argued that phenomenal qualities act as 'mental paints', in the sense that, just like paints on a canvas, although they are not themselves about anything, they can be structurally arranged so as to depict things. After first analysing phenomenal intentionality into the two aspects of *impressions* and *taking*, I then proceeded to argue that both

aspects are constituted by the structural organisation of mental paints. In Chapters 6 and 7, I laid out the account for impressions. In Chapter 6, I elucidated the notion of depiction through consideration of depictive artefacts, and thereby laid the groundwork for treating impressions as phenomenal depictions. At the same time, I identified a positive reason for pursuing my account of phenomenal intentionality: namely, that it is perfectly equipped to explain the dependence of a depiction's *aboutness* upon *experienced structure* in terms of the deeper dependence of phenomenal intentionality upon the structure of the basic, non-intentional phenomenal elements of experience itself. Then, In Chapter 7, I clarified the sense in which impressions are phenomenal depictions, before showing it to be plausible that the phenomenal character of experience has the capacity to instantiate such depictions. I showed that the brain can plausibly support a mental paint palette of the sort required to supply the basic phenomenal elements of all possible experiences, and then explained how, by virtue of the spatiotemporal form of perceptual experience, these phenomenal building blocks are brought to instantiate the depictive structural arrangements which constitute *impressions*. In Chapter 8, I accounted for *taking* in terms of the higher-order structural organisation of *impressions* themselves—namely their being temporally sequenced into a systematically related cluster or series. Having articulated the complete account of *impressions* and *taking* primarily in relation to perceptual experience, I then showed that the account can feasibly be extended across the board, to account for phenomenally intentional experiences beyond perceptual experience.

9.2 Areas for Development and Further Inquiry

In this final section, I would like to briefly highlight two promising areas for further development and inquiry in relation to the structural organisation account of the constitution of phenomenal intentionality.

9.2.1 Incorporating Tracking Relations

Because the structural organisation account treats our ordinary conscious experiences as complex structures made up of various mental paints bearing various different experientially manifest relationships to one-another, there is room for tracking relations to be constructively incorporated into the account. If we take the mental paints involved in veridical perceptual experience to individually track different sorts of properties in the world, then we can thereby explain how the structure of veridical perceptual experience, and hence the non-relational phenomenal intentionality of veridical perceptual experience, corresponds systematically with how things actually are in the world. We can maintain that phenomenal intentionality is non-relationally constituted by structurally organised mental paints, but allow that, during veridical perception, the relevant structural organisation is the result of mental paints piggybacking on tracking relations between the brain and world. This sort of view could enable us to secure the benefits of tracking representationalism, in terms of hooking conscious experience up to the external world during veridical perception, whilst retaining the benefits of a non-relational phenomenal intentionality view for conscious experience in general.

9.2.2 A More Nuanced Origin Story for Intentionality in General

In Chapter 1, I mentioned what Kriegel (2013, p. 1) calls “the *Naturalist-Externalist Research Program*, or NERP.” Although I have, in this thesis, been concerned exclusively with the intentionality of conscious experiences, NERP is a research program concerning intentionality *in general*. And as Kriegel (p. 2) explains, “the basic idea shared by all theories within NERP is that intentionality is injected into the world with the appearance in nature of a certain kind of tracking relation.” Another research program, equally concerned with intentionality *in general*, which is much more in-line with the thrust of this thesis, is what Kriegel (p. 1) calls “the *Phenomenal Intentionality Research Program*, or PIRP.” Within PIRP, “[t]he cornerstone would be the idea that intentionality is injected into the world with the appearance of a certain kind of phenomenal character” (p. 3).

If we incorporate tracking relations into the structural organisation account of the constitution of phenomenal intentionality, as described in the previous subsection, then this paves the way for an account of the origins of intentionality in general which integrates the core ideas of PIRP and NERP, resulting in a more nuanced origin story of intentionality. By holding that there is a form of intentionality, in conscious experience, which only appears when phenomenal character is structurally organised in the appropriate way, we do justice to the guiding intuition of PIRP, that conscious experience brings something of crucial importance into the world, when it comes to intentionality. At the same time, however, we recognise the crucial importance of tracking relations in facilitating this contribution from conscious experience, for it is only in virtue of the tracking relations obtaining that conscious

experience is structurally organised in the appropriate way. If such a synthesis of PIRP and NERP can be made to work, it could provide us with a much more satisfying account of how intentionality has emerged from a complex interplay between the mind and the world.

Bibliography

- Bayne, T. (2009). Perception and the Reach of Phenomenal Content. *Philosophical Quarterly*, 59(236), 385-404.
- Bayne, T., & Montague, M. (2011). Cognitive Phenomenology: An Introduction. In T. Bayne & M. Montague (Eds.) *Cognitive Phenomenology* (pp. 2-35). Oxford: Oxford University Press.
- Block, N. (1990). Inverted Earth. *Philosophical Perspectives*, 4, 53-79.
- Block, N. (1995). On a Confusion about a Function of Consciousness. *Behavioural and Brain Sciences*, 18(2), 227-247.
- Block, N. (1996). Mental Paint and Mental Latex. *Philosophical Issues*, 7, 19-49.
- Block, N. (2003). Mental Paint. In M. Hahn & B. T. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 165-200). Cambridge, Massachusetts: MIT Press.
- Bordini, D. (2017). Is There Introspective Evidence for Phenomenal Intentionality? *Philosophical Studies*, 174(5), 1105-1126.
- Bourget, D., & Mendelovici, A. (2014). Tracking Representationalism: William Lycan, Fred Dretske and Michael Tye. In A. Bailey (Ed.), *Philosophy of Mind: The Key Thinkers* (pp. 209-235). London; New York: Bloomsbury.
- Bourget, D., & Mendelovici, A. (2017). Phenomenal Intentionality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017 ed.). Retrieved 21st October, 2018, from <https://plato.stanford.edu/entries/phenomenal-intentionality/>
- Canmore. (n.d.). *Edinburgh, the Mound, City Relief Map*. Retrieved 26th October, 2018, from <http://canmore.org.uk/site/306311>
- Cassetteboy. (2014, October 1) *Cassetteboy - Cameron's Conference Rap*. [Video file]. Retrieved 5th July, 2018, from <https://youtu.be/OYBumQHPAeU>
- Chalmers, D. J. (2004). The Representational Character of Experience. In B. Leiter (Ed.), *The Future for Philosophy* (pp. 153-181). Oxford: Clarendon Press.
- Coates, P., & Coleman, S. (2015). Introduction: The Nature of Phenomenal Qualities. In P. Coates & S. Coleman (Eds.), *Phenomenal Qualities: Sense, Perception, and Consciousness* (pp. 1-32). New York, NY: Oxford University Press.

- Cohen, W. (1957). Spatial and Textural Characteristics of the Ganzfeld. *The American Journal of Psychology*, 70(3), 403-410.
- Crane, T. (1998). Intentionality as the Mark of the Mental. *Philosophy* (Supplement 43), 229-251.
- Crane, T. (2000). Introspection, Intentionality, and the Transparency of Experience. *Philosophical Topics*, 28(2), 49-67.
- Culliney, K. (2013). *Is that toast I smell? Federation of Bakers grabs attention with perfume*. Retrieved 3rd July, 2018, from <https://www.bakeryandsnacks.com/Article/2013/10/04/Federation-of-Bakers-Eau-de-Toast-perfume-to-promote-bread>
- Cullison, A. (2010). What Are Seemings? *Ratio*, 23(3), 260-274.
- Dainton, B. (2017). Supplement to 'Temporal Consciousness': The Specious Present: Further Issues. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2017 ed.). Retrieved 21st October 2018 from <https://plato.stanford.edu/entries/consciousness-temporal/specious-present.html>
- Dennett, D. C. (1969). *Content and Consciousness*. London: Routledge and Kegan Paul.
- Dennett, D. C. (1988). Quining Qualia. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in Contemporary Science* (pp. 42-78). Oxford: Clarendon Press.
- Dennett, D. C., & Haugeland, J. C. (2004). Intentionality. In R. L. Gregory (Ed.) *The Oxford Companion to the Mind*. Oxford: Oxford University Press.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, Massachusetts: MIT Press.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge, Massachusetts: MIT Press.
- Farkas, K. (2013). Constructing a World for the Senses. In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 99-115). New York, NY: Oxford University Press.
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, Massachusetts: MIT Press.

- Frankish, K. (2016). Illusionism as a Theory of Consciousness. *Journal of Consciousness Studies*, 23(11-12), 11-39.
- Frey, C. (2013). Phenomenal Presence. In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 71-92). New York, NY: Oxford University Press.
- Gen.u.ine.ness. (2016, March 19). *the Fat Duck, Bray – Magical Heston*. [Blog post] Retrieved 3rd July, 2018, from <https://genuiness.wordpress.com/2016/03/19/the-fat-duck-bray-magical-heston/>
- Gow, L. (2016). The Limitations of Perceptual Transparency. *The Philosophical Quarterly*, 66(265), 723-744.
- Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives*, 4, 31-52.
- Hatfield, G. (2007). The Reality of Qualia. *Erkenntnis*, 66(1), 133-168.
- Helton, G. (2016). Recent Issues in High-Level Perception. *Philosophy Compass*, 11(12), 851-862.
- Hochberg, J. (1984). The Perception of Pictorial Representations. *Social Research*, 51(4), 841-841.
- Hoerl, C. (2009). Time and Tense in Perceptual Experience. *Philosophers' Imprint*, 9(12), 1-18.
- Hoerl, C. (2018). Experience and Time: Transparency and Presence. *Ergo*, 5(5), 127-151.
- Hopkins, R. (1995). Explaining Depiction. *The Philosophical Review*, 104(3), 425-455.
- Horgan, T. (2012). Introspection about Phenomenal Consciousness: Running the Gamut from Infallibility to Impotence. In D. Smithies & D. Stoljar (Eds.), *Introspection and Consciousness* (pp. 405-421). New York ; Oxford: Oxford University Press.
- Horgan, T., & Tienson, J. (2002). The Intentionality of Phenomenology and the Phenomenology of Intentionality. In D. J. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings* (pp. 520-533). New York: Oxford University Press.

- Horgan, T., Tienson, J., & Graham, G. (2004). Phenomenal Intentionality and the Brain in a Vat. In R. Schantz (Ed.), *The Externalist Challenge* (pp. 297-318). Berlin; New York: Walter de Gruyter.
- Jackendoff, R. (1987). *Consciousness and the Computational Mind*. Cambridge, Massachusetts; London: MIT Press in cooperation with the British Psychological Society.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.
- Jacob, P. (2014). Intentionality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2014 ed.). Retrieved 21st October, 2018, from <https://plato.stanford.edu/entries/intentionality/>
- James, W. (1890). *The Principles of Psychology* (Vol. 1). London: Macmillan.
- Johnston, M. (2007). Objective Mind and the Objectivity of Our Minds. *Philosophy and Phenomenological Research*, 75(2), 233-268.
- JTD. (2013, August 6). *L'ARTISAN PARFUMEUR TRAVERSÉE DU BOSPHORE, 2010*. Retrieved 3rd July, 2018, from <http://scenthurdle.com/lartisan-parfumeur-traversee-du-bosphore-2010/>
- Kennedy, M. (2009). Heirs of Nothing: The Implications of Transparency. *Philosophy and Phenomenological Research*, 79(3), 574-604.
- Kind, A. (2001). Qualia Realism. *Philosophical Studies*, 104(2), 143-162.
- Kind, A. (2003). What's so Transparent about Transparency? *Philosophical Studies*, 115(3), 225-244.
- Kind, A. (2008). How to Believe in Qualia. In E. L. Wright (Ed.), *The Case for Qualia* (pp. 285-298). Cambridge, Massachusetts: The MIT Press.
- Kind, A. (2010). Transparency and Representationalist Theories of Consciousness. *Philosophy Compass*, 5(10), 902-913.
- Kind, A. (2014). The Case Against Representationalism about Moods. In U. Kriegel (Ed.), *Current Controversies in Philosophy of Mind* (pp. 113-134). New York: Routledge.
- Kriegel, U. (2007). The Phenomenologically Manifest. *Phenomenology and the Cognitive Sciences*, 6(1), 115-136.

- Kriegel, U. (2011). *The Sources of Intentionality*. New York; Oxford : Oxford University Press.
- Kriegel, U. (2013). The Phenomenal Intentionality Research Program. In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 1-26). New York, NY: Oxford University Press.
- Kriegel, U. (2014). Two Notions of Mental Representation. In U. Kriegel (Ed.), *Current Controversies in Philosophy of Mind* (pp. 161-179). New York: Routledge.
- Kripke, S. A. (1980). *Naming and Necessity*. Oxford: Blackwell.
- Lo, V. (2014, October 23). *Marrakech Intense*. Retrieved 3rd July, 2018, from <http://www.asiaspa.com/3143/20141023123451/>
- Loar, B. (2017). Phenomenal Intentionality as the Basis of Mental Content. In K. Balog & S. Beardman (Eds.), *Consciousness and Meaning: Selected Essays* (pp. 291-318). Oxford: Oxford University Press. Originally published (2003). In M. Hahn & B. T. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 229-258). Cambridge, Massachusetts; London: MIT Press.
- Logue, H. (2013). Visual Experience of Natural Kind Properties: Is There Any Fact of the Matter? *Philosophical Studies*, 162(1), 1-12.
- Lormand, E. (1996). Nonphenomenal Consciousness. *Noûs*, 30(2), 242-261.
- Lycan, W. G. (1996). *Consciousness and Experience*. Cambridge, Massachusetts; London: MIT Press.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W.H. Freeman.
- Martin, M. G. F. (1992). Sight and Touch. In T. Crane (Ed.), *The Contents of Experience: Essays on Perception* (pp. 196-215). Cambridge: Cambridge University Press.
- Martin, M. G. F. (1998). Setting Things Before the Mind. *Royal Institute of Philosophy Supplement*, 43, 157.
- Martin, M. G. F. (2002). The Transparency of Experience. *Mind & Language*, 17(4), 376-425.

- Masrour, F. (2013). Phenomenal Objectivity and Phenomenal Intentionality: In Defense of a Kantian Account. In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 116-136). New York, NY: Oxford University Press.
- McGinn, C. (1988). Consciousness and Content. *Proceedings of the British Academy*, 74, 225-245.
- Mendelovici, A. (2014). Pure Intentionalism about Moods and Emotions. In U. Kriegel (Ed.), *Current Controversies in Philosophy of Mind* (pp. 135-157). New York: Routledge.
- Mendelovici, A. (2018). *The Phenomenal Basis of Intentionality*. New York, NY: Oxford University Press.
- Mendelovici, A., & Bourget, D. (2014). Naturalizing Intentionality: Tracking Theories Versus Phenomenal Intentionality Theories. *Philosophy Compass*, 9(5), 325-337.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, Massachusetts; London: MIT Press.
- Molyneux, B. (2009). Why Experience Told Me Nothing about Transparency. *Noûs*, 43(1), 116-136.
- Montague, M. (2010). Recent Work on Intentionality. *Analysis*, 70(4), 765-782.
- Montague, M. (2016). *The Given: Experience and its Content*. Oxford: Oxford University Press.
- Moore, G. (1903). The Refutation of Idealism. *Mind*, 12, 433-453.
- Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review*, 83(4), 435-450.
- Nida-Rümelin, M. (2007). Transparency of Experience and the Perceptual Model of Phenomenal Awareness. *Philosophical Perspectives*, 21, 429-455.
- Nida-Rümelin, M. (2008). Phenomenal Character and the Transparency of Experience. In E. L. Wright (Ed.), *The Case for Qualia* (pp. 309-324). Cambridge, Massachusetts: MIT Press.
- Nida-Rümelin, M. (2016). The illusion of illusionism. *Journal of Consciousness Studies*, 23(11-12), 160-171.
- Nudds, M. (2009). *Sounds and Space*. Oxford: Oxford University Press.

- O'Shaughnessy, B. (1985) Seeing the Light. *Proceedings of the Aristotelian Society*, 85, 193-218.
- Papineau, D. (2014). Sensory Experience and Representational Properties. *Proceedings of the Aristotelian Society*, 114(1pt1), 1-33.
- Pautz, A. (2008). The Interdependence of Phenomenology and Intentionality. *The Monist*, 91(2), 250-272.
- Pautz, A. (2013). Does Phenomenology Ground Mental Content? In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 194-234). New York, NY: Oxford University Press.
- Peacocke, C. (1987). Depiction. *The Philosophical Review*, 96(3), 383.
- Pitt, D. (2004). The Phenomenology of Cognition Or What Is It Like to Think That P? *Philosophy and Phenomenological Research*, 69(1), 1-36.
- Price, R. (2009). Aspect-Switching and Visual Phenomenal Character. *The Philosophical Quarterly*, 59(236), 508-518.
- Prinz, J. J. (2012). *The Conscious Brain: How Attention Engenders Experience*. New York, NY: Oxford University Press.
- Reid, T. (1785). *Essays on the Intellectual Powers of Man*. Edinburgh: Printed for John Bell, Parliament Square, and G. G. J. & J. Robinson, London.
- Reid, T. (1850). *Essays on the Intellectual Powers of Man*. J. Walker (Abr. Ed.). Cambridge: John Bartlett.
- Reid, T. (1970) *An Inquiry into the Human Mind*. T. Duggan (Ed.). Chicago: University of Chicago Press. (Originally published in 1764).
- Rey, G. (1998). A Narrow Representationalist Account of Qualitative Experience. *Philosophical Perspectives*, 32(12), 435-457.
- Rosenthal, D. (1991). The Independence of Consciousness and Sensory Quality. *Philosophical Issues*, 1, 15-36.
- Rosenthal, D. (1994). Identity Theories. In S. Guttenplan (Ed.), *A Companion to the Philosophy of Mind* (pp. 348-355). Oxford: Blackwell.

- Rosenthal, D. (2009). Higher-Order Theories of Consciousness. In A. Beckermann, B. P. McLaughlin & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 239-253). Oxford: Clarendon Press.
- Rosenthal, D. (2010). How to Think about Mental Qualities. *Philosophical Issues*, 20, 368-393.
- Rosenthal, D. (2015). Quality Spaces and Sensory Modalities. In P. Coates & S. Coleman (Eds.), *Phenomenal Qualities: Sense, Perception, and Consciousness* (pp. 33-65). New York, NY: Oxford University Press.
- Searle, J. (1994). Intentionality. In S. Guttenplan (Ed.), *A Companion to the Philosophy of Mind* (pp. 379-386). Oxford: Blackwell.
- Searle, J. (1995). The Mystery of Consciousness: Part II. *The New York Review of Books*, 42(18), 54.
- Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge, Massachusetts; London: MIT Press.
- Sedivy, S. (1996). Conventional Naturalism: A Perceptualist Account of Pictorial Representation. *International Studies in the Philosophy of Science*, 10(2), 103-125.
- Sellars, W. (1963). Philosophy and the Scientific Image of Man. In *Science, Perception and Reality* (pp. 1-40). London: Routledge & Kegan Paul.
- Siegel, S. (2006). Which Properties are Represented in Perception? In T.S. Gendler & J. Hawthorne (Eds.), *Perceptual Experience* (pp. 481-503). New York, NY: Oxford University Press.
- Siegel, S. (2016). The Contents of Perception. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 ed.). Retrieved 24th October, 2018, from <https://plato.stanford.edu/archives/win2016/entries/perception-contents/>
- Siewert, C. (2004). Is Experience Transparent? *Philosophical Studies*, 117(1), 15-41.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton, N.J.; Chichester: Princeton University Press.
- Simakov, O., Takeshi, K., Ferdinand, M., Jerry, J., Ryo, K., Therese, M., et al. (2015). Hemichordate Genomes and Deuterostome Origins. *Nature*, 527(7579), 459-465.

- Sims, F. (2016, January 27). *Heston Blumenthal on the new experience at the reopened Fat Duck*. Retrieved 3rd July, 2018, from <https://www.thecaterer.com/articles/364705/heston-blumenthal-on-the-new-experience-at-the-reopened-fat-duck>
- Smith, D. W. (2018). Phenomenology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018 ed.). Retrieved 21st October, 2018, from <https://plato.stanford.edu/entries/phenomenology/>
- Smithies, D. (2013). The Nature of Cognitive Phenomenology. *Philosophy Compass*, 8(8), 744-754.
- Stampe, D. W. (1977). Toward a Causal Theory of Linguistic Representation. *Midwest Studies In Philosophy*, 2(1), 42-63.
- Stoljar, D. (2004). The Argument from Diaphanousness. *Canadian Journal of Philosophy*, 30, 341-390.
- Strawson, G. (1994). *Mental Reality*. Cambridge, Massachusetts; London: MIT Press.
- Strawson, G. (2008). *Real Materialism and Other Essays*. Oxford: Clarendon Press.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Massachusetts; London: MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, Massachusetts; London: MIT Press.
- Tye, M. (2014). Transparency, Qualia Realism and Representationalism. *Philosophical Studies*, 170(1), 39-57.
- Whiting, D. (2016). On the Appearance and Reality of Mind. *The Journal of Mind and Behavior*, 37(1), 47-70.
- Wikipedia Contributors. (2018). Madame Tussauds. In *Wikipedia, The Free Encyclopedia*. Retrieved 5th July, 2018, from https://en.wikipedia.org/w/index.php?title=Madame_Tussauds&oldid=863533137
- Wittgenstein, L. (1961). *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul.
- Young, B. D., Keller, A., & Rosenthal, D. (2014). Quality-Space Theory in Olfaction. *Frontiers in Psychology*, 5, 1-15.

Zahavi, D. (2003). Intentionality and Phenomenality: A Phenomenological Take on the Hard Problem. *Canadian Journal of Philosophy*, 29, 63-92.